

From Independence to Expansion and Back Again

Tobias Christiani
IT University of Copenhagen,
Denmark
tobc@itu.dk

Rasmus Pagh
IT University of Copenhagen,
Denmark
pagh@itu.dk

Mikkel Thorup
University of Copenhagen,
Denmark
mikkel2thorup@gmail.com

ABSTRACT

We consider the following fundamental problems:

- Constructing k -independent hash functions with a space-time tradeoff close to Siegel's lower bound.
- Constructing representations of unbalanced expander graphs having small size and allowing fast computation of the neighbor function.

It is not hard to show that these problems are intimately connected in the sense that a good solution to one of them leads to a good solution to the other one. In this paper we exploit this connection to present efficient, recursive constructions of k -independent hash functions (and hence expanders with a small representation). While the previously most efficient construction (Thorup, FOCS 2013) needed time quasipolynomial in Siegel's lower bound, our time bound is just a logarithmic factor from the lower bound.

Categories and Subject Descriptors

E.1 [Data]: Data Structures; E.2 [Data]: Data Storage Representations—*Hash-table representations*; G.2.2 [Discrete Mathematics]: Graph Theory; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems; G.3 [Mathematics of Computing]: Probability and Statistics—*Probabilistic algorithms, Random number generation*

General Terms

Algorithms, Performance, Theory

Keywords

bipartite expanders; k -independence; hash functions

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
STOC'15, June 14–17, 2015, Portland, Oregon, USA.
Copyright © 2015 ACM 978-1-4503-3536-2/15/06 ...\$15.00.
DOI: <http://dx.doi.org/10.1145/2746539.2746620>.

1. INTRODUCTION

'Not all those who wander are lost.' — Bilbo Baggins.

The problem of designing explicit unbalanced expander graphs with near-optimal parameters is of major importance in theoretical computer science. In this paper we consider bipartite graphs with edge set $E \subset U \times V$ where $|U| \gg |V|$. Vertices in U have degree d and expansion is desired for subsets $S \subset U$ with $|S| \leq k$ for some parameter k . Such expanders have numerous applications (e.g. hashing [22], routing [1], sparse recovery [13], membership [2]), yet coming up with explicit constructions that have close to optimal parameters has proved elusive. At the same time it is easy to show that choosing E at random will give a graph with essentially optimal parameters. This means that we can efficiently and with a low probability of error produce a description of an optimal unbalanced expander that takes space proportional to $|U|$. Storing a complete description is excessive for most applications that, provided access to an explicit construction, would use space proportional to $|V|$. On the other hand, explicit constructions can be represented using constant space, but the current best explicit constructions have parameters d and $|V|$ that are polynomial in the optimal parameters of the probabilistic constructions [11]. Furthermore, existing explicit constructions have primarily aimed at optimizing the parameters of the expander, with the evaluation time of the neighbor function being of secondary interest, as long as it can be bounded by $\text{poly} \log u$. This evaluation time is excessive in applications that, provided access to the neighbor function of an optimal expander, would use time proportional to d , where d is typically constant or at most logarithmic in $|U|$.

In this paper we focus on optimizing the parameters of the expander while minimizing the space usage of the representation and the evaluation time of the neighbor function. We present randomized constructions of unbalanced expanders in the standard word RAM model. Our constructions have near-optimal parameters, use space close to $|V|$, and support computing the d neighbors of a vertex in time close to d .

Hash functions and expander graphs.

There is a close connection between k -independent hash functions and expanders. A k -independent function with appropriate parameters will, with some probability of failure, represent the neighbor function of a graph that expands on subsets of size k . This is what we refer to as going from independence to expansion, and the fact follows from the standard union bound analysis of probabilistic constructions of expanders. Going in the other direction, from expansion

to independence, was first used by Siegel [22] as a technique for showing the existence of k -independent hash functions with evaluation time that does not depend on k . We follow in Siegel’s footsteps and a long line of work (see e.g. [9] for an overview) that focuses on the space-time tradeoff of k -independent hash functions over a universe of size $u = |U|$.

Ideally, we would like to construct a data structure in the word RAM model that takes as input parameters u , k , and t , and returns a k -independent hash function over U . The hash function should use space $k(u/k)^{1/t}$ and have evaluation time $O(t)$, matching up to constant factors the space-time tradeoff of Siegel’s cell probe lower bound for k -independent hashing [22]. We present the first construction that comes close to matching the space-time tradeoff of the cell probe lower bound when considering both evaluation time and the space usage in k .

Method.

Our work is inspired by Siegel’s graph powering approach [22] and by recent advances in tabulation hashing [24], showing that it is possible to efficiently describe expanders in space much smaller than u . Our main insight is that it is possible to make simple, recursive expander constructions by alternating between strong unbalanced expanders and highly random hash functions. Similarly to previous work, we follow the procedure of letting a k -independent function represent a bipartite graph Γ that expands on subsets of size k . We then apply a graph product to Γ in order to increase the size of the universe covered by the graph while retaining expander properties. At each step of the recursion we return to k -independence by combining the graph product with a table of random bits, leaving us with a new k -independent function that covers a larger universe. By combining the technique of alternating between expansion and independence with a new and more efficient graph product, we can improve upon existing randomized constructions of unbalanced expanders.

1.1 Our contribution

Table 1 compares previous upper and lower bounds on k -independent hashing with our results, as presented in Corollaries 1, 2, and 3. As can be seen, most results present a trade-off between time and space controlled by a parameter t . Tight lower and upper bounds have been known only in the cell probe model, but our new construction nearly matches the cell probe lower bound by Siegel [22].

The time bound for the construction using explicit expanders [11] uses the degree of the expander as a conservative lower bound, based on the possibility that the neighbor function in their construction can be evaluated in constant time in the word RAM model. The time bound that follows directly from their work is $\text{poly log } u$. While the constant factors in the exponent of the space usage of [22, 24] have likely not been optimized, their techniques do not seem to be able to yield space close to the cell probe lower bound.

As can be seen our construction polynomially improves either space or time compared to each of the previously best trade-offs. We also find our construction easier to describe and analyze than the results of [11, 15, 24], with simplicity comparable to that of Siegel’s influential paper [22].

Like all other randomized constructions our data structures comes with an error probability, but this error probability is *universal* in the sense that if the construction works then it provides independent hash values on *every subset* of at

most k elements from U . This is in contrast to other known constructions [10, 19] that give independence with high probability on each *particular* set of at most k elements, but will fail almost surely if independence for a superpolynomial number of subsets is needed.

Applications.

Efficient constructions of highly random functions is of fundamental interest with many applications in computer science. A k -independent function can, without changing the analysis, replace a fully random function in applications that only rely on k -subsets of inputs mapping to random values. We can therefore view k -independent functions as space and randomness efficient alternatives to fully random functions, capable of providing compact representations of complex structures such as expander graphs over very large domains. Apart from the construction of expander graphs with a small description, as an example application, k -independent functions with a universal error probability can be used to construct “real-time” dictionaries that are able to handle extremely long (in expectation) sequences of insertion and deletion operations in constant time per operation before failing.

Let $\tau > 1$ be a constant parameter. We use a k -independent hash function with $k = w^{O(\tau)}$ to split a set of n machine words of w bits into $O(n)$ subsets such that each subset has size at most k , with probability at least $1 - 2^{-w^\tau}$. Handling each subset with Thorup’s recent construction of dictionaries for sets of size $w^{O(\tau)}$ using time $O(\tau)$ per operation [21] we get a dynamic dictionary in which, with high probability, every operation in a sequence of length $\ell < 2^{O(w^\tau)}$ takes constant time. In comparison the hash functions of [8, 10, 19] can only guarantee that sequences of length $\ell < \text{poly}(n)$ operations, where $n < 2^w$, succeed with high probability. The splitting hash function needs space $u^{\Omega(1)}$, which might exceed the space usage of an individual dictionary, but this can be seen as a shared resource that is used for many dictionaries (in which case we bound the total number of operations before failure).

2. BACKGROUND AND OVERVIEW

In the analysis of randomized algorithms we often assume access to a fully random function of the form $f : [u] \rightarrow [r]$ where $[n]$ denotes the set $\{0, 1, \dots, n - 1\}$. To represent such a function we need a table with u entries of $\log r$ bits. This is impractical in applications such as hashing based dictionaries where we typically have that $u \gg r$ and the goal is to use space $O(r)$ to store r elements of $[u]$. Fortunately, the analysis that establishes the performance guarantees of a randomized algorithm can often be modified to work even in the case where the function f has weaker randomness properties.

One such concept of limited randomness is k -independence, first introduced to computer science in the 1970s through the work of Carter and Wegman on universal hashing [4]. A family of functions from $[u]$ to $[r]$ is k -independent if, for every subset of $[u]$ of cardinality at most k , the output of a random function from the family evaluated on the subset is independent and uniformly distributed in $[r]$. Trivially, the family of all functions from $[u]$ to $[r]$ is k -independent, but representing a random function from this family uses too much space. It was shown in [14] that for every finite field \mathbb{F} the family of functions that consist of all polynomials over \mathbb{F}

Table 1: Space-time tradeoffs for k -independent hash functions

Reference	Space	Time
Polynomials [14, 5]	k	$O(k)$
Preprocessed polynomials [15]	$k^{1+\varepsilon}(\log u)^{1+o(1)}$	$(\text{poly } \log k)(\log u)^{1+o(1)}$
Expanders [11] + [22]	$k^{1+\varepsilon}d^2$	$d = O((\log u)(\log k))^{1+1/\varepsilon}$
Expander powering [22]	$k^{(1-\varepsilon)t}u^\varepsilon + u^{1/t}$	$O(1/\varepsilon)^t$
Double tabulation [24]	$k^{5t} + u^{1/t}$	$O(t)$
Recursive tabulation [24]	$\text{poly } k + u^{1/t}$	$O(t^{\log t})$
Corollary 1	$ku^{1/t}t^3$	$O(t^2 + t^3(\log k)/(\log u))$
Corollary 2	$k^2u^{1/t}t^2$	$O(t \log t + t^2(\log k)/(\log u))$
Corollary 3*	$ku^{1/t}t$	$O(t \log t)$
Cell probe lower bound [22]	$k(u/k)^{1/t}$	$t < k$ probes
Cell probe upper bound [22]	$k(u/k)^{1/t}t$	$O(t)$ probes

Table notes: Space-time tradeoffs for k -independent hash functions from a domain of size u , with the trade-off controlled by a parameter t . Time bounds in the last two rows are number of cell probes, and remaining rows refer to the word RAM model with word size $\Theta(\log u)$. Leading constants in the space bounds are omitted. We use t to denote an arbitrary positive integer parameter that controls the trade-off, and We use ε to denote an arbitrary positive constant. *Corollary 3 relies on the assumption $k = u^{O(1/t)}$.

of degree at most $k - 1$ is k -independent. A function from this family can be represented using near-optimal space [6] by storing the k coefficients of the polynomial. The mapping defined by a function f from a k -independent polynomial family over $\mathbb{F} = \{x_1, x_2, \dots, x_u\}$ takes the form

$$\begin{bmatrix} f(x_1) \\ f(x_2) \\ \vdots \\ f(x_u) \end{bmatrix} = \begin{bmatrix} x_1^0 & x_1^1 & \dots & x_1^{k-1} \\ x_2^0 & x_2^1 & \dots & x_2^{k-1} \\ \vdots & \vdots & \vdots & \vdots \\ x_u^0 & x_u^1 & \dots & x_u^{k-1} \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_{k-1} \end{bmatrix}. \quad (1)$$

The k -independence of the polynomial family follows from properties of the Vandermonde matrix: every subset of k rows is linearly independent. The problem with this construction is that the Vandermonde matrix is dense, resulting in an evaluation time of $\Omega(k)$ if we simply store the coefficients of the polynomial. The lower bounds by Siegel [22], and later Larsen [17], as presented in Table 1, show that a data structure for evaluating a polynomial of degree $k - 1$ using time $t < k$ must use space at least $k(u/k)^{1/t}$. The data structure of [15] presents a step in this direction, but is still far from the lower bound for k -independent functions.

The quest for k -independent families of functions with evaluation time $t < k$ can be viewed as attempts to construct compact representations of sparse matrices that fill the same role as the Vandermonde matrix. We are interested in compact representations that support fast computation of the sparse row associated with an element $x \in [u]$. An example of a sparse matrix with these properties is the adjacency matrix of a bipartite expander graph with sufficiently strong expansion properties. For the purposes of constructing k -independent hash functions we are primarily interested in expanders that are highly unbalanced.

Expander hashing.

Prior constructions of fast and highly random hash functions has followed Siegel’s approach of combining expander graphs with tables of random words. If Γ is a k -unique expander graph (see Definition 1) then we can construct a k -independent function by composing it with a simple tabulation function h . This approach would yield optimal k -independent hash functions if we had access to explicit expanders with optimal parameters that could be evaluated in time proportional to the left outdegree. Unfortunately, no explicit construction of a k -unique expander with optimal parameters is known.

Siegel [22] addresses this problem by storing a smaller randomly generated k -unique expander, say, one that covers a universe of size $u^{1/t}$. By the k -independent hashing lower bound, if an expander with $|U| = u^{1/t}$ has degree d , then in order for it to be k -unique it must have a right hand side of size $|V| \geq k(u^{1/t}/k)^{1/d}$. To give a space efficient construction of a k -unique expander that covers a universe of size u , Siegel repeatedly applies the Cartesian product to the graph. Applying the Cartesian product t times to a k -unique expander results in a graph that remains k -unique but with the left degree and size of the left and right vertex sets raised to the power t . Using space $u^{1/t}$ to store an expander with degree t , it follows from the lower bound that the expander resulting from repeatedly applying the Cartesian product must have

$$|V'| \geq (k(u^{1/t}/k)^{1/d})^t = k^{(1-1/d)t}u^{1/d}.$$

Setting $d = 1/\varepsilon$, the randomly generated k -unique expander that forms the basis of the construction has degree $O(1/\varepsilon)$, leading to the expression in Table 1. Since we need to store $|V'|$ random words in a table in order to create a k -independent hash function, Siegel’s graph powering approach

offers a space-time tradeoff that is far from the lower bound from our perspective where both u , k , and t are parameters to the hash function.

Thorup [24] shows that, for the right choice of parameters, a simple tabulation hash function is likely to form a compact representation of a k -unique expander. A simple tabulation function takes a string $x = (x_1, x_2, \dots, x_c)$ of c characters from some input alphabet $\langle n \rangle = \{0, 1\}^n$, and returns a string of d characters from some output alphabet $\langle m \rangle = \{0, 1\}^m$. The simple tabulation function $h : \langle n \rangle^c \rightarrow \langle m \rangle^d$ is evaluated by taking the exclusive-or of c table-lookups

$$h(x) = h_1(x_1) \oplus h_2(x_2) \oplus \dots \oplus h_c(x_c)$$

where $h_i : \langle n \rangle \rightarrow \langle m \rangle^d$ is a random function. The advantage of a simple tabulation function compared to a fully random function is that we only need to store the random character tables h_1, h_2, \dots, h_c . Thorup is able to show that for $d \geq 6c$ a simple tabulation function is k -unique with a low probability of failure when $k \leq (2^m)^{1/5c}$. Setting $n = m$ and composing the k -unique expander resulting from a single application of simple tabulation with another simple tabulation function, Thorup first constructs a hash function with space usage $u^{1/c}$, independence $u^{\Omega(1/c^2)}$, and evaluation time $O(c)$. He then presents a second trade-off with space $u^{1/c}$, independence $u^{\Omega(1/c)}$, and time $O(c^{\log c})$ that comes from applying simple tabulation recursively to the output of a simple tabulation function. Similar to Siegel's upper bound, the space usage of Thorup's upper bounds with respect to k is much larger than the lower bound as can be seen from Table 1 where the space-time tradeoff of his results have been parameterized in terms of the independence k .¹

Explicit constructions.

The literature on explicit constructions has mostly focused on optimizing the parameters of the expander, with the evaluation time of the neighbor function being of secondary interest, as long as it is bounded by $\text{poly } \log u$. As can be seen from Siegel's cell probe lower and upper bounds, optimal constructions of k -independent hash functions have evaluation time in the range $t = 1$ to $t = \log u$. Therefore, an explicit construction, even if we had one with optimal parameters, would without further guarantees on the running time not be enough to solve our problem of constructing efficient expanders. Here we briefly review the construction given by Guruswami et al. [11]. It is, to our knowledge, currently the best explicit construction of unbalanced bipartite expanders in terms of the parameters of the graph. Their construction and its analysis is, similarly to the polynomial hash function in equation (1), algebraic in nature and inspired by techniques from coding theory, in particular Parvaresh-Vardy codes and related list-decoding algorithms [20]. In their construction, a vertex x is identified with its Reed-Solomon message polynomial over a finite field \mathbb{F} . The i th neighbor of x is found by taking a sequence of powers of the message polynomial over an extension field, evaluating each of the resulting polynomials in the i th element of \mathbb{F} , and concatenating the output. In contrast, the constructions presented in this paper only use the subset of standard word RAM

¹It should be noted that Thorup's analysis is not tuned to optimize the polynomial dependency on k , and that he gives stronger concrete parameters for some realistic parameter settings.

instructions that can be implemented in AC^0 . In Table 1 we have assumed that we can evaluate their neighbor function in constant time as a conservative lower bound on the performance of their construction in the word RAM model. Other highly unbalanced explicit constructions given in [3, 23] offer a tradeoff where either one of d or $|V|$ is quasipolynomial in the lower bound. In comparison, the construction by Guruswami et al. is polynomial in both of these parameters.

3. OUR CONSTRUCTIONS

In this section we present three randomized constructions of efficient expanders in the word RAM model. Each construction offers a different tradeoff between space, time, and the probability of failure. We present our constructions as data structures, with the randomness generated by the model during an initialization phase. The initialization time of our data structures is always bounded by their space usage, and to simplify the exposition we therefore only state the latter. Alternatively, our constructions could be viewed directly as randomized algorithms, taking as input a list of parameters, a random seed, and a vertex $x \in [u]$ and returning the list of neighbors of x . The hashing corollaries presented in Table 1 follow directly from our three main theorems using Siegel's expander hashing technique.

3.1 Model of computation

The algorithms presented in this paper are analyzed in the standard word RAM model with word size w as defined by Hagerup [12], modeling what can be implemented in a standard programming language like C [16]. In order to show how our algorithms benefit from word-level parallelism we use w as a parameter in the analysis. To simplify the exposition we impose the natural restriction that, for a given choice of parameters to a data structure, the word size is large enough to address the space used by the data structure. In other words, our results are stated with w as an unrestricted parameter, but are only valid when we actually have random access in constant time.

The data structures we present require access to a source of randomness in order to initialize the character tables of simple tabulation functions. To accommodate this we augment the model with an instruction that uses constant time to generate a uniformly random and independent integer in $[r]$ where $r \leq 2^w$. We note that our constructions use only the subset of arithmetic instructions required for evaluating a simple tabulation function, i.e, standard bit manipulation instructions, integer addition, and subtraction. Our results therefore hold in a version of the word RAM model that only uses instructions that can be implemented in AC^0 , known in the literature as the restricted model [12] or the Practical RAM [18].

3.2 Notation and definitions

Let $\langle n \rangle = \{0, 1\}^n$ denote the alphabet of n -bit strings, and let $x = (x_1, x_2, \dots, x_c) \in \langle n \rangle^c$ denote a string of n -bit characters of length c . We define a concatenation operator \parallel that takes as input two characters $x \in \langle n \rangle$ and $y \in \langle m \rangle$, and concatenates them to form $x \parallel y \in \langle n + m \rangle$. The concatenation operator can also be applied to strings of equal length where it performs component-wise concatenation. Given characters $x \in \langle n \rangle^d$ and $y \in \langle m \rangle^d$ the concatenation $x \parallel y$ is an element of $\langle n + m \rangle^d$ with the i th component of $x \parallel y$ defined by $(x \parallel y)_i = x_i \parallel y_i$.

We will present word RAM data structures that represent functions of the form $\Gamma : \langle n \rangle^c \rightarrow \langle m \rangle^d$. The function Γ defines a d -regular bipartite graph with input set $\langle n \rangle^c$ and output set $\{1, 2, \dots, d\} \times \langle m \rangle$. For $S \subseteq \langle n \rangle^c$ we overload Γ and define $\Gamma(S) = \{(i, \Gamma(x)_i) \mid x \in S\}$, i.e., $\Gamma(S)$ is the set of outputs of S . We are interested in constructing functions where every subset S of inputs of size at most k contains an input that has many unique neighbors, formally:

Definition 1. Let $\Gamma : \langle n \rangle^c \rightarrow \langle m \rangle^d$ be a function satisfying the following property:

$$\forall S \subseteq \langle n \rangle^c, |S| \leq k, \exists x \in S : |\Gamma(\{x\}) \setminus \Gamma(S \setminus \{x\})| > l.$$

Then, for $l > 0$ we say that Γ is k -unique. If further $l > d/2$ we say that Γ is k -majority-unique.

For completeness we define the concept of k -independence:

Definition 2. Let k be a positive integer and let \mathcal{F} be a family of functions from U to R . We say that \mathcal{F} is a k -independent family of functions if, for every choice of $l \leq k$ distinct keys x_1, \dots, x_l and arbitrary values y_1, \dots, y_l , then, for f selected uniformly at random from \mathcal{F} we have that

$$\Pr[f(x_1) = y_1 \wedge f(x_2) = y_2 \wedge \dots \wedge f(x_k) = y_k] = |R|^{-k}.$$

Simple tabulation functions are an important tool in our constructions. Our data structures can be made to consist entirely of simple tabulation functions and our evaluation algorithms can be viewed as a sequence of adaptive calls to this collection of simple tabulation functions.

Definition 3. Let (R, \oplus) denote an abelian group. A *simple tabulation function* $h : \langle n \rangle^c \rightarrow R$ is defined by

$$h(x) = \bigoplus_{i=1}^c h_i(x_i)$$

where each *character table* $h_i : \langle n \rangle \rightarrow R$ is a k -independent function.

In this paper we consider simple tabulation functions with character tables that operate either on bit strings under the exclusive-or operation, $R = (\langle m \rangle, \oplus)$, or on sets of non-negative integers modulo some integer r , $R = ([r], +)$.

3.3 From k -uniqueness to k -independence

In his seminal paper Siegel [22] showed how a k -unique function can be combined with a table of random elements in order to define a k -independent family of functions. In his paper on the expansion properties of tabulation hash functions, Thorup [24, Lemma 2] used a slight variation of Siegel's technique that makes use of the position-sensitive structure of the bipartite graph defined by $\Gamma : \langle n \rangle^c \rightarrow \langle m \rangle^d$. This is the version we state here.

Lemma 1 (Siegel [22], Thorup [24]). *Let $\Gamma : \langle n \rangle^c \rightarrow \langle m \rangle^d$ be k -unique and let $h : \langle m \rangle^d \rightarrow R$ be a simple tabulation function. Then $h \circ \Gamma$ defines a family of k -independent functions. We sample a function from the family by sampling the character tables of h .*

3.4 From k -independence to k -uniqueness

A k -independent function has the same properties as a fully random function when considering k -subsets of inputs. Randomized constructions of k -unique functions only need

to consider k -subsets of inputs. We can therefore use the standard analysis of randomized constructions of bipartite expanders to show that, for the right choice of parameters, a k -independent function is likely to be k -unique. For completeness we provide an analysis here. In our exposition it will be convenient to parameterize the k -uniqueness or k -majority-uniqueness of our constructions in terms of a positive integer κ such that $k = 2^\kappa$.

Lemma 2. *For every choice of positive integers c, n, κ let $\Gamma : \langle n \rangle^c \rightarrow \langle m \rangle^d$ be a 2^κ -independent function. Then,*

- for $m \geq n + \kappa + 1$ and $d \geq 4c$ we have that Γ is 2^κ -unique with probability at least $1 - 2^{-2^{cn}}$.
- for $m \geq n + \kappa + 4$ and $d \geq 8c$ we have that Γ is 2^κ -majority-unique with probability at least $1 - 2^{-2^{cn}}$.

Proof. We will give the proof for k -majority-uniqueness. The proof for k -uniqueness uses the same technique. By a standard argument based on the pigeonhole principle, for Γ to be k -majority-unique it suffices that for all $S \subseteq \langle n \rangle^c$ with $|S| \leq k$ we have that $|\Gamma(S)| > (3/4)d|S|$. Given that Γ is k -independent, we will now bound the probability that there exists a subset S with $|S| \leq k$ such that $|\Gamma(S)| \leq (3/4)d|S|$. For every pair of sets (S, B) satisfying that $S \subseteq \langle n \rangle^c$ with $|S| \leq k$ and $B \subseteq \{1, 2, \dots, d\} \times \langle m \rangle$ with $|B| = (3/4)d|S|$, the probability that $\Gamma(S) \subseteq B$ is given by $\prod_{i=1}^d (|B_i|/2^m)^{|S|}$ where $B_i = \{(i, y) \in B\}$. By the inequality of the arithmetic and geometric means we have that

$$\prod_{i=1}^d \left(\frac{|B_i|}{2^m} \right)^{|S|} \leq \left(\frac{|B|}{d2^m} \right)^{d|S|}.$$

This allows us to ignore the structure of B , and obtain a union bound that matches that of the standard non-compartmentalized probabilistic construction of bipartite expanders. The probability that Γ fails to be k -majority-unique is upper bounded by

$$\sum_{i=2}^k \binom{2^{cn}}{i} \left(\frac{d2^m}{(3/4)di} \right) \left(\frac{(3/4)di}{d2^m} \right)^{di}.$$

For every choice of positive integers c, n, κ , for $m \geq n + \kappa + 4$ and $d \geq 8c$ we get a probability of failure less than $2^{-2^{cn}}$. \square

3.5 A simple k -unique function

In this section we introduce a simple construction of a k -unique function of the form $\Gamma : \langle n \rangle^c \rightarrow \langle m \rangle^d$. We obtain Γ as the last in a sequence $\Gamma_1, \Gamma_2, \dots, \Gamma_c$ of k -unique functions $\Gamma_i : \langle n \rangle^i \rightarrow \langle m \rangle^d$. Each Γ_i for $i > 1$ is defined in terms of Γ_{i-1} . At the bottom of the recursion we tabulate a k -independent function $\Gamma_1 : \langle n \rangle \rightarrow \langle m \rangle^d$. In the general step we apply Γ_{i-1} to the length $i - 1$ prefix of the key $(x_1, x_2, \dots, x_{i-1})$, concatenate the result vector component-wise with the i th character x_i , and apply a simple tabulation function $h_i : \langle m + n \rangle^d \rightarrow \langle m \rangle^d$. The recursion is therefore given by

$$\Gamma_i = h_i \circ (\Gamma_{i-1} \parallel \mathbf{I}^{(d)}) \quad (2)$$

where $\mathbf{I}^{(d)} : \langle n \rangle \rightarrow \langle n \rangle^d$ is the repeated identity function. The following theorem summarizes the properties of Γ in the word RAM model.

Theorem 1. *There exists a randomized data structure that takes as input positive integers c, n, κ and initializes a function $\Gamma : \langle n \rangle^c \rightarrow \langle n + \kappa + 1 \rangle^{4c}$. In the word RAM model with word size w the data structure satisfies the following:*

- The space usage is $O(2^{2n+\kappa} c^3 (n + \kappa) / w)$.
- The evaluation time of Γ is $O(c^2 + c^3 (n + \kappa) / w)$.
- The probability that Γ is 2^κ -unique is at least $1 - 2^{-cn}$.

Proof. Set $m = n + \kappa + 1$ and $d = 4c$. We initialize Γ by tabulating a k -independent function $\Gamma_1 : \langle n \rangle \rightarrow \langle m \rangle^d$ and simple tabulation functions h_2, h_3, \dots, h_c . In total we need to store c functions that each have $O(c)$ character tables with $O(2^{2n+\kappa})$ entries of $O(c(n + \kappa))$ bits. The space usage is therefore $O(2^{2n+\kappa} c^3 (n + \kappa) / w)$. The same bound holds for the time to initialize the data structure.

The evaluation time of Γ can be found by considering the recursion $\Gamma_i = h_i \circ (\Gamma_{i-1} \parallel \mathbb{I}^{(d)})$. At each of the c steps we perform $O(c)$ lookups and take the exclusive-or of $O(c)$ bit strings of length $O(c(n + \kappa))$. The total evaluation time is therefore $O(c^2 + c^3 (n + \kappa) / w)$.

Consider the function $\Gamma_i = h_i \circ (\Gamma_{i-1} \parallel \mathbb{I}^{(d)})$. Conditioned on Γ_{i-1} being k -unique, it is easy to see that $(\Gamma_{i-1} \parallel \mathbb{I}^{(d)})$ is k -unique, and by Lemma 1 we have that Γ_i is k -independent. For our choice of parameters, according to Lemma 2 the probability that Γ_i fails to be k -unique is less than 2^{-2cn} . Therefore, Γ is k -unique if $\Gamma_1, \Gamma_2, \dots, \Gamma_c$ are k -unique. This happens with probability at least $1 - c2^{-2cn} \geq 1 - 2^{-cn}$. \square

Combining Theorem 1 and Lemma 1, we get k -independent hashing in the word RAM model. We state our result in terms of a data structure that represents a family of functions \mathcal{F} . The family is defined as in Lemma 1 and represented by a particular instance of a function Γ , constructed using Theorem 1, together with the parameters of a family of simple tabulation functions.

Corollary 1. *There exists a randomized data structure that takes as input positive integers $u, r = u^{O(1)}, k, t$ and selects a family of functions \mathcal{F} from $[u]$ to $[r]$. In the word RAM model with word length w the data structure satisfies the following:*

- The space used to represent \mathcal{F} , as well as a function $f \in \mathcal{F}$, is $O(ku^{1/t} t^2 (\log u + t \log k) / w)$.
- The evaluation time of f is $O(t^2 + t^2 (\log u + t \log k) / w)$.
- With probability at least $1 - 1/u$ we have that \mathcal{F} is a k -independent family.

Proof. We apply Theorem 1, setting $c = 2t, n = \lceil (\log u) / 2t \rceil, \kappa = \lceil \log k \rceil$. This gives a function $\Gamma : \langle n \rangle^c \rightarrow \langle n + \kappa + 1 \rangle^{4c}$ that is k -unique over $[u]$ with probability at least $1 - 1/u$. To sample a function from the family we follow the approach of Lemma 1 and compose Γ with a simple tabulation function $h : \langle n + \kappa + 1 \rangle^{4c} \rightarrow [r]$. The space used to store Γ follows directly from Theorem 1 and dominates the space used by h . Similarly, the evaluation time of $h \circ \Gamma$ is dominated by the time it takes to evaluate Γ . \square

Remark. For every integer $\tau \geq 1$ we can construct a family $\mathcal{F}^{(\tau)}$ that is k -independent with probability at least $1 - u^{-\tau}$

at the cost of increasing the space usage and evaluation time by a factor τ . The family is defined by

$$\mathcal{F}^{(\tau)} = \{f = \bigoplus_{i=1}^{\tau} f_i \mid f_i \in \mathcal{F}_i\}$$

where each \mathcal{F}_i is constructed independently.

Remark. The recursion in equation (2) is well suited for sequential evaluation. To see this, note that once we have evaluated Γ on a key $x = (x_1, x_2, \dots, x_c)$, a change in the last character only changes the last step of the recursion. It follows that we can generate k -independent variables using amortized time $O(t)$ and space close to $O(ku^{1/t})$. To our knowledge, this presents the best space-time tradeoff for the generation of k -independent variables in the case where we do not have access to multiplication over a suitable finite field as in [7].

3.6 A divide and conquer approach

In this section we introduce a data structure for representing a k -majority-unique function that offers a faster evaluation time at the cost of using more space. As in the simple construction from Theorem 1 we use the technique of alternating between expansion and independence, but rather than reading a single character at the time, we view the key as composed of two characters $x = (x_1, x_2)$ and recurse on each. In the previous section we increased the size of the domain of our k -unique function by concatenating part of the key, forming the k -unique function $\Gamma \parallel \mathbb{I}^{(d)}$. If we use only a few large characters this approach becomes very costly in terms of the space required to store the simple tabulation function h in the composition $h \circ (\Gamma \parallel \mathbb{I}^{(d)})$. To be able to efficiently recurse on large characters we show that the function $\Upsilon((x_1, x_2)) = \Gamma(x_1) \parallel \Gamma(x_2)$ is k -majority-unique.

Lemma 3. *Let $\Gamma : \langle n \rangle^c \rightarrow \langle m \rangle^d$ be a k -majority-unique function. Then $\Gamma \parallel \Gamma : \langle n \rangle^c \times \langle n \rangle^c \rightarrow \langle 2m \rangle^d$ is k -unique.*

Proof. To ease notation we define $\Upsilon = \Gamma \parallel \Gamma$. Let $x = (x_1, x_2)$ denote an element of $\langle n \rangle^c \times \langle n \rangle^c$. For $S \subseteq \langle n \rangle^c \times \langle n \rangle^c$ define $S_{1,a} = \{x \in S \mid x_1 = a\}$. The following holds for every $x = (x_1, x_2) \in S$.

$$\begin{aligned} & |\Upsilon(\{x\}) \setminus \Upsilon(S \setminus \{x\})| \\ &= |\Upsilon(\{x\}) \setminus (\Upsilon(S \setminus S_{1,x_1}) \cup \Upsilon(S_{1,x_1} \setminus \{x\}))| \\ &= |(\Upsilon(\{x\}) \setminus \Upsilon(S \setminus S_{1,x_1})) \cap (\Upsilon(\{x\}) \setminus \Upsilon(S_{1,x_1} \setminus \{x\}))| \\ &\geq |\Upsilon(\{x\}) \setminus \Upsilon(S \setminus S_{1,x_1})| + |\Upsilon(\{x\}) \setminus \Upsilon(S_{1,x_1} \setminus \{x\})| \\ &\quad - |\Upsilon(\{x\})|. \end{aligned} \tag{3}$$

We will show that for every $S \subseteq \langle n \rangle^c \times \langle n \rangle^c$ with $|S| \leq k$ there exists a key $(x_1, x_2) \in S$ such that $|\Upsilon(\{x\}) \setminus \Upsilon(S \setminus \{x\})| > 0$. We begin by choosing the first component of x . Let $\pi_j(S) = \{x_j \mid x \in S\}$ denote the set of j th components of S . By the k -majority-uniqueness of Γ , considering the set $\pi_1(S)$, we have that

$$\exists x_1 \in \pi_1(S) : \forall x \in S_{1,x_1} : |\Upsilon(\{x\}) \setminus \Upsilon(S \setminus S_{1,x_1})| > d/2.$$

Fix x_1 with this property and consider the choice of x_2 . By the k -majority-uniqueness of Γ , considering the set $\pi_2(S)$, we have that

$$\begin{aligned} & \forall x_1 \in \pi_1(S) : \exists x_2 \in \pi_2(S_{1,x_1}) : \\ & |\Upsilon(\{x\}) \setminus \Upsilon(S_{1,x_1} \setminus \{x\})| > d/2. \end{aligned}$$

We can therefore always find a key $(x_1, x_2) \in S$ such that both $|\Upsilon(\{x\}) \setminus \Upsilon(S \setminus S_{1,x_1})| > d/2$, $|\Upsilon(\{x\}) \setminus \Upsilon(S_{1,x_1} \setminus \{x\})| > d/2$ are satisfied. The result follows from equation (3) where we use the fact that $|\Upsilon(\{x\})| = d$. \square

We will give a recursive construction of a k -majority-unique function of the form $\Gamma_i : \langle n \rangle^{2^i} \rightarrow \langle m \rangle^{2^{i+3}}$. Let $h_i : \langle 2m \rangle^{2^{i+2}} \rightarrow \langle m \rangle^{2^{i+3}}$ be a simple tabulation function. For $i > 0$ the recursion takes the following form.

$$\Gamma_i = h_i \circ (\Gamma_{i-1} \parallel \Gamma_{i-1}). \quad (4)$$

At the bottom of the recursion we tabulate a k -independent function Γ_0 .

Theorem 2. *There exists a randomized data structure that takes as input positive integers λ, n, κ and initializes a function $\Gamma : \langle n \rangle^{2^\lambda} \rightarrow \langle n + \kappa + 4 \rangle^{2^{\lambda+3}}$. In the word RAM model with word length w the data structure satisfies the following:*

- The space usage is $O(2^{2(n+\kappa+\lambda)}(n+\kappa)/w)$.
- The evaluation time of Γ is $O(2^\lambda(\lambda + 2^\lambda(n+\kappa)/w))$.
- With probability at least $1 - 2^{-2n+1}$ we have that Γ is 2^κ -majority-unique.

Proof. Let $m = n + \kappa + 4$. We initialize Γ by tabulating Γ_0 and the character tables of the simple tabulation functions $h_1, h_2, \dots, h_\lambda$ where $h_i : \langle 2m \rangle^{2^{i+2}} \rightarrow \langle m \rangle^{2^{i+3}}$. In total we have $O(2^\lambda)$ tables with $O(2^{2(n+\kappa)})$ entries of $O(2^\lambda(n+\kappa))$ bits, resulting in a total space usage of $O(2^{2(n+\kappa+\lambda)}(n+\kappa)/w)$.

Let $T(i)$ denote the evaluation time of Γ_i . For $i = 0$ we can evaluate Γ_0 by performing a single lookup in $O(1)$ time. For $i > 0$ evaluating $h_i \circ (\Gamma_{i-1} \parallel \Gamma_{i-1})$ takes two evaluations of Γ_{i-1} followed by evaluating h_i on their concatenated output using $O(2^i(1 + 2^i(n+\kappa)/w))$ operations. The recurrence takes the form

$$T(i) \leq \begin{cases} 2T(i-1) + O(2^i(1 + 2^i(n+\kappa)/w)) & \text{if } i > 0 \\ O(1) & \text{if } i = 0 \end{cases}$$

The solution to the recurrence is $O(2^i(i + 2^i(n+\kappa)/w))$.

We now turn our attention to the probability that $\Gamma_i = h_i \circ (\Gamma_{i-1} \parallel \Gamma_{i-1})$ fails to be k -majority-unique. Conditional on Γ_{i-1} being k -majority-unique, by Lemma 3 we have that $(\Gamma_{i-1} \parallel \Gamma_{i-1})$ is k -unique and composing it with h_i gives us a k -independent function. For our choice of parameters, according to Lemma 2 the probability that Γ_i fails to be k -majority-unique is less than $2^{-2^{i+1}n}$. Therefore, Γ is k -majority-unique if $\Gamma_0, \Gamma_1, \dots, \Gamma_\lambda$ are k -majority-unique. This happens with probability at least $1 - \sum_{i=0}^\lambda 2^{-2^{i+1}n} \geq 1 - 2^{-2n+1}$. \square

Remark. The recursion in equation (4) is well suited for parallelization. If we have c processors working in lock-step with some small shared memory we can evaluate Γ with domain $\langle n \rangle^c$ in time $O(c)$.

Corollary 2. *There exists a randomized data structure that takes as input positive integers $u, r = u^{O(1)}, k, t$ and selects a family of functions \mathcal{F} from $[u]$ to $[r]$. In the word RAM model with word length w the data structure satisfies the following:*

- The space used to represent \mathcal{F} , as well as a function $f \in \mathcal{F}$, is $O(ku^{1/t}t(\log u + t \log k)/w)$.
- The evaluation time of f is $O(t \log t + t(\log u + t \log k)/w)$.
- With probability at least $1 - u^{-1/t}$ we have that \mathcal{F} is a k -independent family.

Proof. Apply Theorem 2 with parameters $\lambda = \lceil \log t \rceil + 1$, $n = \lceil (\log u)/2t \rceil + 1$, and $\kappa = \lceil \log k \rceil$. This gives is a function Γ that is k -unique over $[u]$ with probability at least $1 - u^{-1/t}$. The family \mathcal{F} is defined by the composition of Γ with a suitable simple tabulation function following the approach of Lemma 1. \square

3.7 Balancing time and space

Theorem 1 yielded a k -unique function over $\langle n \rangle^c$ with an evaluation time of about $O(c^2)$ while using linear space in k . Theorem 2 resulted in an evaluation time of about $O(c \log c)$, using quadratic space in k . Under a mild restriction on k , the two techniques can be combined to obtain an evaluation time of $O(c \log c)$ and linear space in k . We take the construction from Theorem 2 as our starting point, but instead of tabulating the character tables of h_1, \dots, h_λ we replace them with more space efficient k -independent functions that we construct using Theorem 1.

Theorem 3. *There exists a randomized data structure that takes as input positive integers $\lambda, n, \kappa = O(n)$ and initializes a function $\Gamma : \langle n \rangle^{2^\lambda} \rightarrow \langle n + \kappa + 4 \rangle^{2^{\lambda+3}}$. In the word RAM model with word size w the data structure satisfies the following:*

- The space usage is $O(2^{n+\kappa+2\lambda}n/w)$.
- The evaluation time of Γ is $O(2^\lambda(\lambda + 2^\lambda n/w))$.
- With probability at least $1 - 2^{-n+1}$ we have that Γ is 2^κ -majority-unique.

Proof. At the top level, the recursion underlying Γ takes the same form as in Theorem 2.

$$\Gamma_i = h_i \circ (\Gamma_{i-1} \parallel \Gamma_{i-1}).$$

The functions $h_i : \langle 2m \rangle^{2^{i+2}} \rightarrow \langle m \rangle^{2^{i+3}}$ are simple tabulation functions with $m = n + \kappa + 4$. Each h_i is constructed from 2^{i+2} character tables $h_{i,j} : \langle 2m \rangle \rightarrow \langle m \rangle^{2^{i+3}}$. Theorem 2 only assumes that the character tables $h_{i,j}$ are k -independent functions. We will apply Theorem 1 to construct a function Υ that we for each character table $h_{i,j}$ compose with a simple tabulation function $g_{i,j}$ in order to construct $h_{i,j}$. By the restriction that $\kappa = O(n)$ we have that $m = O(n)$. We set the parameters of Υ to $\hat{c} = O(1)$, $\hat{n} = \lceil n/2 \rceil$, $\hat{\kappa} = \kappa$ such that $\langle 2m \rangle$ can be embedded in $\langle \hat{n} \rangle^{\hat{c}}$. Furthermore, Υ uses $O(2^{n+\kappa}n/w)$ words of space, can be evaluated in $O(1)$ operations, and is k -unique with probability at least $1 - 2^{-n-1}$. Because Υ has $O(1)$ output characters, the time to evaluate $h_{i,j} = g_{i,j} \circ \Upsilon$ is no more than a constant times the word length of the output of $h_{i,j}$. The time to evaluate Γ therefore only increases by a constant factor compared to the evaluation time in Theorem 2.

The probability of failure of Γ to be k -majority-unique is the same as in Theorem 2, provided that Υ does not fail to be k -unique. This gives a total probability of failure of less than $2^{-2n+1} + 2^{-n-1} < 2^{-n+1}$.

We only store a single Υ and the character tables of $g_{i,j}$ that we use to simulate the character tables $h_{i,j}$. From the parameters of Υ we have that $g_{i,j}$ uses $O(1)$ character tables with $O(2^{n+\kappa})$ entries of $O(2^i n/w)$ words. The space usage is dominated by the $O(2^\lambda)$ character tables of h_λ that use space $O(2^{n+\kappa+2\lambda} n/w)$ in total. \square

Corollary 3. *There exists a randomized data structure that takes as input positive integers u , $r = u^{O(1)}$, t , $k = u^{O(1/t)}$ and selects a family of functions \mathcal{F} from $[u]$ to $[r]$. In the word RAM model with word length w the data structure satisfies the following:*

- The space used to represent \mathcal{F} , as well as a function $f \in \mathcal{F}$, is $O(ku^{1/t} t(\log u)/w)$.
- The evaluation time of f is $O(t \log t + t(\log u)/w)$.
- With probability at least $1 - u^{-1/t}$ we have that \mathcal{F} is a k -independent family.

Proof. Apply Theorem 3 with parameters $\lambda = \lceil \log t \rceil$, $n = \lceil (\log u)/t \rceil + 1$, and $\kappa = \lceil \log k \rceil$. This gives is a function Γ that is k -unique over $[u]$ with probability at least $1 - u^{-1/t}$. The family \mathcal{F} is defined by the composition of Γ with a suitable simple tabulation function following the approach of Lemma 1. \square

4. CONCLUSION

We have presented new constructions of k -independent hash functions that come close to Siegel’s lower bound on the space-time tradeoff for such functions. An interesting open problem is whether the gap to the lower bound can be closed. From the perspective of efficient expanders it would be very interesting to achieve space $o(k)$ while preserving computational efficiency. Of course, such a result is not possible via k -independence.

5. ACKNOWLEDGEMENTS

The research of Tobias Christiani and Rasmus Pagh has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. 614331.

Mikkel Thorup’s research is partly supported by Advanced Grant DFF-0602-02499B from the Danish Council for Independent Research under the Sapere Aude research career programme.

We thank the STOC reviewers for insightful comments that helped us improve the exposition.

6. REFERENCES

- [1] A. Z. Broder, A. M. Frieze, and E. Upfal. Static and dynamic path selection on expander graphs: a random walk approach. *Random Structures & Algorithms*, 14(1):87–109, 1999.
- [2] H. Burhman, P. B. Miltersen, J. Radhakrishnan, and S. Venkatesh. Are bitvectors optimal? *SIAM J. Comput.*, 31(6):1723–1744, 2002.
- [3] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson. Randomness conductors and constant-degree lossless expanders. In *Proc. STOC ’02*, pages 659–668, 2002.
- [4] J. L. Carter and M. N. Wegman. Universal classes of hash functions. In *Proc. STOC ’77*, pages 106–112, 1977.
- [5] J. L. Carter and M. N. Wegman. New hash functions and their use in authentication and set equality. *J. Comput. System Sci.*, 22(3):265–279, 1981.
- [6] B. Chor, O. Goldreich, J. Hastad, J. Freidmann, S. Rudich, and R. Smolensky. The bit extraction problem or t-resilient functions. In *Proc. FOCS ’85*, pages 396–407, 1985.
- [7] T. Christiani and R. Pagh. Generating k -independent variables in constant time. In *Proc. FOCS ’14*, pages 196–205, 2014.
- [8] M. Dietzfelbinger and F. Meyer auf der Heide. A new universal class of hash functions and dynamic hashing in real time. In *Proc. ICALP ’90*, pages 6–19, 1990.
- [9] M. Dietzfelbinger. On randomness in hash functions (invited talk). In *Proc. STACS ’12*, pages 25–28, 2012.
- [10] M. Dietzfelbinger and P. Woelfel. Almost random graphs with simple hash functions. In *Proc. STOC ’03*, pages 629–638, 2003.
- [11] V. Guruswami, C. Umans, and S. Vadhan. Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes. *J. ACM*, 56(4):20:1–20:34, 2009.
- [12] T. Hagerup. Sorting and searching on the word RAM. In *Proc. STACS ’98*, pages 366–398, 1998.
- [13] P. Indyk and A. Gilbert. Sparse recovery using sparse matrices. *Proc. IEEE*, 98(6):937–947, 2010.
- [14] A. Joffe. On a set of almost deterministic k -independent random variables. *Ann. Prob.*, 2(1):161–162, 1974.
- [15] K. S. Kedlaya and C. Umans. Fast modular composition in any characteristic. In *Proc. FOCS ’08*, pages 146–155, 2008.
- [16] B. W. Kernighan and D. M. Ritchie. *The C Programming Language*. Prentice Hall, Englewood Cliffs, New Jersey, second edition, 1988.
- [17] K. G. Larsen. Higher cell probe lower bounds for evaluating polynomials. In *Proc. FOCS ’12*, pages 293–301, 2012.
- [18] P. B. Miltersen. Lower bounds for static dictionaries on RAMs with bit operations but no multiplications. In *Proc. ICALP ’96*, pages 442–453, 1996.
- [19] A. Pagh and R. Pagh. Uniform hashing in constant time and optimal space. *SIAM J. Comput.*, 38(1):85–96, 2008.
- [20] F. Parvaresh and A. Vardy. Correcting errors beyond the Guruswami–Sudan radius in polynomial time. In *Proc. FOCS ’05*, pages 285–294, 2005.
- [21] M. Pătraşcu and M. Thorup. Dynamic integer sets with optimal rank, select, and predecessor search. In *Proc. FOCS ’14*, pages 166–175, 2014.
- [22] A. Siegel. On universal classes of extremely random constant-time hash functions. *SIAM J. Comput.*, 33(3):505–543, 2004.
- [23] A. Ta-Shma, C. Umans, and D. Zuckerman. Lossless condensers, unbalanced expanders, and extractors. *Combinatorica*, 27(2):213–240, 2007.
- [24] M. Thorup. Simple tabulation, fast expanders, double tabulation, and high independence. In *Proc. FOCS ’13*, pages 90–99, 2013.