The Complexity of Distributions

Emanuele Viola
College of Computer and Information Science, Northeastern University, Boston, MA 02115.
Email: viola@ccs.neu.edu

Abstract— Complexity theory typically studies the complexity of computing a function $h(x) : \{0,1\}^m \to \{0,1\}^n$ of a given input x. We advocate the study of the complexity of generating the distribution $h(x)$ for uniform x, given random bits. Our main results are:

(1) Any function $f : \{0,1\}^\ell \to \{0,1\}^n$ such that (i) each output bit f_i depends on $o(\log n)$ input bits, and (ii) $\ell \leq \log_2 \binom{n}{\alpha n} + n^{0.99}$, has output distribution $f(U)$ at statistical distance $\geq 1 - 1/n^{0.49}$ from the uniform distribution over n-bit strings of hamming weight αn.

We also prove lower bounds for generating $(X, b(X))$ for boolean b, and in the case in which each bit f_i is a small-depth decision tree.

These lower bounds seem to be the first of their kind; the proofs use anti-concentration results for the sum of random variables.

(2) Lower bounds for generating distributions imply succinct data structures lower bounds. As a corollary of (1), we obtain the first lower bound for the membership problem of representing a set $S \subseteq [n]$ of size αn, in the case where $1/\alpha$ is a power of 2: If queries "$i \in S$?" are answered by non-adaptively probing $o(\log n)$ bits, then the representation uses $\geq \log_2 \binom{n}{\alpha n} + \Omega(\log n)$ bits.

(3) Upper bounds complementing the bounds in (1) for various settings of parameters.

(4) Uniform randomized AC^0 circuits of $\mathrm{poly}(n)$ size and depth $d = O(1)$ with error ϵ can be simulated by uniform randomized AC^0 circuits of $\mathrm{poly}(n)$ size and depth $d+1$ with error $\epsilon + o(1)$ using $\leq (\log n)^{O(\log \log n)}$ random bits.

Previous derandomizations [Ajtai and Wigderson '85; Nisan '91] increase the depth by a constant factor, or else have poor seed length.

Keywords-distribution; lower bounds; data structure; pseudorandomness; k-wise independent distributions

1. Introduction

Complexity theory, with some notable exceptions, typically studies the complexity of computing a function $h(x) : \{0,1\}^m \to \{0,1\}^n$ of a given input x. We advocate the study of the complexity of generating the output distribution $h(x)$ for random x, given random bits. This question can be studied for a variety of computational models. In this work we focus on restricted models such as small bounded-depth circuits with unbounded fan-in (AC^0) or bounded fan-in (NC^0).

An interesting example of a function h for which computing $h(x)$ is harder than generating its output distribution is $h(x) := (x, \mathrm{parity}(x))$, where $\mathrm{parity}(x) := \sum_i x_i \mod 2$. Whereas small AC^0 circuits cannot compute parity (cf. [21]), Babai [4] and Boppana and Lagarias [8] show a function f whose output distribution equals that of $(x, \sum_i x_i \mod 2)$ for random $x \in \{0,1\}^n$, and each output bit f_i depends on just 2 input bits (so $f \in \mathrm{NC}^0$):

$$f(x_1, x_2, \ldots, x_n) := (x_1, x_1 + x_2, x_2 + x_3, \ldots, x_{n-1} + x_n,\ x_n). \quad (1)$$

This construction is useful for proving average-case lower bounds, see [4] and [5, Corollary 22].

Later, Impagliazzo and Naor [22] extend the construction (1) to show that small AC^0 circuits can even generate $(x, b(x))$ for more complicated functions, such as inner product $b(x) = x_1 \cdot x_2 + x_3 \cdot x_4 + \cdots + x_{n-1} \cdot x_n$. They use this to construct cryptographic pseudorandom generators computable by poly-size AC^0 circuits based on the hardness of the subset-sum problem, and similar techniques are useful in constructing depth-efficient generators based on other assumptions [2], [35].

We mention that cryptography provides several candidate functions h for which computing $h(x)$ is harder than generating its output distribution (e.g., take h^{-1} to be a one-way permutation). However, in this work we focus on unconditional results.

The work by Mossel, Shpilka, and Trevisan [28] provides another example of the power of NC^0 circuits in generating distributions: NC^0 circuits can generate small-bias distributions with non-trivial stretch.

The surprising nature of the above constructions, and their usefulness (for example for average-case lower bounds and pseudorandom generators) raises the challenge of understanding the complexity of generating distributions, and in particular proving lower bounds:

Challenge 1.1. *Exhibit an explicit map $b : \{0,1\}^n \to \{0,1\}$ such that the distribution $(X, b(X)) \in \{0,1\}^{n+1}$ cannot be generated by $\mathrm{poly}(n)$-size AC^0 circuits given random bits.*

Current lower-bounding techniques appear unable to tackle questions such as Challenge 1.1 (which, to our knowledge, is open even for DNFs). As we have seen, standard "hard functions" b such as parity and inner product have the property that $(X, b(X))$ can be generated exactly by small AC^0 circuits. Along the way, in this work we point out that the same holds for any symmetric b (e.g., $\mathrm{majority}$, $\mod 3$) (up to an exponentially small error). In fact, weaker models often suffice.

Supported by NSF grant CCF-0845003.

This suggests that our understanding of even these simple models is incomplete, and that pursuing the above direction may yield new proof techniques.

1.1. Our results

In this work we prove several "first-of-their-kind" lower bounds for generating distributions. We also complement these with upper bounds, and establish connections to other areas such as succinct data structures, derandomization, and switching networks.

Lower bounds.: We aim to bound from below the statistical (a.k.a. variation) distance Δ between a distribution D on n bits and the output distribution of a "simple" function $f : \{0,1\}^\ell \to \{0,1\}^n$ over random input $U \in \{0,1\}^\ell$:

$$\begin{aligned}\Delta(f(U), D) &:= \max_{T \subseteq \{0,1\}^n} \left| \Pr_U[f(U) \in T] - \Pr_D[D \in T] \right| \\ &= \frac{1}{2} \sum_a |\Pr[f(U) = a] - \Pr[D = a]|.\end{aligned}$$

In addition to being a natural measure, small statistical distance (as opposed to equality) is sufficient in typical scenarios (e.g., pseudorandomness). Moreover, this work shows that statistical distance lower bounds imply lower bounds for succinct data structure problems, and uses this implication to derive a new lower bound for a central data structure problem (Corollary 1.7).

The next convenient definition generalizes NC^0 (which corresponds to $d = O(1)$).

Definition 1.2. A function $f : \{0,1\}^\ell \to \{0,1\}^n$ is *d-local* if each output bit f_i depends on $\leq d$ input bits.

Our first lower bound is for generating the uniform distribution $D_{=\alpha}$ over n-bit strings with αn ones (i.e., hamming weight αn). This distribution arises frequently. For example, we will see that it is related to generating $(X, b(X))$ for symmetric b, and to the membership problem in data structures.

Theorem 1.3 (Lower bound for generating "$= \alpha$" locally)**.** For any $\alpha \in (0,1)$ and any $\delta < 1$ there is $\epsilon > 0$ such that for all sufficiently large n for which αn is an integer:

Let $f : \{0,1\}^\ell \to \{0,1\}^n$ be an $(\epsilon \log n)$-local function where $\ell \leq \log_2 \binom{n}{\alpha n} + n^\delta$.

Let $D_{=\alpha}$ be the uniform distribution over n-bit strings with αn ones.

Then $\Delta(f(U), D_{=\alpha}) \geq 1 - O(1/n^{\delta/2})$.

For $\alpha = 1/2$, Theorem 1.3 matches the 1-local identity function $f : \{0,1\}^n \to \{0,1\}^n$, $f(u) := u$, achieving $\Delta(U, D_{=1/2}) \leq 1 - O(1/\sqrt{n})$ (a standard bound, see Fact 2.2). For $\alpha < 1/2$, upper bounds are a bit more involved. There are $\mathrm{poly}\log(n)$-local functions again achieving statistical distance $\leq 1 - O(1/\sqrt{n})$. We refine this to also obtain input length $\ell = \log_2 \binom{n}{\alpha n} + n/\mathrm{poly}\log n$ (see [38]).

For generating $(X, b(X))$ for boolean b obviously no lower bound larger than $1/2$ holds. We establish $1/2 - o(1)$ for the function which checks if the hamming weight of X modulo p is between 0 and $(p-1)/2$. We call it "majority modulo p," majmod for short.

Theorem 1.4 (Lower bound for generating $(X, \mathrm{majmod}\ X)$ locally)**.** For any $\delta < 1$ there is $\epsilon > 0$ such that for all sufficiently large n: Let $p \in [0.25 \log n, 0.5 \log n]$ be a prime number, and let $\mathrm{majmod} : \{0,1\}^n \to \{0,1\}$ be defined as

$$\mathrm{majmod}(x) = 1 \Leftrightarrow \sum_{i \leq n} x_i \mod p \in \{0, 1, \ldots, (p-1)/2\}.$$

Let $f : \{0,1\}^\ell \to \{0,1\}^{n+1}$ be an $(\epsilon \log n)$-local function where $\ell \leq n + n^\delta$.

Then $\Delta(f(U), (X, \mathrm{majmod}\ X)) \geq 1/2 - O(1/\log n)$.

Theorem 1.4 is tight up to the $O(.)$: it can be verified that $\Pr_X[\mathrm{majmod}(X) = 0] = 1/2 - \Theta(1/\log n)$, hence $\Delta((X,1), (X, \mathrm{majmod}\ X)) \leq 1/2 - O(1/\log n)$. Moreover, we show a $\mathrm{poly}\log(n)$-local function with statistical distance $\leq 1/n$ (see [38]).

Theorems 1.3 and 1.4 may hold even when the input length ℓ is unbounded, but it is not clear to us how to prove such statistical bounds in those cases. However we can prove weaker statistical bounds when the input length ℓ is unbounded, and these hold even against the stronger model where each bit of the function f is a *decision tree*. We call such a function *forest*, to distinguish it from a function computable by a single decision tree.

Definition 1.5. A function $f : \{0,1\}^\ell \to \{0,1\}^n$ is a d-forest if each bit f_i is a decision tree of depth d.

A d-forest function is also 2^d local, so the previous theorems yield bounds for $d = (\log(\epsilon \log n))$-forests. We prove bounds for $d = \epsilon \log n$ with a different argument.

Theorem 1.6 (Lower bound for generating "$= 1/2$" or $(X, \mathrm{majority}\ X)$ by forest)**.** Let $f : \{0,1\}^* \to \{0,1\}^n$ be a d-forest function. Then:

(1) $\Delta(f(U), D_{=1/2}) \geq 2^{-O(d)} - O(1/n)$, where $D_{=1/2}$ is the uniform distribution over n-bit strings with $n/2$ ones.

(2) $\Delta(f(U), (X, \mathrm{majority}\ X)) \geq 2^{-O(d)} - O(1/n)$.

A similar bound to (1) also holds for generating $D_{=\alpha}$; we pick $\alpha = 1/2$ for simplicity.

Theorem 1.6 complements the existence of d-forest functions achieving statistical distance $O(1/n)$ where $d = O(\log n)$ for (1) and $d = O(\log^2 n)$ for (2). (In fact, $d = O(\log n)$ may hold for both, see [38].) We obtain such functions by establishing a simple connection with results on *switching networks*, especially by Czumaj et al. [13]: we prove they imply forest upper bounds. These upper bounds are not explicit; explicit upper bounds are known for $d = \mathrm{poly}\log n$, see [38].

For AC^0 circuits, there are constructions that are both explicit and achieve exponentially small error. In particular, building on results by Matias and Vishkin [27] and Hagerup [20], we exhibit AC^0 circuits of size $\mathrm{poly}(n)$ and depth $O(1)$ whose output distribution has statistical distance $1/2^n$ from the distribution $(X, \sum_i X_i) \in \{0,1\}^n \times \{0,1,\ldots,n\}$ for uniform $X \in \{0,1\}^n$.

The above lower bounds are obtained via new proof techniques also using anti-concentration results for the sum of random variables. We provide an overview of the proof of Theorem 1.3 in §2.

Motivation: Succinct data structures lower bounds.: Succinct data structures aim to compress data using a number of bits close to the information-theoretic minimum while at the same time supporting interesting queries. For a number of problems, tight bounds are known, cf. [33], [37], [34], [15]. But there remains a large gap for the notable *membership* problem which asks to store a subset x of $[n]$ of size ℓ (think of x as an n-bit string of weight ℓ) in $\lceil \log_2 \binom{n}{\ell} \rceil + r$ bits, where r is as small as possible, while being able to answer the query "is i in x" by reading few bits of the data structure [10], [31], [32], [33], [37]. In particular, previous to this work there was no lower bound in the case when $\ell := \alpha n$ for $1/\alpha = 2^a$ a fixed power of two. Note that the lower bound in [37] does not apply to that case; intuitively, that is because the techniques there extend to the problem of succinctly storing arrays over the alphabet $[1/\alpha]$, but when $1/\alpha = 2^a$ no lower bound holds there: using a bits per symbol yields redundancy $r = 0$.

Using different techniques, as a corollary to our lower bound for generating the "$= \alpha$" distribution (Theorem 1.3) we obtain the first lower bound for the membership problem in the case where the set-size is a power-of-two fraction of the universe.

Corollary 1.7 (Lower bound for membership)**.** *For any $\alpha \in (0,1)$ there is $\epsilon > 0$ such that for all large enough n for which αn is an integer:*

Suppose one can store subsets x of $[n]$ of size αn in $m := \lceil \log_2 \binom{n}{\alpha n} \rceil + r$ bits, while answering "is i in x" by non-adaptively reading $\leq \epsilon \log n$ bits of the data structure. Then $r \geq 0.49 \log n$.

Again, Corollary 1.7 is tight for $\alpha = 1/2$ up to the constant 0.49, since $\log_2 \binom{n}{n/2} = n - \Theta(\log n)$, and using $m = n$ bits the problem is trivial. For $\alpha < 1/2$ it is not clear what lower bound on r one should expect, as surprising upper bounds hold for related problems [10], [32], [15]. In particular, the recent work by Dodis, Pătraşcu, and Thorup [15] yields $r = 1$ for storing arrays (non-adaptively reading $O(\log n)$ bits). It remains to be seen whether their techniques apply to the membership problem too.

We obtain Corollary 1.7 from Theorem 1.3 by establishing the simple and general fact that lower bounds for generating distributions somewhat close to a distribution D imply succinct data structure lower bounds for storing support(D). The following claim formalizes this for the membership problem, where $D = D_{=\alpha}$ is the uniform distribution over n-bit strings with αn ones.

Claim 1.8. *Suppose one can store subsets x of $[n]$ of size αn in $m := \lceil \log_2 \binom{n}{\alpha n} \rceil + r$ bits, while answering "is i in x" by non-adaptively reading q bits of the data structure. Then there is a q-local function $f : \{0,1\}^m \to \{0,1\}^n$ such that $\Delta(f(U), D_{=\alpha}) \leq 1 - 2^{-r-1}$.*

Proof: The i-th output bit of f is the algorithm answering "is i in x." Feed f random bits. With probability $\binom{n}{\alpha n} / 2^{\lceil \log_2 \binom{n}{\alpha n} \rceil + r} \geq 1/2^{r+1}$ the input is uniform over encodings of subsets of $[n]$ of size αn, in which case the statistical distance is 0. If we distinguish in every other case, the distance is at most $1 - 1/2^{r+1}$. ∎

Similar considerations extend to adaptive bit-probes and cell probes, corresponding to forest functions (in the latter case, over the alphabet $[n]$ instead of $\{0,1\}$). While one could prove lower bounds for data structures without using this approach, Claim 1.8 and Corollary 1.7 appear to suggest an uncharted direction. Finally, we note that none of the upper bounds mentioned earlier is an obstacle to using Claim 1.8, since those upper bounds use input length that is larger than the information-theoretic minimum by a quantity polynomial in the statistical distance gap, while for Claim 1.8 a logarithmic dependence suffices. Whether the lower bounds for generating $D_{=\alpha}$ can be improved in this case is an interesting open problem.

Pseudorandom generators.: The ability to generate a distribution efficiently has obvious applications in pseudorandomness which we now elaborate upon. The ultimate goal of derandomization of algorithms is to remove, or reduce, the amount of randomness used by a randomized algorithm while incurring the *least possible* overhead in other resources, such as time. Typically, this is achieved by substituting the needed random bits with the output of a pseudorandom generator. There are two types of generators. Cryptographic generators [6], [39] (a.k.a. Blum-Micali-Yao) use less resources than the algorithm to be derandomized. In fact, computing these generators can even be done in the restricted circuit class NC^0 [2]. However, unconditional instantiations of these generators are rare, and in particular we are unaware of any unconditional cryptographic generator with large stretch, a key feature for derandomization. By contrast, Nisan-Wigderson generators [30] use more resources than the algorithm to be derandomized, and this looser notion of efficiency allows for more unconditional results [29], [30], [26], [36]. Moreover, all of these results yield generators with large, superpolynomial stretch.

In particular, Nisan [29] shows a generator that fools small AC^0 circuits of depth d with exponential stretch, or seed length $\log^{O(d)} n$. As mentioned above, this generator

uses more resources than the circuits to be derandomized. Specifically, it computes the parity function on $\geq \log^d n$ bits, which requires AC^0 circuits that have either depth $\geq d$ or superpolynomial size. Thus, if one insists on polynomial-size circuits, the derandomized circuit, consisting of the circuit computing the generator and the original circuit, has depth at least twice that of the original circuit. This constant factor blow-up in depth appears necessary for Nisan-Wigderson constructions.

In this work we present a derandomization which only blows up the depth by 1, and uses a number of random bits close to Nisan's (an improvement in the tools we use would let us match the number of random bits in Nisan's result).

Theorem 1.9 (Depth-efficient derandomization of AC^0)**.** The following holds for every d. Let $f : \{0,1\}^* \to \{0,1\}^*$ be computable by uniform randomized AC^0 circuits of $\mathrm{poly}(n)$-size and depth d with error ϵ. Then f is computable by uniform randomized AC^0 circuits of $\mathrm{poly}(n)$-size and depth $d+1$ with error $\epsilon + o(1)$ using $\leq (\log n)^{O(\log \log n)}$ random bits.

Theorem 1.9 is proved by exhibiting a generator whose output looks random to small AC^0 circuits, and yet each of its output bits can be computed by a DNF, i.e., a depth-2 circuit (of size $n^{O(d)}$). Some evidence that such a generator may exist comes from Example (1), which implies a generator mapping $n-1$ bits to n bits that can be shown to look random to AC^0 circuits, and yet each output bit just depends on 2 inputs bits. However, the seed length of this generator is very poor, and it is not clear how to improve on it. Intuitively, one would like to be able to generate the output distribution of Nisan's generator [29] more efficiently than shown in [29]. We were not able to do so, and we raise this as another challenge. (Some recent progress on this question appears in [25].)

For Theorem 1.9, we notice that the recent line of work by Bazzi, Razborov, and Braverman [9] shows that any distribution that is $(k := \log^c n)$-wise independent looks random to small AC^0 circuits of depth d, for a certain constant $c = c(d) \geq d$.

We show how such distributions can be generated by DNFs. Although the constructions of k-wise independent distributions in [11], [1], [19] all require iterated sums of k bits, which for $k := \log^c n$ is unfeasible in our setting, we follow an approach of Mossel, Shpilka, and Trevisan [28] and give an alternative construction using unique-neighbor expanders. Specifically, we use the recent unique-neighbor expanders by Guruswami, Umans, and Vadhan [18].

More related work and discussion.: A result (we already mentioned briefly) by Applebaum, Ishai, Kushilevitz [2] shows, under standard assumptions, that there are pseudorandom distributions computable by NC^0 circuits. Their result is obtained via a generic transformation that turns a distribution D into another "padded" distribution D' that is computable in NC^0 and at the same time maintains interesting properties, such as pseudorandomness (but not stretch). The techniques in [2] do not seem to apply to distributions such as $(x, \sum_i x_i)$, and they destroy stretch, which prevents them from obtaining Theorem 1.9 (regardless of the stretch of the original generator, the techniques in [2] always produce a generator with sublinear stretch).

Under an assumption on the hardness of decoding random linear codes, the same authors show in [3] how to construct generators computable in NC^0 that have linear stretch. Their construction requires generating in NC^0 a uniform "noise vector" $e \in \{0,1\}^n$. They consider two types of noise vectors. The first type is when e has hamming weight exactly pn (think $p = 1/4$), i.e. e comes from the distribution $D_{=pn}$. The results in this paper show that it is impossible to generate such an e in NC^0, regardless of the input length, except with constant statistical distance, see a remark in [38] related to Theorem 1.6. The second type of noise vector is when e is obtained by setting each bit to 1 independently with probability p. This distribution can be trivially generated in NC^0 when $p = 2^{-t}$, using tn bits of randomness, which is much larger than the entropy of the distribution. This loss in randomness is problematic for pseudorandom generator constructions, but the authors of [3] make up for it by applying an extractor. (They use an extractor computable in NC^0 that is implied by [28]). Whether such a noise vector can be generated in NC^0 using randomness close to optimal is an interesting open question.

It is perhaps worthwhile to pause to make a philosophical remark. While the above mentioned works [2], [3] show that, under various assumptions, one can locally generate distributions on n bits with small entropy that look random to any polynomial-time test, by contrast our results show that one cannot locally generate a distribution that is close to being uniform over n-bit strings with $n/2$ ones, which superficially seems a less demanding goal.

Dubrov and Ishai [16] also address the problem of generating distributions, but focus on the randomness complexity, as opposed to our work which emphasizes the complexity of the generation process.

Recently and after a preliminary version [38] of this work, Lovett and the author [25] prove that small AC^0 circuits cannot generate the uniform distribution over any good error-correcting codes. This result does not solve Challenge 1.1 – it does not apply to distributions like $(X, b(X))$ – although it does answer a question asked in a preliminary version of this work [38].

Organization: In §2 we provide the intuition and the proof of our lower bound for generating the "$= \alpha$" distribution locally (Theorem 1.3). The lower bound for generating $(X, \mathrm{majmod}\, X)$ locally (Theorem 1.4) is in §3. The proofs of the other results are omitted due to space restrictions.

2. Intuition and proof of lower bound for generating "$= \alpha$" locally

In this section we prove our lower bound for generating the "$= \alpha$" distribution, restated next.

Theorem 1.3 (Lower bound for generating "$= \alpha$" locally)**.** *(Restated.)* For any $\alpha \in (0,1)$ and any $\delta < 1$ there is $\epsilon > 0$ such that for all sufficiently large n for which αn is an integer:

Let $f : \{0,1\}^\ell \to \{0,1\}^n$ be an $(\epsilon \log n)$-local function where $\ell \leq \log_2 \binom{n}{\alpha n} + n^\delta$.

Let $D_{=\alpha}$ be the uniform distribution over n-bit strings with αn ones.

Then $\Delta(f(U), D_{=\alpha}) \geq 1 - O(1/n^{\delta/2})$.

2.1. Intuition for the proof of Theorem 1.3.

We now explain the ideas behind the proof of Theorem 1.3. For simplicity, we consider the case $\ell = n$ and $\alpha = 1/2$, that is, we want to prove that any $(\epsilon \log n)$-local function $f : \{0,1\}^n \to \{0,1\}^n$ has output distribution $f(U)$ for uniform $U \in \{0,1\}^n$ at statistical distance $\geq 1 - 1/n^{\Omega(1)}$ from the distribution $D_{=1/2}$ uniform over n-bit strings with $n/2$ ones. For simplicity, we denote the latter by $D = D_{=1/2}$.

We start with two warm-up scenarios:

Low-entropy scenario.: Suppose that f is the constant function $f(u) := 0^{n/2}1^{n/2}$. In this case, the simple test

$$T_F := \text{support}(f) = \{0^{n/2}1^{n/2}\}$$

gives $\Pr_U[f(U) \in T_F] = 1$ and $\Pr[D \in T_F] = 1/\binom{n}{n/2} \ll 1/n$, proving the theorem.

We call this the "low-entropy" scenario because $f(U)$ has low entropy.

Anti-concentration scenario.: Suppose that $f(u) := u$. In this case we consider the test

$$T_S := \overline{\text{support}(D)} = \{z : \sum_i z_i \neq n/2\}.$$

Note $\Pr[D \in T_S] = 0$ by definition, while $\Pr_U[f(U) \in T_S] = \Pr[\sum_i U_i \neq n/2] = \binom{n}{n/2}/2^n \geq 1 - O(1/\sqrt{n})$ by a standard bound (Fact 2.2). (Taking T_S to be the complement of the support of D, rather than the support itself, is useful when pasting tests together.)

We call this the "anti-concentration" scenario because the bound $\Pr[\sum_i U_i \neq n/2] \geq 1 - O(1/\sqrt{n})$ is an instance of the general anti-concentration phenomenon that the sum of independent, non-constant, uniform random variables is unlikely to equal any fixed value. Specifically, the bound is a special case ($S_i = U_i \in \{0,1\}$) of the following anti-concentration inequality by Littlewood and Offord (later we use the general case).

Fact 2.1 (Littlewood-Offord anti-concentration [23], [17])**.** Let $S_1, S_2, \ldots, S_t$ be t independent random variables, where S_i is uniform over $\{a_i, b_i\}$ for $a_i \neq b_i$. Then for any integer c, $Pr[\sum_i S_i = c] \leq O(1/\sqrt{t})$.

Having described the two scenarios, we observe that each of them, taken by itself, is not sufficient. This is because the output distribution of the low-entropy function $f(u) = 0^{n/2}1^{n/2}$ has the same probability of passing the anti-concentration test T_S as the distribution D, and similarly in the other case.

We would like to use a similar approach for a generic f. The first step is to partition the input bits u of f as $u = (x, y)$ and rewrite (up to a permutation)

$$f(u) = f(x,y) = h(y) \circ g_1(x_1, y) \circ g_2(x_2, y) \circ \cdots \circ g_s(x_s, y),$$

where each function g_i depends on only the single bit x_i of x (but arbitrarily on y), and has small range: $g_i(x_i, y) \in \{0,1\}^{O(d)} = \{0,1\}^{O(\epsilon \log n)}$. A greedy approach allows for such a decomposition with $|x| = s \geq \Omega(n/d^2) = n/\text{poly}\log n$. Specifically, by an averaging argument a constant fraction of the input bits are adjacent to $\leq O(d)$ output bits. We iteratively collect such a bit x_i and move in y the $\leq O(d^2)$ other input bits adjacent to any of the input bits x_i is adjacent to.

To reduce to the previous scenarios, fix y. Two things can happen: either $\geq \sqrt{n}$ of the functions g_i are fixed, i.e., do not depend on x_i anymore, or at least $s - \sqrt{n} = n/\text{poly}\log n$ take two different values over the choice of x_i. We think of the first case as the low-entropy scenario. Indeed, for this y the output distribution of $f(x, y)$ has small support, and we can hardwire it in the test. Here we use that the input length n of f is close to the information-theoretic minimum necessary to generate D, which is $n - \Theta(\log n)$, and hence removing $\sqrt{n}$ bits of entropy yields a tiny support where D is unlikely to land.

In the second case, intuitively, we would like to use anti-concentration, since we have independent random variables $g_1(x_1, y), g_2(x_2, y), \ldots, g_s(x_s, y)$. Specifically, we let $S_i := \sum_k (g_i(x_i, y))_k$ denote the sum of the bits of g_i, and would like to apply the Littlewood-Offord inequality to argue that $f(U)$ is likely to pass the anti-concentration test T_S, which checks if the hamming weight of f is $\neq n/2$. However, the following problem arises. It may be the case that, for example,

$$g_i(0, y) = 01, \text{ and } g_i(1, y) = 10,$$

corresponding to the constant variable $S_i \equiv 1$. In this case, the value of g_i is not fixed, hence this is not a low-entropy scenario, but on the other hand it does not contribute to anti-concentration, since $S_i \equiv 1$. In fact, precisely such functions g_i arise when running this argument on the function that generates the uniform distribution over n-bit strings with an even number of ones, which can be done with locality 2 via the construction (1) in §1.

We solve this problem as follows. We add to our test the check T_0 that $\leq 2\sqrt{n}$ of the blocks of output bits corresponding to g_i are all 0. Since the blocks are small

(recall $g_i \in \{0,1\}^{O(\epsilon \log n)}$), the distribution D will have $\geq n^{0.99}$ such blocks with high probability, and so will almost never pass T_0.

Consider however what happens with $f(x,y)$, for a fixed y. If $\leq 2\sqrt{n}$ functions $g_i(x_i,y)$ can output all zeros (for some $x_i \in \{0,1\}$), then $f(x,y) \in T_0$ for every x, and we are again done. Otherwise, since $\leq \sqrt{n}$ functions g_i are fixed, we have $2\sqrt{n} - \sqrt{n} = \sqrt{n}$ functions $g_i(x_i,y)$ that take two different values over $x_i \in \{0,1\}$, and one of the two is all zero. That means that the other value is not all zero, and hence has a sum of bits $a_i > 0$. We are now in the position to apply the Littlewood-Offord anti-concentration inequality, since we have $\geq \sqrt{n}$ independent variables S_i, each uniform over $\{0, a_i\}$ for $a_i \neq 0$. The inequality guarantees that $f(x,y) \in T_S$ with probability $\geq 1 - 1/n^{\Omega(1)}$, and this concludes the overview of the proof of Theorem 1.3.

We now proceed with the formal proof. We use several times the following standard approximation of the binomial by the binary entropy function $H(x) = x\log_2(1/x) + (1-x)\log_2(1/(1-x))$:

Fact 2.2 (Lemma 17.5.1 in [12])**.** *For $0 < p < 1, q = 1-p$, and n such that np is an integer,*

$$\frac{1}{\sqrt{8npq}} \leq \binom{n}{pn} \cdot 2^{-H(p)n} \leq \frac{1}{\sqrt{\pi npq}}.$$

2.2. Proof of Theorem 1.3

We begin by bounding some parameters in a way that is convenient for the proof. First, we assume without loss of generality that $\alpha \leq 1/2$ (otherwise, complement the output of f). Next, we bound $\ell = \Theta(H(\alpha)n)$. For this, first note that if $\ell \leq \log\binom{n}{\alpha n} - \log n$ then the size of the range of f is at most a $1/n$ fraction of the support of $D_{=1/2}$, and the result follows. Hence $\ell \geq \log\binom{n}{\alpha n} - \log n$. Fact 2.2 gives $|\log_2\binom{n}{\alpha n} - H(\alpha)n| \leq O(\log n)$, for n large. Hence, $\ell = \Theta(H(\alpha)n)$.

Now consider the bipartite graph with the n output nodes on one side and the ℓ input nodes on the other, where each output node is connected to the d input nodes it is a function of. Without loss of generality, each input node has degree at least 1 (otherwise, run this proof with ℓ the number of input bits actually used by f).

Claim 2.3. *There is a set I of $s := |I| \geq \Omega(H(\alpha)^2 n/d^2)$ input bits such that (i) each input bit in I has degree at most $b = O(d/H(\alpha))$, and (ii) each output bit is adjacent to at most one input bit in I.*

Proof: The average degree of an input node is dn/ℓ. By a Markov argument, at least $\ell/2$ input nodes have degree $\leq 2dn/\ell = O(d/H(\alpha))$. Let K be the set of these nodes. We obtain $I \subseteq K$ greedily as follows: Put a $v \in K$ in I, then remove from K any other input node adjacent to one of the outputs that v is adjacent to. Repeat until $K = \emptyset$.

Since each output node has degree d, for each node put in I we remove $\leq (d-1)\cdot O(d/H(\alpha)) = O(d^2/H(\alpha))$ others. So we collect at least $(\ell/2)/(1 + O(d^2/H(\alpha))) = \Omega(H(\alpha)^2 n/d^2)$. ∎

Let I by the set given by Claim 2.3, and without loss of generality let $I = [s] = \{1, 2, \ldots, s\}$. For an input node $u_i \in [s]$, let B_i be the set of output bits adjacent to u_i. Note $1 \leq |B_i| \leq O(d/H(\alpha))$ (the first inequality holds because input nodes have degree ≥ 1).

By dividing an input $u \in \{0,1\}^\ell$ in (x,y) where x are the first s input bits and y are the other $\ell - s$, and by permuting output bits, we rewrite f as

$$f(x,y) = h(y) \circ g_1(x_1,y) \circ g_2(x_2,y) \circ \cdots \circ g_s(x_s,y),$$

where g_i has range $\{0,1\}^{|B_i|}$.

Definition 2.4. We say that a function g_i is *y-fixed* if $g_i(0,y) = g_i(1,y)$, i.e., after fixing y it does not depend on x_i anymore.

For a string $z \in \{0,1\}^n$, we denote by z_{B_i} the projection of z on the bits of B_i, so that $f(x,y)_{B_2} = g_2(x_2,y)$, for example.

Definition of the statistical test.: The statistical test $T \subseteq \{0,1\}^n$ which will witness the claimed statistical distance is the union of three tests:

$$\begin{aligned}
T_F :=& \{z : \exists (x,y) : f(x,y) = z \\
&\quad \text{and } \geq 2n^\delta \text{ functions } g_i(x_i,y) \text{ are } y\text{-fixed}, i \in [s]\},\\
T_0 :=& \{z : z_{B_i} = 0^{|B_i|} \text{ for } \leq 3n^\delta \text{ indices } i \in [s]\},\\
T_S :=& \{z : \sum_i z_i \neq \alpha n\};\\
T :=& T_F \bigcup T_0 \bigcup T_S.
\end{aligned}$$

We now prove that the output of f is likely to pass the test, while a uniform string of weight αn is not.

Claim 2.5. $\Pr_u[f(u) \in T] \geq 1 - O(1/n^{\delta/2})$.

We recall the Littlewood-Offord anti-concentration inequality.

Fact 2.1 (Littlewood-Offord anti-concentration [23], [17])**.** *(Restated.)* Let $S_1, S_2, \ldots, S_t$ be t independent random variables, where S_i is uniform over $\{a_i, b_i\}$ for $a_i \neq b_i$. Then for any integer c, $Pr[\sum_i S_i = c] \leq O(1/\sqrt{t})$.

To prove this fact, reduce to the case $a_i \leq 0, b_i > 0$. Then generate $\sum S_i$ by first permuting variables, and then setting exactly the first S of them to the smallest values of their domains, where S is binomially distributed. Since for every permutation there is at most one value of S yielding sum c, and each value has probability $\leq O(1/\sqrt{t})$, the result follows.

Proof of Claim 2.5: Write again an input u to f as $u = (x,y)$. We prove that for every y we have $\Pr_x[f(x,y) \in$

$T] \geq 1 - O(1/n^{\delta/2})$, which implies the claimed bound. Fix any y.

If $\geq 2n^\delta$ functions $g_i(x_i, y)$ are y-fixed, then $\Pr_x[f(x,y) \in T_F] = 1$.

Also, if there are $\leq 3n^\delta$ indices $i \in [s]$ such that $g_i(x_i, y) = 0^{|B_i|}$ for some x_i, then clearly for any x the string $f(x,y)$ satisfies $f(x,y)_{B_i} = g_i(x_i, y) = 0^{|B_i|}$ for $\leq 3n^\delta$ indices i. In this case, $\Pr_x[f(x,y) \in T_0] = 1$.

Therefore, assume both that there are $\leq 2n^\delta$ functions $g_i(x_i, y)$ that are y-fixed, and that there are $\geq 3n^\delta$ indices i such that $g_i(x_i, y) = 0^{|B_i|}$ for some x_i. Consequently, there is a set $J \subseteq [s]$ of $\geq 3n^\delta - 2n^\delta = n^\delta$ indices i such that $g_i(x_i, y)$ is not y-fixed and $g_i(x_i, y) = 0^{|B_i|}$ for some $x_i \in \{0,1\}$. The key idea is that for the other value of $x_i \in \{0,1\}$ the value of $g_i(x_i, y)$ must have hamming weight bigger than 0, and therefore it contributes to anti-concentration.

Specifically, fix all bits in x except those in J, and denote the latter by x_J. We show that for any such fixing, the probability over the choice of the bits x_J that the output falls in T_S, i.e. $Pr_{x_J}[\sum_{k \leq n} f(x,y)_k \neq \alpha n]$, is at least $1 - O(1/n^{\delta/2})$. To see this, note that, for $i \in J$, the sum S_i of the bits in $g_i(x_i, y)$ (i.e., $S_i := \sum_{k \leq |B_i|} g_i(x_i, y)_k$) is 0 with probability $1/2$ over x_i and strictly bigger than 0 with probability $1/2$ (since $0^{|B_i|}$ is the only input with sum 0); moreover, the variables S_i are independent. Writing the sum of the bits in $f(x,y)$ as $a + \sum_{i \in J} S_i$ for some integer a which does not depend on x_J, we have

$$\begin{aligned}\Pr_{x_J \in \{0,1\}^{|J|}}[f(x,y) \neq \alpha n] &= \Pr_{x_J \in \{0,1\}^{|J|}}[\sum_{i \in J} S_i \neq \alpha n - a] \\ &\geq 1 - O(1/n^{\delta/2}),\end{aligned}$$

where the last inequality is by Fact 2.1. ∎

Claim 2.6. *Let $D = D_{=\alpha}$ be the uniform distribution over n-bit strings of hamming weight αn. Then $\Pr_D[D \in T] \leq 1/n$.*

The proof gives the stronger bound $\Pr_D[D \in T] \leq 1/2^{n^\gamma}$, for a $\gamma > 0$ depending on δ.

Proof of Claim 2.6: By a union bound,

$$\Pr_D[D \in T] \leq \Pr_D[D \in T_F] + \Pr_D[D \in T_0] + \Pr_D[D \in T_S].$$

We separately show that each term is at most $1/(3n)$.

First, $\Pr_D[D \in T_S] = 0$ by definition of D.

Also, $\Pr_D[D \in T_F] = |T_F| / \binom{n}{\alpha n}$. Note each string in T_F can be described by a string of $|y| + |x| - 2n^\delta$ bits, where the first $|y|$ are interpreted as a value for y, and the remaining $|x| - 2n^\delta$ are interpreted as values for the variables x_i corresponding to functions $g_i(x_i, y)$ that are not y-fixed. Hence,

$$|T_F| \leq 2^{|y|+|x|-2n^\delta} = 2^{\ell - 2n^\delta} \leq 2^{\log \binom{n}{\alpha n} - n^\delta},$$

and

$$\Pr_D[D \in T_F] \leq 2^{-n^\delta} \leq 1/(3n),$$

for large enough n.

Finally, we bound $\Pr_D[D \in T_0]$. There are several ways of doing this; the following is self-contained. For $i \in [s]$, let N_i be the event $D_{B_i} \neq 0^{|B_i|}$, over the choice of D. Let $t := 3n^\delta$ be as in the definition of T_0. We have:

$$\begin{aligned}\Pr_D[D \in T_0] \leq &\Pr[\exists J \subseteq [s], |J| = s - t, \\ &\quad \text{such that } N_i \text{ holds for all } i \in J] \\ \leq &\binom{s}{t} \max_{J \subseteq [s], |J| = s-t} \Pr[N_i \text{ for all } i \in J] \\ \leq &\binom{s}{t} \max_{J \subseteq [s], |J| = n/\log^2 n} \Pr[N_i \text{ for all } i \in J],\end{aligned} \tag{2}$$

where in the last inequality we use that $s - t = \Omega(H(\alpha)^2 n/d^2) - 3n^\delta \geq n/\log^2 n$ for $\delta < 1$, sufficiently small ϵ, and sufficiently large n, using that $d \leq \epsilon \log n$. Let

$$m := n/\log^2 n.$$

We now bound $\max_{J \subseteq [s], |J| = m} \Pr[N_i \text{ for all } i \in J]$. Without loss of generality, let the maximum be achieved for $J = \{1, 2, \ldots, m\}$. Write

$$\begin{aligned}&\Pr[N_i \text{ for all } i \leq m] = \\ &\Pr[N_1] \cdot \Pr[N_2 | N_1] \cdot \cdots \cdot \Pr[N_m | N_{m-1} \wedge \ldots \wedge N_1].\end{aligned} \tag{3}$$

We proceed by bounding $\Pr[N_k | N_{k-1} \wedge \ldots \wedge N_1]$ for any $k \leq m$. Recall that each set B_i has size $\leq b = O(d/H(\alpha))$. So the event $N_{k-1} \wedge \ldots \wedge N_1$ depends on $\leq (k-1)b$ bits. If we condition on any value of $(k-1)b$ bits, the probability that N_k is not true, i.e. that $D_{B_k} = 0^{|B_k|}$, is at least

$$\begin{aligned}\prod_{j=0}^{b-1} \frac{(1-\alpha)n - (k-1)b - j}{n - (k-1)b - j} &\geq \left(\frac{(1-\alpha)n - kb}{n}\right)^b \\ &\geq 1/3^b \geq 1/n^{O(\epsilon/H(\alpha))},\end{aligned}$$

using our initial assumption $\alpha \leq 1/2$, and that $k \leq m = n/\log^2 n$ and $b = O(d/H(\alpha)) = O(\epsilon \log n / H(\alpha))$, so $kb = o(n)$. Hence, $\Pr[N_k | N_{k-1} \wedge \ldots \wedge N_1] \leq 1 - 1/n^{O(\epsilon/H(\alpha))}$.

Plugging this bound in Equation (3), we obtain

$$\begin{aligned}\Pr[N_i \text{ for all } i \leq m] &\leq \left(1 - 1/n^{O(\epsilon/H(\alpha))}\right)^m \\ &\leq e^{-n^{1 - O(\epsilon/H(\alpha))}/\log^2 n} \leq e^{-n^{(1+\delta)/2}},\end{aligned}$$

for sufficiently small ϵ and large n (recall $\delta < 1$).

Plugging this bound back in Equation (2) we get

$$\Pr_D[D \in T_0] \leq (es/t)^t e^{-n^{(1+\delta)/2}} \leq n^{3n^\delta} e^{-n^{(1+\delta)/2}} \leq 1/(3n),$$

for large enough n. ∎

To conclude the proof of the theorem, note that the combination of the two claims gives $\Delta(f(U), D) \geq 1 - O(1/n^{\delta/2}) - 1/n = 1 - O(1/n^{\delta/2})$. □

This proof actually shows that for any $\tau > 0$ and $\delta < 1$, we can pick the same ϵ for any $\alpha \in (\tau, 1-\tau)$.

3. Lower bound for generating $(X, \text{majmod } X)$ locally

In this section we prove our lower bound for generating $(X, \text{majmod } X)$, restated next.

Theorem 1.4 (Lower bound for generating $(X, \text{majmod } X)$ locally)**.** *(Restated.)* For any $\delta < 1$ there is $\epsilon > 0$ such that for all sufficiently large n: Let $p \in [0.25 \log n, 0.5 \log n]$ be a prime number, and let $\text{majmod} : \{0,1\}^n \to \{0,1\}$ be defined as

$$\text{majmod}(x) = 1 \Leftrightarrow \sum_{i \leq n} x_i \mod p \in \{0, 1, \ldots, (p-1)/2\}.$$

Let $f : \{0,1\}^\ell \to \{0,1\}^{n+1}$ be an $(\epsilon \log n)$-local function where $\ell \leq n + n^\delta$.

Then $\Delta(f(U), (X, \text{majmod } X)) \geq 1/2 - O(1/\log n)$.

Intuition for the proof of Theorem 1.4.: The proof follows closely that of the lower bound for generating the "$= \alpha n$" distribution (Theorem 1.3). The main difference is that we use anti-concentration modulo p to argue that the number of ones in the input is uniform modulo p, and thus the output is correct with probability about $1/2$.

The problem in the proof of Theorem 1.3 that unfixed functions g_i can take two values with the same hamming weight translates here in the problem that g_i can take two values with the same weight modulo p. Locality is used to guarantee that the output length of g_i is smaller than p, and so if one of the two values of g_i is all zero the other one must be different modulo p.

3.1. Proof of Theorem 1.4

The beginning of the proof is like that of Theorem 1.3: we write (up to a permutation of the input and output bits):

$$f(x, y) = h(y) \circ g_1(x_1, y) \circ g_2(x_2, y) \circ \cdots \circ g_s(x_s, y),$$

where g_i has range $\{0,1\}^{|B_i|}$ (B_i denotes the output bits of g_i, so that $f(x,y)_{B_i} = g_i(x_i, y)$) for $1 \leq |B_i| \leq O(d)$, and $s \geq \Omega(n/d^2)$.

For notational simplicity, we assume that the last bit of f does not get permuted; so f_{n+1} is still the bit corresponding to majmod.

Definition of the statistical test.: Let

$$\begin{aligned}
T_F :=& \{z \in \{0,1\}^{n+1} : \exists (x,y) : f(x,y) = z \\
& \text{and } \geq 2n^\delta \text{ functions } g_i(x_i, y) \text{ are } y\text{-fixed}, i \in [s]\}, \\
T_0 :=& \{z : z_{B_i} = 0^{|B_i|} \text{ for } \leq 3n^\delta \text{ indices } i \in [s]\}, \\
T_S :=& \{(z', b) \in \{0,1\}^n \times \{0,1\} : \\
& \left(\sum_i z'_i \mod p \in \{0, 1, \ldots, (p-1)/2\} \right) \\
& \text{xor } (b = 1)\} \quad \text{(that is, } T_S = \text{"wrong answer");} \\
T :=& T_F \bigcup T_0 \bigcup T_S.
\end{aligned}$$

We now prove that the output of f passes the test with probability $1/2 - O(1/\log n)$, while $(X, \text{majmod}(X))$ passes the test with probability $1/n$.

Claim 3.1. $\Pr_u[f(u) \in T] \geq 1/2 - O(1/\log n)$.

The proof uses the following well-known fact, which can be thought of as an anti-concentration result for the sum of random variables modulo p.

Fact 3.2. *Let $a_1, a_2, \ldots, a_t$ be t integers not zero modulo p. The statistical distance between the distribution $\sum_{i \leq t} a_i x_i \mod p$ for uniform $x \in \{0,1\}^t$ and the uniform distribution over $\{0, 1, \ldots, p-1\}$ is at most $\sqrt{p} e^{-t/p^2}$.*

Proof using various results: By [7, Claim 33], the statistical distance is at most

$$\sqrt{p} \max_{a \neq 0} |E_{x \in \{0,1\}^t}[e(a \sum_{i \leq t} a_i x_i)] - E_{U_p}[e(aU_p)]|,$$

where $e(x) := e^{2\pi\sqrt{-1}x/p}$ and U_p is the uniform distribution over $\{0, 1, \ldots, p-1\}$. Fix any $a \neq 0$. By [24, Lemma 12] $|E_{x \in \{0,1\}^t}[e(a \sum_{i \leq t} a_i x_i)] \leq e^{-t/p^2}$; also, $E_{U_p}[e(aU_p)] = 0$. ■

Proof of Claim 3.1: Write again an input u to f as $u = (x, y)$. We prove that for every y we have $\Pr_x[f(x,y) \in T] \geq 1/2 - O(1/\log n)$, which implies the claimed bound. Fix any y.

If $\geq 2n^\delta$ functions $g_i(x_i, y)$ are y-fixed, then $\Pr_x[f(x,y) \in T_F] = 1$.

Also, if there are $\leq 3n^\delta$ indices $i \in [s]$ such that $g_i(x_i, y) = 0^{|B_i|}$ for some x_i, then clearly for any x the string $f(x,y)$ satisfies $f(x,y)_{B_i} = g_i(x_i, y) = 0^{|B_i|}$ for $\leq 3n^\delta$ indices i. In this case, $\Pr_x[f(x,y) \in T_0] = 1$.

Therefore, assume both that there are $\leq 2n^\delta$ functions $g_i(x_i, y)$ that are y-fixed, and that there are $\geq 3n^\delta$ indices i such that $g_i(x_i, y) = 0^{|B_i|}$ for some x_i. Consequently, there is a set $J \subseteq [s]$ of $\geq 3n^\delta - 2n^\delta = n^\delta$ indices i such that $g_i(x_i, y)$ is not y-fixed and $g_i(x_i, y) = 0^{|B_i|}$ for some $x_i \in \{0,1\}$. The key idea is that for the other value of $x_i \in \{0,1\}$ the value of $g_i(x_i, y)$ must have hamming weight different from 0 modulo p, and therefore it contributes to anti-concentration.

Specifically, note that g_s is the only function that may affect the output bit f_{n+1}, corresponding to majmod. If present, remove s from J. Fix all bits in x except those in J, and denote the latter by x_J. We show that for any such fixing, the probability over the choice of the bits x_J that the output falls in T_S is $\geq 1/2 - O(1/\log n)$. To see this, note that, for $i \in J$, the sum S_i of the bits in $g_i(x_i, y)$ (i.e., $S_i := \sum_{k \leq |B_i|} g_i(x_i, y)_k$) is 0 with probability $1/2$ over x_i, and $a_i \neq 0 \mod p$ with probability $1/2$. This is because the maximum sum is

$$|B_i| = O(d) = O(\epsilon \log n) < p$$

for sufficiently small ϵ. Moreover, the variables S_i are independent. Writing the sum of the first n bits of $f(x, y)$ as $a + \sum_{i \in J} S_i$ for some integer a which does not depend on x_J, we have by Fact 3.2 that, over the choice of x_J, the statistical distance between the sum of the first n bits of f and the uniform distribution U_p over $\{0, 1, \ldots, p-1\}$ is at most

$$\sqrt{p} e^{-(n^\delta - 1)/p^2} \leq 1/n,$$

since $p = O(\log n)$. Because the last bit $b := f_{n+1}(x, y)$ is fixed (independent from x_J), and

$$\begin{aligned} \Pr_{U_p}[U_p \in \{0, 1, \ldots, (p-1)/2\}] &= 1/2 - 1/(2p) \\ &= 1/2 - \Theta(1/\log n), \quad (4) \end{aligned}$$

we have

$$\begin{aligned} \Pr_{x_J}[f(x, y) \in T_S] &\geq 1/2 - O(1/\log n) - 1/n \\ &\geq 1/2 - O(1/\log n). \quad (5) \end{aligned}$$

■

Claim 3.3. *Let $D = (X, \text{majmod}(X))$ for uniform $X \in \{0,1\}^n$. Then $\Pr_D[D \in T] \leq 1/n$.*

The proof gives a stronger, exponential bound.

Proof of Claim 3.3: By a union bound,

$$\Pr_D[D \in T] \leq \Pr_D[D \in T_F] + \Pr_D[D \in T_0] + \Pr_D[D \in T_S].$$

We separately show that each term is at most $1/(3n)$.

First, $\Pr_D[D \in T_S] = 0$ by definition of D.

Also, $\Pr_D[D \in T_F] = |T_F|/2^n$. Note each string in T_F can be described by a string of $|y| + |x| - 2n^\delta$ bits, where the first $|y|$ are interpreted as a value for y, and the remaining $|x| - 2n^\delta$ are interpreted as values for the variables x_i corresponding to functions $g_i(x_i, y)$ that are not y-fixed. Hence,

$$|T_F| \leq 2^{|y|+|x|-2n^\delta} = 2^{\ell - 2n^\delta} \leq 2^{n - n^\delta},$$

and

$$\Pr_D[D \in T_F] \leq 2^{-n^\delta} \leq 1/(3n),$$

for large enough n.

Finally, we bound $\Pr_D[D \in T_0]$. For any $i \in [s]$,

$$\Pr_{X \in \{0,1\}^n}[X_{B_i} = 0^{|B_i|}] = 1/2^{|B_i|} = 1/2^{O(d)} = 1/n^{O(\epsilon)}.$$

Moreover, these events are independent for different i. Hence, recalling that $s = \Omega(n/d^2) \geq n/\log^2 n$, we have:

$$\begin{aligned} \Pr_D[D \in T_0] &\leq \binom{s}{3n^\delta} (1 - 1/n^{O(\epsilon)})^{s - 3n^\delta} \\ &\leq n^{3n^\delta} e^{-n^{1-O(\epsilon)}/\log^2 n} \leq 1/(3n) \end{aligned}$$

for a sufficiently small ϵ and large enough n. ■

To conclude the proof of the theorem, note that the combination of the two claims gives $\Delta(f(U), (X, \text{majmod } X)) \geq 1/2 - O(1/\log n) - 1/n = 1/2 - O(1/\log n)$.

4. Lower bounds for generating by decision trees

In this section we prove Theorem 1.6, (1). We make use of the following lemma (the Pailey-Zygmund inequality could be used instead, see [38]).

Lemma 4.1 ([14])**.** *There is a constant k such that for large enough n and any k-wise independent distribution $X \in \{0,1\}^n$, with probability ≥ 0.49 the variable X has strictly less than $n/2$ ones.*

Proof of Theorem 1.6, (1): Let k be the constant from Lemma 4.1. Suppose the distribution $X := f(U)$ is k-wise independent. Then by Lemma 4.1 $\Pr[\sum_i X_i < n/2] \geq 0.49$. The statistical test which checks if the output bits sum to $n/2$ proves the claim in this case.

Otherwise, there are k output bits of f that are not uniformly distributed over $\{0,1\}^k$. We claim that, for any y, the probability k output bits evaluate to y equals $A/2^{kd}$ for an integer A. To see this, note that the k output bits can be computed with a decision tree of depth dk (e.g., use the decision tree for the first bit, then use the decision tree for the second, and so on). Since the probability of outputting a value y in a decision tree is the sum over all leaves labeled with y of the probabilities of reaching that leaf, and each leaf has probability $a/2^{kd}$ for some integer a, the result follows.

Therefore, if these k bits are not uniform, there there must be an output value that has probability at least $1/2^k + 1/2^{kd}$.

But over $D_{=1/2}$, this output combination of the k bits has probability at most

$$\frac{1}{2} \cdot \frac{n/2}{n-1} \cdot \cdots \cdot \frac{n/2}{n-(k-1)} \quad \leq \quad \frac{1}{2^k} + O(1/n).$$

So, checking if these k bits equal y we get statistical distance $\geq 1/2^{O(d)} - O(1/n)$. ■

Acknowledgments.: I am very grateful to Rajmohan Rajaraman and Ravi Sundaram for extensive collaboration on this project. I also thank Artur Czumaj and Shachar Lovett for useful discussions, and the anonymous referees for helpful feedback.

REFERENCES

[1] N. Alon, L. Babai, and A. Itai, "A fast and simple randomized algorithm for the maximal independent set problem," *Journal of Algorithms*, vol. 7, pp. 567–583, 1986.

[2] B. Applebaum, Y. Ishai, and E. Kushilevitz, "Cryptography in NC^0," *SIAM J. Comput.*, vol. 36, no. 4, pp. 845–888, 2006.

[3] ——, "On pseudorandom generators with linear stretch in nc^0," *Computational Complexity*, vol. 17, no. 1, pp. 38–69, 2008.

[4] L. Babai, "Random oracles separate PSPACE from the polynomial-time hierarchy," *Inform. Process. Lett.*, vol. 26, no. 1, pp. 51–53, 1987.

[5] R. Beigel, "The polynomial method in circuit complexity," in *8th Annual Structure in Complexity Theory Conference*. IEEE, 1993, pp. 82–95.

[6] M. Blum and S. Micali, "How to generate cryptographically strong sequences of pseudo-random bits," *SIAM J. on Computing*, vol. 13, no. 4, pp. 850–864, Nov. 1984.

[7] A. Bogdanov and E. Viola, "Pseudorandom bits for polynomials," *SIAM Journal on Computing*, vol. 39, no. 6, pp. 2464–2486, 2010.

[8] R. Boppana and J. Lagarias, "One-way functions and circuit complexity," *Inform. and Comput.*, vol. 74, no. 3, pp. 226–240, 1987.

[9] M. Braverman, "Poly-logarithmic independence fools AC^0 circuits," in *24th Conference on Computational Complexity (CCC)*. IEEE, 2009.

[10] H. Buhrman, P. B. Miltersen, J. Radhakrishnan, and V. Srinivasan, "Are bitvectors optimal?" *SIAM J. Comput.*, vol. 31, no. 6, pp. 1723–1744, 2002.

[11] B. Chor and O. Goldreich, "On the power of two-point based sampling," *Journal of Complexity*, vol. 5, no. 1, pp. 96–106, 1989.

[12] T. Cover and J. Thomas, *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.

[13] A. Czumaj, P. Kanarek, M. Kutylowski, and K. Lorys, "Delayed path coupling and generating random permutations via distributed stochastic processes," in *Symposium on Discrete Algorithms (SODA)*, 1999, pp. 271–280.

[14] I. Diakonikolas, P. Gopalan, R. Jaiswal, R. A. Servedio, and E. Viola, "Bounded independence fools halfspaces," in *50th Symposium on Foundations of Computer Science (FOCS)*, 2009.

[15] Y. Dodis, M. Pătraşcu, and M. Thorup, "Changing base without losing space," in *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, 2010.

[16] B. Dubrov and Y. Ishai, "On the randomness complexity of efficient sampling," in *38th Annual ACM Symposium on Theory of Computing (STOC)*, 2006, pp. 711–720.

[17] P. Erdős, "On a lemma of Littlewood and offord," *Bull. Amer. Math. Soc.*, vol. 51, pp. 898–902, 1945.

[18] V. Guruswami, C. Umans, and S. P. Vadhan, "Unbalanced expanders and randomness extractors from parvaresh–vardy codes," *J. ACM*, vol. 56, no. 4, 2009.

[19] D. Gutfreund and E. Viola, "Fooling parity tests with parity gates," in *8th International Workshop on Randomization and Computation (RANDOM)*. Springer, 2004, pp. 381–392.

[20] T. Hagerup, "Fast parallel generation of random permutations," in *18th Colloquium on Automata, Languages and Programming (ICALP)*, 1991, pp. 405–416.

[21] J. Håstad, *Computational limitations of small-depth circuits*. MIT Press, 1987.

[22] R. Impagliazzo and M. Naor, "Efficient cryptographic schemes provably as secure as subset sum," *Journal of Cryptology*, vol. 9, no. 4, pp. 199–216, Fall 1996. [Online]. Available: citeseer.nj.nec.com/impagliazzo96efficient.html

[23] J. Littlewood and A. Offord, "On the number of real roots of a random algebraic equation," *III. Rec. Math. [Mat. Sbornik] N.S.*, vol. 12, pp. 277–286, 1943.

[24] S. Lovett, O. Reingold, L. Trevisan, and S. P. Vadhan, "Pseudorandom bit generators that fool modular sums," in *APPROX-RANDOM*, 2009, pp. 615–630.

[25] S. Lovett and E. Viola, "Bounded-depth circuits cannot sample good codes," 2010, manuscript.

[26] M. Luby, B. Veličković, and A. Wigderson, "Deterministic approximate counting of depth-2 circuits," in *2nd Israeli Symposium on Theoretical Computer Science (ISTCS)*, 1993, pp. 18–24.

[27] Y. Matias and U. Vishkin, "Converting high probability into nearly-constant time-with applications to parallel hashing," in *23rd ACM Symposium on Theory of Computing (STOC)*, 1991, pp. 307–316.

[28] E. Mossel, A. Shpilka, and L. Trevisan, "On epsilon-biased generators in NC^0," *Random Struct. Algorithms*, vol. 29, no. 1, pp. 56–81, 2006.

[29] N. Nisan, "Pseudorandom bits for constant depth circuits," *Combinatorica*, vol. 11, no. 1, pp. 63–70, 1991.

[30] N. Nisan and A. Wigderson, "Hardness vs randomness," *J. Computer & Systems Sciences*, vol. 49, no. 2, pp. 149–167, 1994.

[31] R. Pagh, "Low redundancy in static dictionaries with constant query time," *SIAM J. Comput.*, vol. 31, no. 2, pp. 353–363, 2001.

[32] ——, "On the cell probe complexity of membership and perfect hashing," in *33rd Annual Symposium on Theory of Computing (STOC)*. ACM, 2001, pp. 425–432.

[33] M. Pătraşcu, "Succincter," in *49th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2008.

[34] M. Pătraşcu and E. Viola, "Cell-probe lower bounds for succinct partial sums," in *21th Symposium on Discrete Algorithms (SODA)*, 2010.

[35] E. Viola, "On constructing parallel pseudorandom generators from one-way functions," in *20th Annual Conference on Computational Complexity (CCC)*. IEEE, 2005, pp. 183–197.

[36] ——, "Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates," *SIAM Journal on Computing*, vol. 36, no. 5, pp. 1387–1403, 2007. [Online]. Available: http://link.aip.org/link/?SMJ/36/1387/1

[37] ——, "Bit-probe lower bounds for succinct data structures," in *41th Annual Symposium on the Theory of Computing (STOC)*. ACM, 2009.

[38] ——, "The complexity of distributions," in *51th Symposium on Foundations of Computer Science (FOCS)*, 2010, preliminary version titled "Are all distributions easy?" (2009).

[39] A. Yao, "Theory and applications of trapdoor functions," in *23rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1982, pp. 80–91.