Deterministic Parallel Algorithms for Fooling Polylogarithmic Juntas and the Lovász Local Lemma

DAVID G. HARRIS, Department of Computer Science, University of Maryland

Many randomized algorithms can be derandomized efficiently using either the method of conditional expectations or probability spaces with low (almost-) independence. A series of papers, beginning with Luby (1993) and continuing with Berger and Rompel (1991) and Chari et al. (2000), showed that these techniques can be combined to give deterministic parallel algorithms for combinatorial optimization problems involving sums of w -juntas. We improve these algorithms through derandomized variable partitioning, reducing the processor complexity to essentially independent of w and time complexity to linear in w .

As a key subroutine, we give a new algorithm to generate a probability space which can fool a given set of neighborhoods. Schulman (1992) gave an NC algorithm to do so for neighborhoods of size $w \leq O(\log n)$. Our new algorithm is in NC¹, with essentially optimal time and processor complexity, when $w = O(\log n)$; it remains in NC up to $w = \text{polylog}(n)$. This answers an open problem of Schulman.

One major application of these algorithms is an NC algorithm for the Lovász Local Lemma. Previous NC algorithms, including the seminal algorithm of Moser and Tardos (2010) and the work of Chandrasekaran et al. (2013), required that (essentially) the bad-events could span only $O(\log n)$ variables; we relax this to $\text{polylog}(n)$ variables. We use this for an NC² algorithm for defective vertex coloring, which works for arbitrary degree graphs.

CCS Concepts: • **Theory of computation** → **Pseudorandomness and derandomization**; • **Computing methodologies** → *Combinatorial algorithms*;

Additional Key Words and Phrases: Derandomization, NC algorithms, method of conditional expectations, Lovász local lemma

ACM Reference format:

David G. Harris. 2018. Deterministic Parallel Algorithms for Fooling Polylogarithmic Juntas and the Lovász Local Lemma. *ACM Trans. Algorithms* 14, 4, Article 47 (August 2018), 24 pages.
<https://doi.org/10.1145/3230651>

1 INTRODUCTION

Many algorithms can be formulated as optimization problems, in which we seek to maximize or minimize a function of the form $S(x) = \sum_j f_j(x)$ over $x \in \{0, 1\}^n$; we refer to the summands f_j as *objective functions*. These may correspond to a scoring function measuring solution quality, or they might be indicators for bad events to avoid. We will consider cases in which each f_j depends on at most w coordinates of x ; this is known as a w -*junta*.

This is an extended version of a paper appearing in the Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA) 2017. Research supported in part by NSF Awards CNS-1010789 and CCF-1422569.

Authors' address: D. G. Harris, Department of Computer Science, University of Maryland, A.V. Williams Building, 8223 Paint Branch Drive, College Park, MD 20742; email: davidgharris29@gmail.com.

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.

© 2018 ACM 1549-6325/2018/08-ART47 \$15.00

<https://doi.org/10.1145/3230651>

This often leads to randomized algorithms with the following structure: if X is drawn from a suitable distribution (say independent fair coins), then $E[S(X)] = \sum_j E[f_j(X)] = S_0$. Obviously, there exists some $x \in \{0, 1\}^n$ with the property $S(x) \geq S_0$. Usually we can find such an x with a randomized algorithm, since a “typical” vector x has the property $S(x) \approx S_0$.

A key derandomization problem is thus to find such x deterministically. There are two main paradigms to do so: conditional expectations and low-independence probability spaces. To use conditional expectations, we gradually assign the bits of X to 0 or 1, ensuring that at each step the conditional expectation $E[S(X)]$ increases. To solve this by low-independence, we draw the random variables X from a probability space which has the same w -wise marginal distributions as the independent space $\{0, 1\}^n$. Each of these methods has disadvantages. The method of conditional expectations is inherently sequential: decisions about some x_i depend on the assignment of $x_1, \dots, x_{i-1}$. The method of low-independence can easily be parallelized, but leads to large processor counts as each element of the probability space requires a separate processor.

A hybrid approach was proposed by Luby in [11] and extended by Berger and Rompel [3], which combines parallelism with low processor complexity. Their key observation is that probability spaces with polylog-wise-independence can be described as linear codes over $GF(2)$, of length $\text{polylog}(n)$. The method of conditional expectations can be applied to the code itself, not directly to the solution vector. The main limitation of this hybrid algorithm is that, at least in its simplest form, it has processor complexity which is *exponential* in w . Berger and Rompel describe a limited number of problem-specific techniques to overcome this. In this article, we will investigate a more general method of dealing with this computational bottleneck, based on a derandomization of random variable partitioning.

1.1 Alternate Derandomization Approaches

We mention three other general approaches to derandomization, and the ways in which they fail to cover some key applications. The first approach is to use a probability space which is ϵ -approximately w -wise-independent (see Definition 3.4). Such a space is significantly smaller than a fully independent space. If the objective functions f_j were simply monomial functions, or more generally had small decision tree complexity, then their expectation would differ only slightly between an ϵ -approximately independent and a fully independent space. However, in many applications, f_j may be significantly more complex and the overall bias can become as large as $2^w \epsilon$ — requiring ϵ to be super-polynomially small for $w = \omega(\log n)$, and requiring the probability space to be too large to explore exhaustively.

A related approach is one of Schulman [15] for generating a probability space which fools a given list of neighborhoods (see Section 2.1 for a formal definition). If Ω fools the neighborhoods corresponding to each f_j , then there is guaranteed to exist some $x \in \Omega$ with $S(x) \geq S_0$; if Ω has small support, then this leads to an efficient algorithm. Although the space Ω can be significantly smaller than a fully w -wise-independent space, it is still super-polynomial for $w = \omega(\log n)$ and so this approach does not give NC algorithms. Fooling neighborhoods will nonetheless be a key building block of our algorithms.

Finally, the derandomization technique of Sivakumar [16] can be applied when the functions f_j are computed via automata on a polynomial-sized state-space. One can build a relatively small probability distribution which fools a polynomial number of such automata. This class of functions, while quite general, nevertheless omits some key applications, most notably rainbow hypergraph coloring (as we discuss later). Another disadvantage of Sivakumar’s method is its high processor complexity (on the order of n^{20} or more).

1.2 Our Contributions and Overview

In Section 2, we present a new algorithm to produce probability spaces fooling a list of neighborhood or a list of Fourier characters over $GF(2)$ (a closely related problem). These are important subroutines needed for the algorithmic approach of Berger and Rompel [3]. The algorithm we develop has significantly lower complexity than previous algorithms. In particular, when the neighborhood size w is $w = O(\log n)$, then we obtain an NC^1 algorithm and when $w = \text{polylog}(n)$ we obtain an NC algorithm.

In addition to their use in the Berger-Rempel framework, these algorithms can be used for some other derandomization problems. For instance, we obtain a near-optimal algorithm to find a codeword of Hamming weight at least $L/2$ in a length- L binary code, a toy derandomization problem introduced by [14].

In Section 3, we consider fooling sums of w -juntas. As we have discussed, the main bottleneck in the Berger-Rempel algorithm [3] is the exponential processor dependency on w . We give an algorithm based on random variable partitioning, which is then derandomized. This approach makes the processor complexity independent of w while giving a *linear* time dependency on w . This allows us to handle, for the first time in NC , many applications with $w = \text{polylog}(n)$.

We describe a sample application in Section 3.5 to rainbow hypergraph coloring. Given a d -uniform hypergraph on m edges, we wish to d -color the vertices so that at least $md!/d^d$ edges see all d colors (as is expected in a uniform random coloring). This was an example application given in [3]. Although [3] did not provide concrete complexity bounds, their algorithm appears to require $O(\log^4 m)$ time and $O(m^{1+\log 2})$ processors. We reduce this to roughly $\tilde{O}(\log^2 m)$ time and $O(m)$ processors. This illustrates how our derandomization procedure has been optimized for processor and time complexity, so that it can be beneficial even for applications with prior NC algorithms.

In Section 4, we consider the seminal Moser-Tardos algorithm for the Lovász Local Lemma [12]. In this setting, one seeks to avoid a set of “bad events,” which are Boolean functions of the variables $x_1, \dots, x_n$. There have been some NC versions of this algorithm, appearing in the original paper of Moser and Tardos along with some extensions in [5, 9]. These algorithms are somewhat limited in the types of problems they can handle, with restrictive conditions on the decision-tree-complexity of the bad-events. We greatly expand the scope of these algorithms to give NC algorithms in almost any application where the bad-events are w -juntas for $w = \text{polylog}(n)$.

In Section 5, we apply our LLL derandomization to two graph theory applications. First, *defective vertex coloring*: given a graph G of maximum degree Δ , we achieve an $\tilde{O}(\log^2 n)$ -time algorithm for a k -defective vertex coloring with $c = O(\Delta/k)$ colors. Notably, although our general LLL algorithm only applies to bad-events which span a polylogarithmic number of variables (in particular here, when $\Delta \leq \text{polylog}(n)$), our coloring algorithm works for arbitrary values of Δ and k . The second application is to domatic partition; here we only get an NC algorithm for graphs of degree $k = \text{polylog}(n)$.

1.3 Notation and Conventions

All our algorithms will be described in the deterministic EREW PRAM model. In this model, we say an algorithm $\mathcal{A}$ has *complexity* (C_1, C_2) if it runs in C_1 time and C_2 processors. In order to focus on the leading-order terms, we often use a looser metric which we refer to as *quasi-complexity*. We say $\mathcal{A}$ has *quasi-complexity* (C_1, C_2) if it has complexity $(\tilde{O}(C_1), C_2^{1+o(1)})$, where we define $\tilde{O}(t) = t(\log t)^{O(1)}$.¹

¹It is very difficult to obtain estimates which are finer than this; small changes in the computational model or the input data (e.g., the register size, the precisions of the real-valued weights, or the atomic arithmetic operations) can change the runtime by hard-to-track polyloglog factors.

We let $[n]$ denote the set $\{1, \dots, n\}$. For any collection of sets $E \subseteq 2^{[n]}$, we define the *width* of E as $\text{width}(E) = \max_{e \in E} |e|$. For a probability space Ω , we use the notation $|\Omega|$ to mean the cardinality of the support of Ω . Given a set X , we write $x \sim X$ to mean that x is drawn from the uniform probability distribution on X . For a Boolean predicate P , we use the Iverson notation where $[P]$ is 1 if P is true and zero otherwise.

We write $GF(q)$ for the finite field with q elements. In particular, the field $GF(2^s)$ can be represented as s -bit binary vectors, and addition in the field is taken mod 2, the same as coordinatewise XOR. We write this addition operation as $\oplus$.

Throughout, $\log x$ refers to the natural logarithm and $\log_2 x$ to the base-two logarithm.

2 FOOLING NEIGHBORHOODS

2.1 Fourier Characters, Neighborhoods, and Codes

Many probability spaces satisfying (approximate) independence conditions are built on top of codes over $GF(2)$. These are closely related to Fourier characters over $GF(2)$. We begin by reviewing some definitions and basic results.

Definition 2.1. A Fourier character over $GF(2)$ is a function $\chi_e : GF(2)^n \rightarrow \{-1, 1\}$ defined by $\chi_e(x) = (-1)^{\sum_{i \in e} x_i}$, for some $e \subseteq [n]$.

We say that a probability space Ω is *unbiased* for e if $\mathbb{E}_{X \sim \Omega}[\chi_e(X)] = \mathbb{E}_{X \sim \{0,1\}^n}[\chi_e(X)]$. This condition trivially holds for $e = \emptyset$ (in which case $\chi_e(X) = 1$ with probability one). For $e \neq \emptyset$, we have $\mathbb{E}_{X \sim \{0,1\}^n}[\chi_e(X)] = 0$ and so the condition is that $\mathbb{E}_{X \sim \Omega}[\chi_e(X)] = 0$ as well.

We say Ω *fools* e if Ω is unbiased for every subset $f \subseteq e$. In this context, the set e is referred to as a *neighborhood* and this condition is also referred to as *fooling neighborhood* e . Likewise, we say Ω is unbiased for (fools, respectively) a list $E = \{e_1, \dots, e_m\}$ if Ω is unbiased for (fools, respectively) each $e_1, \dots, e_m$.

Our notation and definitions will differ slightly from the standard use in coding theory. Given a list of vectors $A(1), \dots, A(n) \in GF(2)^L$, we refer to the list $A = A(1), \dots, A(n)$ as a *code of length L and size n* . We use the following notational shortcut throughout: if $x \in \{0, 1\}^n$ and $e \subseteq [n]$ is a set, then we define $x(e) = \bigoplus_{i \in e} x_i$. If A is a collection of n binary vectors and $e \subseteq [n]$, then $A(e)$ is the binary vector defined coordinatewise by $\bigoplus_{i \in e} A(i)$.

PROPOSITION 2.2. For any Boolean function $g : \{0, 1\}^n \rightarrow \mathbb{R}$, there are weights $\gamma_e \in \mathbb{R}$, where e ranges over all 2^n subsets of $[n]$, such that for all $x \in \{0, 1\}^n$ we have

$$g(x) = \sum_{e \subseteq [n]} \gamma_e \chi_e(x).$$

The weights γ can be determined with quasi-complexity $(n, 2^n)$.

PROOF. This is the Discrete Fourier Transform over $GF(2)$. Define $\gamma_e = 2^{-n} \sum_{y \in \{0,1\}^n} \chi_e(y)g(y)$; these weights can be computed efficiently using the well-known Fast Walsh-Hadamard Transform algorithm. $\square$

PROPOSITION 2.3. If a probability space Ω fools the neighborhood $e \subseteq [n]$, then for all $z \in \{0, 1\}^n$ we have

$$P_{X \sim \Omega} \left(\bigwedge_{i \in e} x_i = z_i \right) = 2^{-|e|}.$$

PROOF. Let $g(x) = [\bigwedge_{i \in e} x_i = z_i]$. By Proposition 2.2, there exist weights γ_f , where f ranges over subsets of e , such that $g(x) = \sum_{f \subseteq e} \gamma_f \chi_f(x)$. Then

$$\mathbb{E}_{X \sim \Omega}[g(X)] = \sum_{f \subseteq e} \gamma_f \mathbb{E}_{X \sim \Omega}[\chi_f(X)] = \sum_{f \subseteq e} \gamma_f \mathbb{E}_{X \sim \{0,1\}^n}[\chi_f(X)] = \mathbb{E}_{X \sim \{0,1\}^n}[g(X)] = 2^{-|e|}.$$

□

The main connection between codes, Fourier characters, and fooling neighborhoods comes from the following construction.

Definition 2.4. Given a code A of length L , define the probability space Ω^A as follows: draw a vector $y \sim GF(2)^L$, and set $X_i = A(i) \cdot y$ for $i = 1, \dots, n$, where $\cdot$ is the inner product over $GF(2)^L$. Note that $|\Omega^A| = 2^L$.

Definition 2.5 (E-unbiased code). The code A is an *E-unbiased code* if $A(e) \neq \vec{0}$ for all non-empty sets $e \in E$.

PROPOSITION 2.6. If A is an *E-unbiased code*, then Ω^A is unbiased for every $e \in E$.

PROOF. Let $e \in E$ with $e \neq \emptyset$. We have

$$\mathbb{E}_{X \sim \Omega^A}[\chi_e(X)] = 2^{-L} \sum_{y \in GF(2)^L} (-1)^{\sum_{i \in e} A(i) \cdot y} = 2^{-L} \sum_{y \in GF(2)^L} (-1)^{A(e) \cdot y} = 0. \quad \square$$

2.2 Unbiased Codes and Fooling Neighborhoods

We begin with an algorithm to construct a code which is *E-unbiased* for a given set $E \subseteq 2^{[n]}$; we later extend this to fool neighborhoods. This algorithm has two phases: first, we show how to find a code which is unbiased for *most* of a given set E ; we then bootstrap this to be unbiased on all of E .

We begin with a simple result about multivariate polynomials over a finite field.

PROPOSITION 2.7. Let $p(z_1, \dots, z_k)$ be a non-zero polynomial over $GF(2^s)$, with degree at most d in each variable separately. For any $\alpha \in GF(2^s)$, note that $p(\alpha, z_2, \dots, z_k)$ is a $k-1$ variable polynomial over $GF(2^s)$. If $\alpha \sim GF(2^s)$, then $p(\alpha, z_2, \dots, z_k) \equiv 0$ with probability at most $d/2^s$.

PROOF. Factor p as

$$p(z_1, \dots, z_k) = \sum_{t_2, \dots, t_k} q_{t_2, \dots, t_k}(z_1) z_2^{t_2} \dots z_k^{t_k},$$

where $t_2, \dots, t_k$ range over non-negative integers. Each such polynomial q has degree d , and they are not all zero (else $p \equiv 0$). Let $t_2, \dots, t_k$ be such that $q_{t_2, \dots, t_k} \neq 0$; with probability at most $d/2^s$ we have $q_{t_2, \dots, t_k}(\alpha) = 0$. But if $q_{t_2, \dots, t_k}(\alpha) \neq 0$, then $p(\alpha, z_2, \dots, z_k)$ has a non-zero coefficient of $z_2^{t_2} \dots z_k^{t_k}$, hence $p(\alpha, z_2, \dots, z_k) \not\equiv 0$. □

PROPOSITION 2.8. Let $E \subseteq 2^{[n]}$. Given integer parameters $k \geq 1, s \geq 0$, there is an algorithm to construct a code A of length s , such that at most $kn^{1/k}2^{-s}|E|$ sets $e \in E$ have $A(e) = 0$. This procedure has quasi-complexity $(k(s + \log(mn)), 2^s W)$, where $m = |E|$ and $W = n + \sum_{e \in E} |e|$.

PROOF. Let Z be the set of formal monomials of the form $z_1^{u_1} \dots z_k^{u_k}$ in the ring $GF(2^s)[z_1, \dots, z_k]$, where $u_1, \dots, u_k \in \{0, \dots, d\}$ and $d = \lceil n^{1/k} - 1 \rceil$. Enumerate Z (in some arbitrary order) as $\mu_1, \dots, \mu_\ell$ where $\ell \geq n$.

For $i \in [n]$, we will define $A(i)$ to be the binary representation of $\mu_i(\alpha_1, \dots, \alpha_k)$, where $\alpha_1, \dots, \alpha_k$ will be chosen suitably from $GF(2^s)$. For any $e \subseteq [n]$, define the polynomial $\mu_e = \sum_{i \in e} \mu_i$. By linearity, $A(e)$ is the binary representation of $\mu_e(\alpha_1, \dots, \alpha_k)$. So we need to select $\alpha_1, \dots, \alpha_k$ so that

there are few sets $e \in E$ with $\mu_e(\alpha_1, \dots, \alpha_k) = 0$. We select $\alpha_1, \dots, \alpha_k$ sequentially, according to the following rule. For $i = 1, \dots, k+1$ let us define

$$E_i = \{e \in E \mid \mu_e(\alpha_1, \alpha_2, \dots, \alpha_{i-1}, z_i, z_{i+1}, \dots, z_k) \neq 0\}.$$

By Proposition 2.7, if α_i is chosen uniformly at random, then in expectation at most $(d/2^s)|E_i|$ sets $e \in E_i$ satisfy $\mu_e(\alpha_1, \alpha_2, \dots, \alpha_i, z_{i+1}, \dots, z_k) \equiv 0$. By enumerating over all possible values of $\alpha_i \in GF(2^s)$ to maximize the size of $|E_{i+1}|$, we ensure that

$$|E_{i+1}| \geq (1 - d/2^s)|E_i|.$$

As $E_1 = E$, at the end of this process we have $|E_{k+1}| \geq |E|(1 - d/2^s)^k \geq |E|(1 - kn^{1/k}/2^s)$, as required. This procedure requires k separate stages. In each stage, we must count $|E_{i+1}|$ for every choice of $\alpha_i \in GF(2^s)$, which requires quasi-complexity $(s + \log mn, 2^s W)$. $\square$

THEOREM 2.9. *Let $E \subseteq 2^{[n]}$. There is an algorithm with quasi-complexity $(\log mn, W)$ to find an E -unbiased code of length $L = \log_2 m + O(\frac{\log m}{\log \log \log m})$, where $m = |E|$ and $W = n + \sum_{e \in E} |e|$.*

PROOF. We will first discuss our algorithm under the assumption that $m \geq n$.

We form the code A by concatenating r separate codes $A_1, \dots, A_r$, each of length s , i.e.,

$$A(\ell) = \left(A_1(\ell)(1), \dots, A_1(\ell)(s), A_2(\ell)(1), \dots, A_2(\ell)(s), \dots, A_r(\ell)(1), \dots, A_r(\ell)(s) \right).$$

The resulting code A has length $L = rs$. We form $A_1, \dots, A_r$ sequentially, according to the following rule. For $i = 1, \dots, r+1$ define

$$E_i = \{e \in E \mid A_1(e) = \dots = A_{i-1}(e) = 0\}.$$

Note that $E_1 = E$. Select each A_i in turn by applying Proposition 2.8 to the set E_i , so that $|E_{i+1}| \leq \epsilon |E_i|$ for $\epsilon = kn^{1/k}2^{-s}$, where we set $k = \lceil \log \log m \rceil$, $s = \lceil \frac{\log m}{\log \log \log m} \rceil$, and $r = \lceil 1 + \frac{\log m}{\log(1/\epsilon)} \rceil$.

At the end of this process, we have $|E_{r+1}| \leq \epsilon^r |E| < 1$, and hence $E_{r+1} = \emptyset$, and hence the resulting code A is E -unbiased.

Using the fact that $n \leq m$, we may compute r as

$$r \leq 2 + \frac{\log m}{\log(1/\epsilon)} \leq 2 - \frac{\log m}{\log\left(2^{-\frac{\log m}{\log \log \log m}} (\log \log m + 1) n^{\frac{1}{\log \log m}}\right)} \leq \frac{\log \log \log m}{\log 2} + O(1).$$

So $L = rs \leq \log_2 m + O(\frac{\log m}{\log \log \log m})$. Overall, this procedure has quasi-complexity $(rk \log m, 2^s W) = (\log(mn), W)$.

Next, we discuss how to modify this procedure when $m < n$. In that case, with a simple pre-processing step of quasi-complexity $(\log mn, W)$, we can identify for each $e \in E$ a coordinate $v_e \in e$. Let $V' = \{v_e \mid e \in E\}$, and define $E' = \{e \cap V' \mid e \in E\}$. Using the above procedure we find a code A' of length $L \leq \log_2 m + O(\frac{\log m}{\log \log \log m})$ which is E' -unbiased. We finish by setting

$$A(\ell) = \begin{cases} A'(\ell) & \text{if } \ell \in V' \\ 0 & \text{otherwise.} \end{cases} \quad \square$$

2.3 Fooling Neighborhoods

If we wish to fool a list of neighborhoods $E \subseteq 2^{[n]}$, we could apply Theorem 2.9 to the set $E' = \{f \mid f \subseteq e \in E\}$. However, even forming E' directly might require exponential work. Instead, we can modify our algorithm to construct an E' -unbiased code, without needing to list E' explicitly.

As in the proof of Proposition 2.8, we associate to each $i \in [n]$ a distinct non-zero monomial μ_i over $GF(2^s)[z_1, \dots, z_k]$, wherein each indeterminate z_i has degree at most $d = \lceil n^{1/k} - 1 \rceil$. We also define $\mu_e = \sum_{i \in e} \mu_i$ for any $e \subseteq [n]$.

We will form the code A as

$$A(\ell) = (\mu_\ell(\alpha_{1,1}, \dots, \alpha_{1,k}), \mu_\ell(\alpha_{2,1}, \dots, \alpha_{2,k}), \dots, \mu_\ell(\alpha_{r,1}, \dots, \alpha_{r,k}))$$

for appropriate values $\alpha_{i,j} \in GF(2^s)$, where $i = 1, \dots, r$ and $j = 1, \dots, k$ (where here we identify elements of $GF(2^s)$ with binary vectors of length s). We also define $\alpha_{(i)} = (\alpha_{i,1}, \dots, \alpha_{i,k})$, and we write α as shorthand for $(\alpha_{1,1}, \dots, \alpha_{r,k})$.

For $i = 1, \dots, r$ and $j = 1, \dots, k$ and $e \in E$, let us define the potential function

$$F_{i,j,e}(\alpha) = \sum_{\substack{f \subseteq e \\ f \neq \emptyset}} [\mu_f(\alpha_{(1)}) = \dots = \mu_f(\alpha_{(i-1)}) = 0 \wedge \mu_f(\alpha_{i,1}, \dots, \alpha_{i,j}, z_{j+1}, \dots, z_k) \equiv 0].$$

The function $F_{i,j,e}(\alpha)$ only depends on $\alpha_{1,1}, \dots, \alpha_{1,k}, \alpha_{2,1}, \dots, \alpha_{2,k}, \dots, \alpha_{i,1}, \dots, \alpha_{i,j}$. Note here that $\mu_f(\alpha_{(1)}), \dots, \mu_f(\alpha_{(i-1)})$ are elements of $GF(2^s)$ while $\mu_f(\alpha_{i,1}, \dots, \alpha_{i,j}, z_{j+1}, \dots, z_k)$ is a polynomial in $k - j$ variables over $GF(2^s)$. This function $F_{i,j,e}(\alpha)$ can be regarded as a type of pessimistic estimator for the number of sets $f \subseteq e$ for which the code A will be biased.

PROPOSITION 2.10. *For any values i, j, e , the function $F_{i,j,e}$ can be computed with complexity $(\text{polylog}(|e|, r, s), \text{poly}(|e|, r, s))$.*

PROOF. Let $w = |e|$. If we associate the collection of subsets of e with binary vectors of length w , then the set of all $f \subseteq e$ which satisfy the given constraint is a linear subspace U , and so $F_{i,j,e}(\alpha)$ has the value $(2^{\text{rank}(U)} - 1)$. Thus, we need to compute the rank of the set of vectors

$$X = \{\mu_\ell(\alpha_{(1)}), \dots, \mu_\ell(\alpha_{(i)}), \mu_\ell(\alpha_{i,1}, \dots, \alpha_{i,j}, z_{j+1}, \dots, z_k) \mid \ell \in e\}.$$

Here $\mu_\ell(\alpha_{i,1}, \dots, \alpha_{i,j}, z_{j+1}, \dots, z_k)$ is regarded as a listing of coefficients.

Let us count the length of each such vector. Each term $\mu_\ell(\alpha_{(i)})$ is an entry of $GF(2^s)$, hence has length s . Each value of ℓ corresponds to a distinct monomial $\mu_\ell(\alpha_{i,1}, \dots, \alpha_{i,j}, z_{j+1}, \dots, z_k)$, so over all we need to keep track of at most w distinct monomials for the polynomial $\mu_\ell(\alpha_{i,1}, \dots, \alpha_{i,j}, z_{j+1}, \dots, z_k)$, for which each coefficient also has length s . In total, the length of a vector $x \in X$ is at most $(r + w)s$, and the number of such vectors is $|X| = w$. There is an NC algorithm to compute matrix rank [13]; thus, this rank calculation has overall complexity $(\text{polylog}(w, r, s), \text{poly}(w, r, s))$. $\square$

THEOREM 2.11. *Let $E \subseteq 2^{[n]}$, $m = |E|$ and $\text{width}(E) = w$. There is an algorithm with quasi-complexity $(w + \log(mn), (m + n) \text{poly}(w))$ to produce a code A of length $L \leq (1 + o(1))(w + \log_2 m)$, such that Ω^A fools E .*

PROOF. We assume $n \leq mw$, as we can simply ignore all coordinates which do not appear in E . For $i = 1, \dots, r$ and $j = 1, \dots, k$ let us define

$$H_{i,j}(\alpha) = \sum_{e \in E} F_{i,j,e}(\alpha).$$

We also define $H_{0,k}(\alpha) = \sum_{e \in E} (2^{|e|} - 1)$, so for $i = 0, \dots, k$, we have

$$H_{i,k}(\alpha) = \sum_{e \in E} \sum_{\substack{f \subseteq e \\ f \neq \emptyset}} [\mu_f(\alpha_{(1)}) = \dots = \mu_f(\alpha_{(i)}) = 0].$$

If $H_{r,k}(\alpha) = 0$, then the code A is unbiased for every $f \subseteq e \in E$. Our strategy will be to loop over $i = 1, \dots, r$ and then $j = 1, \dots, k$, selecting $\alpha_{i,j}$ at each stage to minimize $H_{i,j}(\alpha)$.

We now make a few observations on the sizes of $F_{i,j,e}(\alpha)$. First, $H_0(\alpha) \leq 2^w m$. Also, since $\mu_f(z_1, \dots, z_k) \neq 0$, we always have $F_{i,0,e}(\alpha) = 0$. Next, observe that if $\alpha_{i,j} \sim GF(2^s)$, then for any f with $\mu_f(\alpha_{i,1}, \dots, \alpha_{i,j-1}, z_j, z_{j+1}, \dots, z_k) \neq 0$, Proposition 2.7 gives

$$P(\mu_f(\alpha_{i,1}, \dots, \alpha_{i,j}, z_{j+1}, \dots, z_k) \equiv 0) \leq d/2^s \leq n^{1/k}/2^s.$$

Consequently, when $\alpha_{i,j} \sim GF(2^s)$ and we condition on $\alpha_{1,1}, \dots, \alpha_{i,1}, \dots, \alpha_{i,j-1}$, we have

$$\mathbb{E}[F_{i,j,e}(\alpha)] \leq F_{i,j-1,e}(\alpha) + (n^{1/k}/2^s)(F_{i-1,k,e}(\alpha) - F_{i,j-1,e}(\alpha)).$$

By selecting $\alpha_{i,j}$ to minimize $H_{i,j}(\alpha)$, we thus ensure that

$$H_{i,j}(\alpha) \leq H_{i,j-1}(\alpha) + (n^{1/k}/2^s)(H_{i-1,k}(\alpha) - H_{i,j-1}(\alpha)) \leq H_{i,j-1}(\alpha) + (n^{1/k}/2^s)H_{i-1,k}(\alpha).$$

Since $F_{i,0,e}(\alpha) = 0$, this in turn ensures that $H_{i,k}(\alpha) \leq (kn^{1/k}/2^s)H_{i-1,k}(\alpha)$, so $H_{r,k}(\alpha) < (2^w m)(kn^{1/k}/2^s)^r$. Thus, for

$$r = \left\lceil \frac{\log(2^w m)}{\log(2^s/(kn^{1/k}))} \right\rceil,$$

the code A will fool all of E .

Now set $k = \lceil \log \log(mw) \rceil$ and $s = \lceil \frac{\log(mw)}{\log \log \log(mw)} \rceil$. Using the fact that $n \leq mw$, calculations similar to Theorem 2.9 show that $r \leq (1 + o(1))(w + \log_2 m) \log \log \log(wm) / \log(wm)$. The code A has length $L = rs \leq (1 + o(1))(w + \log_2 m)$.

Next let us examine the complexity of this process. In each iteration, we must evaluate $F_{i,j,e}$ for every $e \in E$ and every $\alpha_{i,j} \in GF(2^s)$. By Proposition 2.10, each evaluation of $F_{i,j,e}$ has complexity $(\text{polylog}(w, \log mn), \text{poly}(w, \log mn))$. Over all possible values $\alpha_{i,j} \in GF(2^s)$ and $e \in E$, this gives a total complexity of $(\text{polylog}(w, \log mn), 2^s m \text{poly}(w, \log mn))$.

There are $rk \leq \tilde{O}(\frac{w+\log m}{\log(mw)})$ iterations, so the overall complexity of this process is $\tilde{O}(\frac{w+\log m}{\log(mw)}) \times \tilde{O}(\log(mn) + (\log w)^{O(1)})$ time and $(m+n)^{1+o(1)} w^{O(1)}$ processors. As $n \leq mw$, this simplifies to $\tilde{O}(w + \log(mn))$ time. $\square$

2.4 Comparison with Previous Algorithms

Let us briefly compare Theorem 2.11 with previous algorithms for fooling neighborhoods. The simplest approach to fool a list $E \subseteq 2^{[n]}$, is to select a code A whose dual code has weight $w + 1$, where $w = \text{width}(E)$. The resulting probability space Ω^A is then w -wise-independent. There are algebraic constructions to do so efficiently; for example, [2] discussed how to use BCH codes in this context for derandomization. Such codes have length roughly $(w/2) \log_2 n$.

An algorithm of Schulman [15] reduces the code-length significantly to $O(w + \log |E|)$. To do so, it generates the set $E' = \{f \subseteq e \mid e \in E\}$ and then uses an algorithm fooling Fourier characters, similar to Theorem 2.9, to generate a code which is unbiased for E' . A similar approach is used in [6], which interleaves other algorithmic steps with the generation of the code. The basic idea of both these works is to form Ω as a product of many independent copies of an ϵ -approximately-independent probability space, where ϵ is constant. In [15], the underlying ϵ -approximately-independent probability space was based on a construction of [14] using Reed-Solomon codes; these have a particularly nice form for derandomizing part of the random seed.

These algorithms have high processor complexity (approximately $O(mn2^w)$), and there are two main reasons for this. First, simply enumerating the set E' requires a large processor count, exponential in w . Second, these algorithms test all possible seeds for the underlying Reed-Solomon code, and this requires a processor complexity exponential in the seed-length of that code.

Theorem 2.11 thus improves in two ways over the previous algorithms. First, it has reduced time and processor complexity; in particular, it answers an open problem posed by Schulman [15] in giving an NC¹ algorithm for $w \leq O(\log(mn))$, and it gives an NC algorithm for $w = \text{polylog}(mn)$. Second, the code size is smaller: it gives $L \leq (1 + o(1))(w + \log_2 m)$, whereas the previous algorithms only guarantee $L \leq O(w + \log m)$.

3 FOOLING SUMS OF JUNTAS

We say that a function $f : \{0, 1\}^n \rightarrow \mathbf{R}$ is a w -junta if there exists a set $Y = \{y_1, \dots, y_w\} \subseteq [n]$, such that

$$f(x_1, \dots, x_n) = f'(x_{y_1}, \dots, x_{y_w})$$

for some function $f' : \{0, 1\}^w \rightarrow \mathbf{R}$.

In this section, we consider a function $S : \{0, 1\}^n \rightarrow \mathbf{R}$ of the form

$$S(x) = \sum_{j=1}^m f_j(x), \tag{1}$$

where each f_j is a w -junta whose value is determined by a variable subset $Y_j \subseteq [n]$. Our goal is to find some $x \in \{0, 1\}^n$ with the property that $S(x) \geq \mathbf{E}_{X \sim \{0, 1\}^n} [S(X)]$. Our algorithm has four main components, which we will describe in turn:

- (1) We show how to apply conditional expectations when the objective functions are Fourier characters.
- (2) We apply Fourier decomposition to the sum (1), thus reducing a sum of w -juntas to a sum of Fourier characters. As we have discussed earlier, this step implemented directly has an exponential processor dependence on w .
- (3) We use (derandomized) random variable partitioning to break the overall sum into w/w' subproblems involving w' -juntas, where $w' = o(\log n)$.
- (4) We introduce an object we refer to as *partial-expectations oracle (PEO)*, which allows us to use conditional expectations to solve these subproblems sequentially. This requires $O(w)$ time, but only $2^{w'} = n^{o(1)}$ processors.

Berger and Rompel [3] discusses a few alternate strategies to mitigate the exponential dependence on w , for example when the underlying variables are drawn from $\{0, 1\}^b$ for $b = \text{polylog}(n)$, or when f_j are indicators of affine functions. But these strategies are not as general as we need for many applications. Our overall algorithm will handle all these situations as special cases.

3.1 Conditional Expectations for Sums of Fourier Characters and Sums of w -juntas

Our approach begins with a subroutine for optimization problems involving sums of Fourier characters. This idea has been used in a number of deterministic algorithms, starting with [11] and more extensively developed in [3, 6, 10]. We present a slightly optimized form.

THEOREM 3.1. *Let $E \subseteq 2^{[n]}$ be given along with associated weights γ_e for every $e \in E$. There is an algorithm with quasi-complexity $(\log mn, W)$ to find $x \in \{0, 1\}^n$ such that*

$$\sum_{e \in E} \gamma_e \chi_e(x) \geq \gamma_{\emptyset},$$

where $m = |E|$ and $W = n + \sum_{e \in E} |e|$.

PROOF. First, use Theorem 2.9 to construct the E -unbiased code A of length $L = O(\log m)$, using quasi-complexity $(\log mn, W)$. For any $y \in GF(2)^L$, let us define

$$G(y) = \sum_{e \in E} \gamma_e (-1)^{A(e) \cdot y}.$$

We want to find $y \in GF(2)^L$ with $G(y) \geq \gamma_0$; we can then produce the desired $x \in \{0, 1\}^n$ by setting $x_i = A(i) \cdot y$ for $i = 1, \dots, n$. Since A is E -unbiased, we have $\mathbb{E}_{y \sim GF(2)^L} [(-1)^{A(e) \cdot y}] = 0$ for every $e \neq \emptyset$. So $\mathbb{E}_{y \sim GF(2)^L} [G(y)] = \gamma_0$, and thus a satisfying $y \in GF(2)^L$ exists.

To find it, we use conditional expectation: we guess chunks of $t = \frac{\log mn}{\log \log mn}$ bits of y at a time, to ensure that the expected value of $\mathbb{E}[G(y)]$ increases. For each such guess, we will compute in parallel the resulting expected value $\mathbb{E}[G(y)]$, when certain bits of y are fixed and the rest remain independent fair coins. We may compute the conditional expectations of a term $(-1)^{A(e) \cdot y}$, using the following observation: suppose that $y_1, \dots, y_k$ are determined while $y_{k+1}, \dots, y_L$ remain independent fair coins. Then, $\mathbb{E}[(-1)^{A(e) \cdot y}] = 0$ if $A(i) = 1$ for any $i \in e \cap \{k+1, \dots, L\}$, and otherwise $(-1)^{A(e) \cdot y}$ is determined by $y_1, \dots, y_k$.

This process requires $\lceil L/t \rceil \leq O(\log \log mn)$ rounds. For each possible value for a t -bit chunk of y , evaluating G has complexity $(\log mn, W^{1+o(1)})$. As $t \leq o(\log mn)$, we get an overall quasi-complexity of $\tilde{O}(\log mn, W)$. $\square$

LEMMA 3.2. *Suppose we have a full listing of the truth-table of each f_j . There is an algorithm to find $x \in \{0, 1\}^n$ satisfying $S(x) \geq \mathbb{E}_{X \sim \{0, 1\}^n} [S(X)]$, using quasi-complexity $(w + \log mn, 2^w m + n)$.*

PROOF. Using Proposition 2.2, transform each f_j as $f_j(x) = \sum_{e \subseteq Y_j} \gamma_{j,e} \chi_e(x)$. This step has quasi-complexity $(w + \log mn, 2^w m)$. We thus have

$$S(x) = \sum_j \sum_{e \subseteq Y_j} \gamma_{j,e} \chi_e(x) = \sum_e \chi_e(x) \left(\sum_j \gamma_{j,e} \right)$$

and $\sum_j \gamma_{j,\emptyset} = \mathbb{E}_{X \sim \{0, 1\}^n} [S(X)]$.

Next, apply Theorem 3.1 to the set $E = \{e \mid e \subseteq Y_j\}$ and associated weights $\sum_j \gamma_{j,e}$. Since $|E| \leq 2^w m$ and $\text{width}(E) \leq w$, this procedure has quasi-complexity $(\log(2^w mn), 2^w m + n)$. $\square$

When $w = \text{polylog}(n)$, this means that Lemma 3.2 gives quasi-NC algorithms. When $w = \Theta(\log n)$, then Lemma 3.2 gives NC algorithms; however, the processor complexity (while polynomial) may be quite large, depending on the size of w .

As a side application of Theorem 3.1 (which is not needed for our overall derandomization approach), let us consider the *heavy-codeword problem*. We are given a code A of length L and size n , presented as a $L \times n$ generator matrix. Our goal is to find a codeword whose weight is at least the expected weight of a randomly chosen codeword. This was introduced as a toy derandomization problem by [14]; this work also gave an algorithm with complexity roughly $(\log Ln, L^2 n^2)$. This was later improved by [6] to complexity $(\log Ln, Ln^2)$. We improve this further to nearly optimal time and processor complexities.

COROLLARY 3.3. *There is an algorithm with quasi-complexity $(\log Ln, Ln)$ to find a heavy codeword.*

PROOF. We suppose without loss of generality that no row of A is all zero. In this case, the expected weight of a codeword is $L/2$. Letting $y_1, \dots, y_L$ denote the rows of A ; we wish to find a vector $x \in \{0, 1\}^n$ such that $y_j \cdot x = 1$ for at least $L/2$ values of j .

Define $S(x) = -\sum_{j=1}^L \chi_{y_j}(x)$. If $S(x) \geq 0$, then x is orthogonal to at least half of $y_1, \dots, y_L$ as desired. So we apply Theorem 3.1, noting that $W \leq Ln$. Since $y_1, \dots, y_L$ are all distinct from zero, we have $y_\emptyset = 0$. $\square$

3.2 Derandomized Variable Partitioning

This step is based on a derandomization technique of [1] using symmetric polynomials and approximately independent probability spaces (also known as small-bias probability spaces). We begin by defining and quoting some results on approximate independent probability spaces.

Definition 3.4. A probability space Ω over $\{0, 1\}^n$ is t -wise, ϵ -approximately independent, if for any indices $1 \leq i_1 < i_2 < \dots < i_t \leq n$, and any bits $y_1, \dots, y_t \in \{0, 1\}^t$, we have

$$P_{x \sim \Omega}(x_{i_1} = y_1 \wedge \dots \wedge x_{i_t} = y_t) \leq (1 + \epsilon)2^{-t}.$$

THEOREM 3.5 ([14]). For any integer $t \geq 1$ and $\epsilon > 0$, there is a t -wise, ϵ -approximately independent probability space Ω of support size $|\Omega| \leq 2^{O(t + \log(1/\epsilon) + \log \log n)}$. The space Ω can be constructed with quasi-complexity $(t + \log(1/\epsilon) + \log n, 2^{O(t + \log(1/\epsilon) + \log \log n)})$.

LEMMA 3.6. Let $E \subseteq 2^{[n]}$, where $m = |E|$ and $w = \text{width}(E)$. One can construct a partition of $[n]$ into R parts $[n] = T_1 \sqcup T_2 \sqcup \dots \sqcup T_R$, for some $R = O(1 + \frac{w(\log \log mn)^5}{\log mn})$, satisfying

$$|f \cap T_k| \leq O\left(\frac{\log mn}{\log \log \log mn}\right) \text{ for all } f \in E, k \in [R].$$

This algorithm has quasi-complexity $(\log w \log(mn), w^{O(1)}(m + n))$.

PROOF. Let $r = \lceil \log_2 \frac{Cw(\log \log mn)^5}{\log mn} \rceil$, where C is a constant to be specified. We will construct binary vectors $y_1, \dots, y_r \in \{0, 1\}^n$ and then define for each $k \in \{0, 1\}^r$, $\ell \in \{0, \dots, r\}$ the sets $T_k^\ell \subseteq [n]$ by

$$T_k^\ell = \{i \in [n] \mid y_1(i) = k(1) \wedge y_2(i) = k(2) \wedge \dots \wedge y_\ell(i) = k(\ell)\}.$$

We will finish by setting $R = 2^r$ and forming the sets $T_1, \dots, T_R$ by $T_k = T_k^r$ where k ranges over $\{0, 1\}^r$. For each $k \in \{0, 1\}^r$, $\ell \in \{0, \dots, r\}$, $f \in E$ we define $H_\ell(f, k) = |f \cap T_k^\ell|$. We will achieve the goal of the theorem if we select $y_1, \dots, y_r$ so that every $f \in E, k \in \{0, 1\}^r$ has $H_r(f, k) \leq t$ for $t = \lceil \frac{\log mn}{\log \log \log mn} \rceil$.

For each $\ell = 0, \dots, r$ let us define the potential function

$$Q_\ell = \sum_{f, k} \binom{H_\ell(f, k)}{t}.$$

Observe that Q_r is an integer; thus, if $Q_r < 1$, then it follows that $Q_r = 0$ and so $H_r(f, k) < t$ for all f, k as desired.

Let Ω be a probability distribution over $GF(2)$ which is t -wise, ϵ -approximately independent, where $\epsilon = 1/r$, according to Definition 3.4. By Theorem 3.5, we have $|\Omega| \leq (mn)^{o(1)} w^{O(1)}$; furthermore, the complexity of generating Ω will be negligible for the overall algorithm.

For $y_\ell \sim \Omega$, each t -tuple of elements in $f \cap T_k^\ell$ has a probability of at most $2^{-t}(1 + \epsilon)$ of surviving to $T_k^{\ell+1}$. This implies that

$$\mathbb{E} \left[\binom{H_{\ell+1}(f, k)}{t} \mid y_1, \dots, y_\ell \right] \leq (1 + \epsilon)2^{-t} \binom{H_\ell(f, k)}{t}$$

and consequently, $\mathbb{E}[Q_{\ell+1} \mid y_1, \dots, y_\ell] \leq (1 + \epsilon)2^{-t}Q_\ell$.

Our algorithm is to select $y_1, \dots, y_r$ sequentially in order to minimize $Q_{\ell+1}$ at each stage ℓ . This ensures that $Q_\ell \leq (1 + \epsilon)2^{-t}Q_{\ell-1}$, and so at the end of the process we have

$$Q_r \leq (1 + \epsilon)^r 2^{-tr} Q_0 = (1 + 1/r)^r R^{-t} \sum_{f,k} \binom{H_0(f,k)}{t} \leq eR^{1-t} m \binom{w}{t} \leq eR^{1-t} m(ew/t)^t.$$

Simple calculations now show that $Q_r < 1$ for C a sufficiently large constant.

We now examine the complexity of this algorithm. There are r stages; in each stage, we must search the probability space Ω and compute Q_ℓ . The potential function Q_ℓ can be computed with quasi-complexity $(\log mn, mwR)$. Note now that $mwR \leq (mn)^{o(1)}mw^{O(1)}$. As $|\Omega| \leq (mn)^{o(1)}w^{O(1)}$, this costs $(mn)^{o(1)}w^{O(1)}(m+n)$ processors and $\tilde{O}(r \log mn) = \tilde{O}(\log w \log mn)$ time. $\square$

3.3 The Partial-Expectations Oracle

As we have discussed, we need implicit access to f_j in order to avoid the exponential dependence on w . A key idea of Berger and Rompel [3] to achieve this is an algorithm capable of determining certain conditional expectations for the objective functions.

Definition 3.7. Algorithm $\mathcal{A}$ is a PEO for the functions f_j , if it is capable of the following operation. Given any $X' \in \{0, 1, ?\}^n$, the algorithm $\mathcal{A}$ computes $F_j = \mathbb{E}_{X \sim \Omega}[f_j(X)]$ for $j = 1, \dots, m$, where the probability distribution Ω is defined by drawing each bit X_i independently, such that if $X'_i = ?$, then X_i is Bernoulli-1/2 and if $X'_i \neq ?$, then $X_i = X'_i$.

We note that this form of PEO is simpler than that used by Berger and Rompel: the latter requires evaluating the conditional expectation of $f_j(X)$ given that X is confined to an affine subspace, while our PEO only requires computing this conditional expectation when individual bits of X are fixed.

We now combine all the ingredients to obtain our conditional expectations algorithm.

THEOREM 3.8. Suppose $S(x) = \sum_{j=1}^m f_j(x)$ for $x \in \{0, 1\}^n$, where each f_j is a w -junta. Suppose we have a PEO for the functions f_j with complexity (C_1, C_2) .

Then there is an algorithm to find a vector x satisfying

$$S(x) \geq \mathbb{E}_{X \sim \{0,1\}^n} S(X),$$

with quasi-complexity $(C_1(1 + \frac{w}{\log mn}), w^{O(1)}C_2)$.

PROOF. We assume $C_1 \geq \Omega(\log mn)$ and $C_2 \geq \Omega(m+n)$ as it requires this complexity to take as input the values j, X' and output F_j . We similarly assume that $n \leq mw$, as variables not involved in any objective function may be ignored.

First apply Lemma 3.6 to determine a partition $[n] = T_1 \sqcup \dots \sqcup T_R$ for some $R = O(1 + \frac{w(\log \log mn)^5}{\log mn})$, such that $|Y_j \cap T_k| \leq w'$ for some $w' \leq O(\frac{\log mn}{\log \log \log mn})$. This stage has quasi-complexity $(\log w \log mn, w^{O(1)}(m+n))$.

Next, for $r = 1, \dots, R$, we seek to determine the bits $\{x_i \mid i \in T_r\}$. Define the function $f'_j(z)$ to be the expected value of $f_j(X)$, when the entries X_i for $i \in T_r$ are set to z_i , the variables X_i for $i \in T_1, \dots, T_{r-1}$ are set to x_i , and the remaining entries of X (for $i \in T_{r+1}, \dots, T_R$) remain fair coins. Each f'_j is a w' -junta and we can determine its truth-table f'_j using $2^{w'}$ invocations of our PEO, where we define $X'_i = ?$ for $i \in T_{r+1} \cup \dots \cup T_R$ and $X'_i \neq ?$ otherwise. This in turn requires $C_1 + \tilde{O}(\log(2^{w'}mn)) \leq \tilde{O}(C_1)$ time and $(mn)^{o(1)}2^{w'}C_2 \leq (mn)^{o(1)}C_2$ processors.

Next, apply Lemma 3.2 to determine a value for the relevant variables in T_r ; this step takes $\tilde{O}(w' + \log mn) \leq \tilde{O}(\log mn)$ time and $(n + 2^{w'}m)^{1+o(1)} \leq (mn)^{o(1)}(m+n)$ processors.

Over all R stages, the total time for this algorithm is $\tilde{O}(RC_1) \leq \tilde{O}(\frac{wC_1}{\log mn} + w + C_1)$. $\square$

We emphasize the low time and processor complexity of this algorithm. For example, if $w = \text{polylog}(mn)$ and $C_1 = \tilde{O}(\log mn)$ (which are typical parameters), then this has quasi-complexity (w, C_2) . Even if $w = \Theta(\log mn)$, this can lead to greatly reduced complexities as compared to the algorithm of [3].

This algorithm requires an appropriate PEO, which must be constructed in a problem-specific way. One important class of objective functions, which was one of the main cases considered by Berger and Rompel [3], is indicator functions for affine spaces; PEO's for such functions can be derived by a rank calculation. We will consider more complicated types of PEO's; one significant difficulty, as we discuss next, is that many objective functions are naturally represented as functions of integer-valued variables (not just isolated bits).

3.4 Non-Binary Variables

Let us consider a slightly more general setting: we have n variables $x_1, \dots, x_n$, each of which is an integer in the range $\{0, \dots, 2^b - 1\}$. Our objective function is again a sum of w -juntas, that is, each $f_j(x)$ depends on at most w coordinates of x . This can easily be reduced to the model we have discussed earlier: we replace each variable x_i with b separate binary variables $x_{i1}, \dots, x_{ib}$. Now each f_j depends on wb bits of the expanded input, and so is a wb -junta.

However, there is a complication. In order to apply Theorem 3.8, we need a PEO for the functions f_j . Thus, we need to compute the expected value of $f_j(x)$, given that certain *bits* of x are fixed to specific values. This can be somewhat awkward, as restricting arbitrary bits of x_i does not necessarily have any natural interpretation when x_i is an integer in the range $\{0, \dots, 2^b - 1\}$. It is often easier to use the strategy of [3], which fixes the bit-levels of $x_1, \dots, x_n$ one at a time. This allows us to use a simpler type of PEO where the pattern of known/unknown bits is more controlled.

For the purposes of the algorithm, we identify the integer set $\{0, \dots, 2^b - 1\}$ with the set of length- b binary vectors; a vector $(x_0, \dots, x_{b-1})$ corresponds to the integer $\sum_{i=0}^{b-1} 2^i x_{b-1-i}$. Note here that x_0 is the *most-significant bit*. Let us define $\mathcal{M}_b$ as the set $\{0, \dots, 2^b - 1\}$ equipped with this bit-based interpretation. Likewise, if $x \in \mathcal{M}_b^n$, we let $x(i, j)$ denote the j^{th} most significant bit of the integer value x_i .

Definition 3.9. We say that $X' \in \{0, 1, ?\}^{nb}$ is *graded* if there is some integer $\ell \in \{0, \dots, b-1\}$ such that for all i, j the following two conditions hold:

- (1) $X'(i, j) \in \{0, 1\}$ for $j = 0, \dots, \ell - 1$.
- (2) $X'(i, j) = ?$ for $j = \ell + 1, \dots, b - 1$.

We say that X' is *fully-graded* if X' satisfies for some integer $\ell \in \{0, \dots, b\}$ the stricter condition that for all i, j the following two conditions hold:

- (1) $X'(i, j) \in \{0, 1\}$ for $j = 0, \dots, \ell - 1$.
- (2) $X'(i, j) = ?$ for $j = \ell, \dots, b - 1$.

An algorithm $\mathcal{A}$ is a *graded PEO* (*fully-graded PEO*, respectively) for the functions f_j if it is a PEO, but only for queries X' which are graded (fully-graded, respectively).

THEOREM 3.10. Suppose that $S(x) = \sum_{j=1}^m f_j(x)$ for $x \in \mathcal{M}_b^n$, where each function f_j is a w -junta, and we have a graded PEO for the functions f_j with complexity (C_1, C_2) .

Then we can find $x \in \{0, 1\}^n$ satisfying $S(x) \geq \mathbb{E}_{X \sim \mathcal{M}_b^n}[S(X)]$, using quasi-complexity $(bc_1(1 + \frac{w}{\log(mn)}), w^{O(1)}C_2)$.

PROOF. We will determine the bits of x in b separate stages; at the ℓ^{th} stage, we determine the bit-level ℓ of each entry x_i . For $\ell = 0, \dots, b-1$, consider the following process. Define the function $f_{\ell,j}(z)$ to be the expected value of $f_j(X)$, when the bit-levels $0, \dots, \ell-1$ of X are taken from the already-determined vector x ; when the bit-level ℓ of X is set to z ; and when the bit-levels $\ell+1, \dots, b-1$ of X are independent fair coins.

Each $f_{\ell,j}$ is a w -junta, and the graded PEO for the functions f_j yields a PEO for the functions $f_{\ell,j}$. Therefore, Theorem 3.8 produces a $z \in \{0, 1\}^n$ with $\sum_j f_{\ell,j}(z) \geq \mathbb{E}_{Z \sim \{0,1\}^n} [\sum_j f_{\ell,j}(Z)]$. $\square$

One important application of non-binary variables concerns derandomizing biased coins. For a vector of probabilities $p \in [0, 1]^n$, consider the probability space with n independent variables $X_1, \dots, X_n$, wherein each X_i is Bernoulli- p_i . We write this more compactly as $X \sim p$. Most of our derandomization results we have proved earlier have assumed that the underlying random bits are independent fair coins (i.e., with $p_1, \dots, p_n = 1/2$).

Definition 3.11. An algorithm $\mathcal{A}$ is a *continuous PEO* for the functions f_j , if it is capable of the following operation. Given any vector $q \in [0, 1]^n$, whose entries are rational numbers with denominator 2^b , the algorithm $\mathcal{A}$ computes $F_j = \mathbb{E}_{X \sim q} [f_j(X)]$ for $j = 1, \dots, m$.

Note that a PEO can be regarded as a special case of a continuous PEO, in which the probability vector q is restricted to the entries $\{0, 1/2, 1\}$.

THEOREM 3.12. Suppose that $S(x) = \sum_{j=1}^m f_j(x)$ for $x \in \{0, 1\}^n$, where each function f_j is a w -junta, and we have a continuous PEO for the functions f_j with complexity (C_1, C_2) .

Let $p \in [0, 1]^n$ be a vector of probabilities, wherein each entry p_i is a rational number with denominator 2^b . Then we can find a vector x satisfying $S(x) \geq \mathbb{E}_{X \sim p} [S(X)]$, using quasi-complexity $(bC_1(1 + \frac{w}{\log mn}), (wb)^{O(1)}C_2)$.

PROOF. Consider the function $f'_j : \mathcal{M}_b^n \rightarrow \mathbb{R}$ defined by

$$f'_j(y_1, \dots, y_n) = f_j([y_1/2^b \leq p_1], \dots, [y_n/2^b \leq p_n]).$$

Each function f'_j depends on w coordinates of y . Furthermore, if certain most-significant bit levels of y are fixed to a certain value and the remaining least-significant bit-levels of y are independent fair coins, then each term $[y_i/2^b \leq p_i]$ is a Bernoulli- q_i variable, where q_i depends on the fixed values of y_i . Therefore, the given continuous PEO for f_j provides a graded PEO for the functions f'_j , with a complexity of $(\log(nb) + C_1, C_2 + nb)$.

Finally, observe that when $Y \sim \mathcal{M}_b^n$, each term $[Y_i/2^b \leq p_i]$ is Bernoulli- p_i ; therefore, we have

$$\mathbb{E}_{Y \sim \mathcal{M}_b^n} \left[\sum_j f'_j(Y) \right] = \mathbb{E}_{X \sim p} \left[\sum_j f_j(X) \right].$$

So Theorem 3.10 produces $y_1, \dots, y_n \in \mathcal{M}_b^n$ with $\sum_j f'_j(y) \geq \mathbb{E}_{X \sim p} [\sum_j f_j(X)]$. Output the vector $x \in \{0, 1\}^n$ defined by $x_i = [y_i/2^b \leq p_i]$. $\square$

This leads to PEOs for the class of functions computed by a read-once branching program (ROBP). In this computational model, the function f is represented as a directed acyclic graph; at each node v , a single variable x_v is read and the program branches to two possible destinations depending on the variable x_v . There is a designated starting vertex and at some designated sink vertices, a real number is output. In addition, every variable label appears at most once on each directed path. This is a quite general class of functions, which includes log-space statistical tests as used by Sivakumar's derandomization [16]. See [4] for further details.

PROPOSITION 3.13. *If a w -junta f can be computed by a ROBP on M states, then it has a continuous PEO with quasi-complexity $(\log b \log w \log Mw, bM^3w^{O(1)})$.*

PROOF. We must calculate the expected value of f , given that the variables $X_1, \dots, X_n$ are independent Bernoulli- q_i . Now observe that, for any states s_1, s_2 , the probability that s_1 goes to s_2 in at most h timesteps is the sum over intermediate states s of the probability that s_1 goes to s in at most $h/2$ timesteps and that s goes to s_2 in at most $h/2$ timesteps; this follows from the definition of an ROBP. Using this relation, one may recursively build the transition matrix for pairs of states $s_1 \rightarrow s_2$ over time horizons $h = 1, 2, 4, \dots, w$. Each such iteration takes time $\tilde{O}(\log b \log Mn)$ and there are $\tilde{O}(\log w)$ iterations. $\square$

3.5 Application: Rainbow Hypergraph Coloring

As a simple example of our derandomization method, let us consider a d -uniform hypergraph H , with m edges and n vertices. We say that an edge e is *rainbow* for a vertex coloring of H , if all its vertices receive distinct colors. A challenge determinization problem is to find a d -coloring with $m \frac{d!}{d^d}$ rainbow edges, which is the expected number in a uniform random coloring. In [2], an NC algorithm was given in the case $d = O(1)$. This was extended by [3] to arbitrary d ; although [3] did not give any concrete time or processor complexity, the complexity appears to be roughly $(\log^4 mn, n + m^{1+\log 2})$.

We significantly improve both the time and processor costs. Note that it requires $\Omega(md + n)$ space to store the hypergraph.

THEOREM 3.14. *There is an algorithm to find a vertex coloring with d colors and at least $m \frac{d!}{d^d}$ rainbow edges, using quasi-complexity $(\log^2 mn, md + n)$.*

PROOF. We begin with simple pre-processing steps using complexity $(\log(mn), md + n)$. First, when $d \geq \log m + C \log \log m$ for a sufficiently large constant C , then it suffices to rainbow-color a single edge, which may be done easily. Second, when $n < md$, then some vertex is not used; we may delete it from the graph. Hence, we assume $m \geq nd$ and $d \leq \log m + O(\log \log m)$.

Given a binary vector $x \in \mathcal{M}_b^n$, we define the associated d -coloring $\phi : V \rightarrow \{0, \dots, d-1\}$ by $\phi_x(v) = \lfloor (d/2^b)x_v \rfloor$. For each edge $e \in H$ let $f_e(x)$ be the indicator function that e is rainbow on the coloring ϕ_x , and define $S(x) = \sum_e f_e(x)$.

As shown in [3], by taking $b = \tilde{O}(\log mn)$, we can ensure that $\mathbb{E}_{X \sim \mathcal{M}_b^n}[S(X)] > \frac{(md!-1)}{d^d}$. Furthermore, since $S(x)$ is an integer and $\frac{(md!-1)}{d^d}$ is a rational number with denominator d^d , when $S(x) > \frac{(md!-1)}{d^d}$ we ensure that $S(x) \geq \frac{md!}{d^d}$. So, if we find $x \in \mathcal{M}_b^n$ with $S(x) \geq \mathbb{E}_{X \sim \mathcal{M}_b^n}[S(X)]$, then this will yield our desired coloring.

Since each f_e is a d -junta, we apply Theorem 3.10 with $w = d$ to construct ϕ_x using total quasi-complexity of $O(bC_1(1 + \frac{d}{\log mn}), n + C_2)$, where (C_1, C_2) is the complexity of PEO for the functions f_e . Observe that $b \leq \tilde{O}(\log mn)$, $d \leq \log m + O(\log \log m)$. In Proposition A.1 (which we defer to the Appendix) we show that $C_1 \leq \tilde{O}(\log mn)$ and $C_2 \leq (m+n)^{1+o(1)}$, so this simplifies to quasi-complexity of $(\log^2 mn, m+n)$. $\square$

4 THE LOVÁSZ LOCAL LEMMA WITH COMPLEX BAD-EVENTS

The Lovász Local Lemma (LLL) is a keystone principle in probability theory which asserts that if one has a probability space Ω and a set $\mathcal{B}$ of “bad-events” in Ω , then under appropriate “local” conditions there is a positive probability that no event in $\mathcal{B}$ occurs. The LLL has numerous applications to combinatorics, graph theory, routing, and so on. The simplest “symmetric” form

of the LLL states that if each bad-event $B \in \mathcal{B}$ has probability $P_\Omega(B) \leq p$ and affects at most d bad-events (including itself), then if $epd \leq 1$, then $P(\bigcap_{B \in \mathcal{B}} \bar{B}) > 0$.

Although the LLL applies to general probability spaces, in most applications a simpler bit-based form suffices, wherein the space Ω has n variables $x_1, \dots, x_n$, which are independently drawn from $\mathcal{M}_b$. In this setting, each $B \in \mathcal{B}$ is a Boolean function f_B on a subset Y_B of the variables. We say that bad-events B, B' affect each other (and write $B \sim B'$) if $Y_B \cap Y_{B'} \neq \emptyset$. We say that $x \in \mathcal{M}_b^n$ avoids $\mathcal{B}$ if $f_B(x) = 0$ for all $B \in \mathcal{B}$.

In a seminal paper [12], Moser and Tardos introduced the following simple randomized algorithm, which we refer to as the MT algorithm, giving efficient randomized constructions for nearly all LLL applications in bit-based probability spaces.

ALGORITHM 1: The Moser-Tardos algorithm

- 1: Generate $x_1, \dots, x_n$ as independent fair coins.
 - 2: **while** some bad-event is true on x **do**
 - 3: Arbitrarily select some true bad-event B
 - 4: For each $i \in Y_B$, draw x_i as an independent fair coin. (We refer to this as *resampling* B .)
-

Under nearly the same conditions as the probabilistic LLL, the MT algorithm terminates in polynomial expected time. Moser and Tardos also gave a parallel (RNC) variant of this algorithm, requiring a slack compared to the LLL criterion.

There are two key analytic techniques introduced by [12] for this algorithm. The first is the idea of a *resampling table*. In the MT algorithm as we have presented it, the new values for each variable are drawn in an online fashion. Instead, one can imagine a fixed table R . This table records, for each variable i , an infinite list of values $R(i, 1), R(i, 2), \dots$, for that variable, which are all independent draws from $\mathcal{M}_b$. When the MT algorithm begins, it sets $x_i = R(i, 1)$ for each variable i ; if a variable x_i needs to be resampled, it sets $x_i = R(i, 2)$, and so forth. Once we have fixed a resampling table R , the MT algorithm can be executed deterministically.

We view the resampling table R as a function $R : [n] \times \mathbb{Z}_+ \rightarrow \mathcal{M}_b$. We define a *slice* to be a set $W \subseteq [n] \times \mathbb{Z}_+$ with the property that each $i \in [n]$ has at most one $j \in \mathbb{Z}_+$ with $(i, j) \in W$. For such a slice W , sorted as $W = \{(i_1, j_1), \dots, (i_k, j_k)\}$ with $i_1 < i_2 < \dots < i_k$, we define the projection π_W by setting $\pi_W(R) = (R(i_1, j_1), \dots, R(i_k, j_k))$.

The other key idea introduced by Moser and Tardos is the *witness tree*, which represents a possible execution path for the MT algorithm leading to a given resampling. This is explained in great detail in [12], which we recommend as an introduction. As a brief summary, suppose we want to explain why some bad-event B was resampled at time t . We form a witness tree τ by first placing a root node labeled by B , and then going in time through the execution log from time $t - 1$ to time 1. For each event B we encounter at time $s < t$, we look in τ to find if there is some node v' labeled by $B' \sim B$. If so, we place a node v labeled by B in the tree as a child of v' ; if there are multiple choices for v' , we always choose the one of greatest depth (if there are multiple choices at greatest depth, we break the tie arbitrarily).

For any witness tree τ and any node $v \in \tau$, we let $L(v) \in \mathcal{B}$ denote the label of v .

Definition 4.1 (Weight and size of witness tree). For a witness tree τ , we define the *size* of τ as the number of nodes in τ and we define the *weight* of τ as $w(\tau) = \prod_{v \in \tau} P_\Omega(L(v))$.

The most important result of [12], which explains why the MT algorithm works, is the *Witness Tree Lemma*.

LEMMA 4.2 ([12]). *The probability that a witness tree τ appears during the execution of the MT algorithm is at most $w(\tau)$.*

To prove this Lemma, [12] shows that τ imposes certain conditions on the resampling table R .

LEMMA 4.3 ([12]). *For any witness tree τ on t nodes, there is a set of slices W_v , indexed by nodes $v \in \tau$, such that*

- (A1) *For $v \neq v'$ we have $W_v \cap W_{v'} = \emptyset$.*
- (A2) *A necessary condition for τ to appear is that $f_{L(v)}(\pi_{W_v}(R)) = 1$ for every $v \in \tau$.*
- (A3) *The sets W_v can be determined from τ with quasi-complexity $(\log nt, nt)$.*
- (A4) *Every $v \in \tau$ has $|W_v| = |Y_{L(v)}|$.*

PROOF. For each node $v \in \tau$ and each $i \in [n]$, let $u_{i,v}$ denote the number of nodes v' which are at greater depth than v and which have $i \in Y_{L(v')}$. Define $W_v = \{(i, u_{i,v} + 1) \mid i \in Y_{L(v)}\}$. $\square$

Lemma 4.2 follows from Lemma 4.3; as the entries of R are fair coins, the probabilities of each event $f_{L(v)}(\pi_{W_v}(R)) = 1$ is $P_\Omega(L(v))$; furthermore, since the sets W_v are non-intersecting, these events are all independent.

4.1 Derandomizing Moser-Tardos

The original paper of Moser and Tardos gave a sequential deterministic algorithm that only worked for a very limited class of LLL instances, for example when d was constant. An NC algorithm was later given in [5], covering a slightly larger class of bad-events. This algorithm required satisfying the LLL criterion with a slack, in particular it required $epd^{1+\epsilon} \leq 1$ for some $\epsilon > 0$, and had a complexity of roughly $(\frac{\log^3 mn}{\epsilon}, m^{O(1/\epsilon)})$. An alternative NC algorithm was provided in [9], which is slightly faster than [5].

These latter algorithms have numerous conditions on the functions f_B ; roughly speaking, they require f_B to have decision-tree complexity of order $\log d$. The clearest example of this problem type is k -SAT, in which bad-event corresponds to a clause being violated. So each bad-event is defined by $x_{i_1} = j_1 \wedge \dots \wedge x_{i_k} = j_k$, a monomial in k variables.

Many other LLL applications, particularly those in which the bad-events are determined by sums of random variables, do not fit into this paradigm; we discuss two examples in Section 5. The hallmark of these types of problems is that the bad-events are *complex* Boolean functions; our focus here will be to give NC algorithms for such problems.

The analysis of [9] is based on an extension of the witness tree to a more general object referred to as a *collectible witness DAG* (CWD). These objects represent in a sense all the ways the MT algorithm could require a long execution time. This requires a great deal of notation to define properly, but the important point for us is that each CWD τ satisfies Lemma 4.3 in the same a witness tree does. We will not discuss the (technical) differences between witness trees and CWD's.

We say that a CWD τ is *compatible* with a resampling table R if condition (A2) of Lemma 4.3 is satisfied, namely, $f_{L(v)}(\pi_{W_v}(R)) = 1$ for every $v \in \tau$. For any set $\mathcal{T}$ of CWD's and a resampling table R , we define $\mathcal{T}^R \subseteq \mathcal{T}$ to be the set of CWD's $\tau \in \mathcal{T}$ compatible with R . We summarize some key results of [9] which are relevant to us.

LEMMA 4.4 ([9]). *Suppose that $epd^{1+\epsilon} \leq 1$ for $\epsilon > 0$, and suppose that the functions f_B can be evaluated with complexity $(U, \text{poly}(m, n))$ where $m = |\mathcal{B}|$.*

Let $K = \frac{c \log(mn/\epsilon)}{\epsilon \log d}$ for some constant $c > 0$. There is a set $\mathcal{T}$ of CWD's with the following properties:

- (T1) $|\mathcal{T}| \leq (mn/\epsilon)^{O(1/\epsilon)}$.
- (T2) Each $\tau \in \mathcal{T}$ has size at most $2K$.
- (T3) If every $\tau \in \mathcal{T}^R$ has size less than K , then an assignment avoiding $\mathcal{B}$ can be found with complexity $(KU + K \log(mn|\mathcal{T}^R|) + \log^2 |\mathcal{T}^R|, \text{poly}(K, m, |\mathcal{T}^R|))$.
- (T4) $\sum_{\tau \in \mathcal{T}, |\tau| \geq K} w(\tau) < 1/2$.
- (T5) $\sum_{\tau \in \mathcal{T}} w(\tau) < O(m)$.
- (T6) The set $\mathcal{T}$ can be enumerated with quasi-complexity $(\frac{\log^2(mn/\epsilon)}{\epsilon}, (mn/\epsilon)^{O(1/\epsilon)})$.

Now consider drawing a resampling table $R(i, j)$ where $i = 1, \dots, n$ and $j = 1, \dots, 2K$. Given any CWD τ , let $f(\tau, R)$ denote the indicator function that τ is compatible with R . We also define

$$S(R) = \frac{1}{Cm} \sum_{\tau \in \mathcal{T}} f(\tau, R) + \sum_{\substack{\tau \in \mathcal{T} \\ |\tau| \geq K}} f(\tau, R) \quad (2)$$

for some constant $C > 0$.

When the entries of R are drawnly independent from $\mathcal{M}_b$, properties (T4) and (T5), ensure that $\mathbb{E}[S(R)] < 1$ for a sufficiently large C . If $S(R) < 1$, the resulting resampling table R has the property that $|\mathcal{T}^R| \leq O(m)$ and that every $\tau \in \mathcal{T}^R$ has size less than K . Therefore, by property (T3), a satisfying assignment can be found in quasi-complexity $(\frac{U \log(mn) + \log^2(mn)}{\epsilon}, \text{poly}(m, n, 1/\epsilon))$. Thus, the problem of finding a satisfying assignment is reduced to the problem of minimizing $S(R)$, which is a sum of juntas.

THEOREM 4.5. *Suppose that $\text{epd}^{1+\epsilon} \leq 1$ for some $\epsilon > 0$. Let $m = |\mathcal{B}|$. Suppose that each bad-event $B \in \mathcal{B}$ is a w -junta (i.e. $|Y_B| \leq w$) on $\mathcal{M}_b^n$.*

Suppose that we have a graded PEO for the collection of functions f_B which has complexity $(C_1, \text{poly}(m, n))$; namely, given any $B \in \mathcal{B}$ as well as a graded partial assignment to the variables in Y_B , it computes the corresponding probability that $f_B = 1$.

Then we can find $x \in \mathcal{M}_b^n$ avoiding $\mathcal{B}$ in quasi-complexity $(wb(C_1 + \frac{\log mn}{\epsilon}) + \frac{C_1 \log mn}{\epsilon}, (mn)^{O(1/\epsilon)})$.

PROOF. If $\epsilon < 1/(mn)$, then we can solve this problem by exhaustive search in $2^n \leq (mn)^{O(1/\epsilon)}$ processors. So let us assume that $\epsilon \geq 1/(mn)$. We also assume $C_1 \geq \Omega(\log mn)$ as this time is required to read the input.

Let $m' = (mn)^{c/\epsilon}$. For an appropriate constant c , the objective function $S(R)$ is by (T1) a sum of at most m' functions $f(\tau, R)$. By (T2), each $\tau \in \mathcal{T}$ has size at most $2K$; property (A4) ensures that each term $f_{L(v)}(\pi_{W_v}(R))$ depends on at most w entries of R , so in all each function $f(\tau, R)$ is a w' -junta for some $w' = O(\epsilon^{-1}w \log mn)$. The total number of variables determining R is $n' = 2nK \leq O(\epsilon^{-1}n \log(mn))$. (We do not need to compute any entries of R beyond this point.)

We claim next that we can form a graded PEO for the functions $f(\tau, R)$ with quasi-complexity (C'_1, C'_2) , where $C'_1 = C_1 + \frac{\log mn}{\epsilon}$ and $C'_2 = (mn)^{O(1/\epsilon)}$. For suppose we are given a graded partial assignment query R' . We can compute the associated projections $\pi_{W_1}(R'), \dots, \pi_{W_v}(R')$ for each τ using $\tilde{O}(\frac{\log mn}{\epsilon})$ time. The probability that any τ is compatible with R' is simply the product of the probabilities of $f_{B_i}(\pi_{W_i}(R'))$. The PEO for f_B allows us to compute these probabilities in parallel with quasi-complexity $((\log K)C_1, (mn)^{O(1/\epsilon)})$. Finally, we multiply the probabilities together in $\tilde{O}(\log K)$ time.

Now apply Theorem 3.8 to find R with $S(R) \leq \mathbb{E}[S(R)] < 1$ using quasi-complexity $\tilde{O}(bC'_1(1 + \frac{w'}{\log m'n'}), (w')^{O(1)}C'_2)$, which simplifies to $(wb(C_1 + \frac{\log mn}{\epsilon}), (mn)^{O(1/\epsilon)})$. Once we have found R , we

use (T3) to find x avoiding $\mathcal{B}$. The PEO can be used to check whether a given bad-event is true, so $U \leq C_1$ and this step requires $O(\frac{C_1 \log mn}{\epsilon})$ time. $\square$

5 APPLICATIONS OF THE LLL

5.1 Defective Vertex Coloring

A k -defective vertex c -coloring of a graph $G = (V, E)$, is an assignment of colors to the vertices such that every vertex v has at most k neighbors with the same color as v . This generalizes proper vertex coloring, in that a proper vertex coloring is a 0-defective coloring. In this section, we give an algorithm which gives a k -defective c -coloring of a graph G of maximum degree Δ with $c = O(\Delta/k)$, for any choice of k in the range $\{1, \dots, \Delta\}$. The main idea, inspired by a similar randomized distributed algorithm of [7], is a degree-splitting step; when Δ is small, this can be achieved efficiently using our deterministic LLL algorithm and when Δ is large, then we can use an alternate algorithm based on simple Chernoff bounds.

LEMMA 5.1. *There is an absolute constant K with the following property. Given a graph G of maximum degree Δ and an integer parameter $j \leq \log_2(\frac{\Delta}{K \log \Delta})$, there is an algorithm with quasi-complexity $(\Delta \log n + \log^2 n, \text{poly}(n))$ to 2^j -color the vertices, so that each vertex v has at most $(\Delta/2^j)(1 + K\sqrt{(2^j/\Delta) \log \Delta})$ neighbors with the same color as v .*

PROOF. Consider applying the LLL to the random process in which each vertex independently and uniformly selects a color (represented as a j -bit string). Each vertex v has a bad-event B_v that it has too many neighbors of its own color. Thus, there are $m = n$ bad-events, and each bad-event involves at most $\Delta + 1$ variables. The number of neighbors of each color is a Binomial random variable with mean at most $\Delta 2^{-j}$. Note that $B_v \sim B_w$ iff v and w are at distance at most 2 in G . So in the sense of the LLL we have $d \leq \Delta^2$.

Let $\delta = c\sqrt{(2^j/\Delta) \log \Delta}$ for some constant c . For K sufficiently large, we ensure that $\delta \leq 1$. Therefore, by the Chernoff bound, B_v has probability at most $e^{-\mu\delta^2/3}$, which is smaller than Δ^{-4} for an appropriate choice of c . So, in the sense of the LLL, we have $p \leq \Delta^{-4}$. These parameters satisfy Theorem 4.5 with $\epsilon = 1/2$. Each bad-event B_v is a Boolean function on at most Δ variables, and a graded PEO can be constructed with quasi-complexity $(\log n, \text{poly}(n))$. Thus, Theorem 4.5 gives the desired goal in quasi-complexity $(j\Delta \log n + \log^2 n, \text{poly}(n))$. Note that $j \leq O(\log \Delta)$, and so it can be dropped from the quasi-complexity bounds. $\square$

THEOREM 5.2. *Let G be a graph with maximum degree Δ and $k \in \{1, \dots, \Delta\}$. Then there is an NC algorithm running in time $\tilde{O}(\log^2 n)$ to obtain a k -defective vertex c -coloring with $c = O(\Delta/k)$.*

PROOF. When $\Delta \geq \log n$, let us consider the random process of assigning every vertex a color uniformly selected from $\frac{\Delta}{\log n}$; a simple Chernoff bound shows that, with high probability, this ensures that each vertex has at most $C \log n$ neighbors of each color class, where $C > 0$ is some sufficiently large constant. This can be derandomized by an algorithm of Sivakumar [16] (among other methods), as there are a polynomial number of “statistical tests” (in this case, the degree of each vertex with respect to each color class) which can be computed in logspace. After this first coloring step, which can be executed in $O(\log^2 n)$ time, we get multiple subgraphs with maximum degree $\Delta \leq C \log n$.

Thus, we can assume that $\Delta \leq C \log n$ for a constant $C > 0$. In this case, we can use iterated applications of Lemma 5.1. Each iteration reduces the degree of the residual graphs by a logarithmic factor, and so the overall running time is close to the running time of a single application of Lemma 5.1. We defer the full proof to Appendix B, as the construction is technical and similar to that of [7]. $\square$

5.2 Domatic Partition

A *domatic partition* of a graph is a c -coloring of the vertices of G with the property that each vertex of G sees all c -colors in its neighborhood (including itself). That is, for any color $\ell = 1, \dots, c$, the color- ℓ vertices form a dominating set of G . An algorithm was given in [8] using the LLL to find a domatic partition with a large number of colors. For simplicity, we specialize their algorithm to k -regular graphs.

THEOREM 5.3. *Let $\eta > 0$ be any fixed constant. There is some constant $K = K_\eta$ with the following property. If G is k -regular with $k > K$, then G has a domatic partition of size $c \geq (1 - \eta) \frac{k}{\log k}$, which can be found using $\tilde{O}_\eta(k \log n + \log^2 n)$ time and $n^{O_\eta(1)}$ processors.*

PROOF. We follow the iterated LLL construction of [8], in which the color of each vertex is an ordered pair $\chi(v) = (\chi_1(v), \chi_2(v))$; here χ_1 is chosen from $c_1 = k / \log^3 k$ colors, and χ_2 is chosen from $c_2 = (1 - \eta) \log^2 k$ colors. In the first phase of the LLL, we will select χ_1 and in the second phase we will select χ_2 . Each vertex chooses its colors uniformly at random among $[c_1], [c_2]$ respectively.²

Now consider the phase I coloring. For each vertex v and each color $j \in [c_1]$, define $N_j(v)$ to be the set of neighbors w with $\chi_1(w) = j$ and let $X_{v,j} = |N_j(v)|$. The expected value of $X_{v,j}$ is $\mu = \log^3 k$. For each vertex v and each color $j \in [c_1]$, we have a bad-event $B_{v,j}$ that $X_{v,j} \leq t_0$ or $X_{v,j} \geq t_1$, where $t_0 = \mu - \phi \log^2 k$ and $t_1 = \mu + \phi \log^2 k$ and ϕ is a large constant.

For ϕ sufficiently large, the Chernoff bound shows that $B_{v,j}$ has probability at most $p \leq k^{-5}$. Furthermore, each bad-event $B_{v,j}$ affects $B_{v',j'}$ only if $\text{dist}(v, v') \leq 2$, so in the sense of the LLL we have $d \leq k^4$. A graded PEO for these bad-events with running time $C_1 = \tilde{O}(\log mn)$. Apply Theorem 4.5 to find χ_1 with quasi-complexity $(k \log n + \log^2 n, n^{O(1)})$.

For each vertex v , each $j \in [c_1]$, and each $j' \in [c_2]$, we have a bad-event $B_{v,j,j'}$ that there is no $w \in N_j(v)$ with $\chi_2(w) = j'$; if all such bad-events are avoided, then the resulting coloring $(\chi_1(v), \chi_2(v))$ gives a domatic partition. The only dependencies now are between bad-events $B_{v,j,j'}$ and $B_{w,j,j''}$, where v, w share a neighbor u with $\chi_1(u) = j$, so $d \leq t_1 k c_2$ and $p \leq (1 - 1/c_2)^{t_0}$.

Set $\epsilon = \eta/2, \phi = 10$. It is straightforward to verify that the criterion $epd^{1+\epsilon} \leq 1$ is satisfied when k is sufficiently large. Thus, Theorem 4.5 gives a coloring avoiding the phase-II bad-events using $\tilde{O}_\eta(k \log n + \log^2 n)$ time and $n^{O_\eta(1)}$ processors. $\square$

APPENDICES

A THE PEO FOR HYPERGRAPH RAINBOW COLORING

PROPOSITION A.1. *Let f_e be the indicator function that edge e is rainbow. Then the collection of functions f_e has a graded PEO with overall quasi-complexity $(\log mn, m + n)$.*

PROOF. It suffices to compute the probability that a given edge e will be rainbow on the coloring ϕ_x for some graded $u \in \{0, 1, ?\}^{db}$ (here u represents the projection of the overall partially graded $x \in \{0, 1, ?\}^{nb}$ to the vertices in e). Since $db = (mn)^{o(1)}$, the processor complexity of this task can be an arbitrarily polynomial in b, d .

We first describe how to do so if u is fully graded; we then modify it to allow u to be merely graded. Suppose the most-significant $\ell \leq b$ bit-levels of the vector y have been determined and the least-significant $b - \ell$ bit-levels of u remain fair coins. We may write u in the form

²If c_1, c_2 are not powers of two, then we cannot directly represent this in our bit-based LLL formulation. However, we can simulate it by drawing values u_1, u_2 from $[2^{r_1}, 2^{r_2}]$ for $r_i = \log_2 c_i + O(\log k)$, and projecting uniformly down to $[c_i]$. This changes the probabilities of the bad-events by a negligible factor of $\text{polylog}(1/k)$. We omit further details for simplicity.

$u_v = (y_v, ?, \dots, ?)$, where $y_v \in \mathcal{M}_\ell$. For each $c \in \mathcal{M}_\ell$, let S_c denote the set of vertices $v \in e$ with $y_v = c$.

For each $k \in \{0, \dots, d-1\}$, let us define G_k to be the set of values $c \in \{0, \dots, 2^\ell - 1\}$ such that some vertex $v \in S_c$ could (depending on the lower order bits of y) be assigned color k . Specifically,

$$G_k = \{c \in \{0, \dots, 2^\ell - 1\} \mid F(2^{b-\ell}c) \leq k \leq F(2^{b-\ell}(c+1) - 1)\},$$

where the function $F : \{0, 2^b - 1\} \rightarrow \{0, \dots, d-1\}$ is defined by $F(x) = \lfloor (d/2^b)x \rfloor$.

Observe that if $y \geq x + (d/2^b)$, then we must have $F(y) > F(x)$. Using this fact, we claim that each G_k is either a singleton set, or a set of two adjacent elements $\{c, c+1\}$. For, suppose not; then there must exist $c_1, c_2 \in G_k$ with $c_2 > c_1 + 1$ and

$$F(2^{b-\ell}c_1) \leq k \leq F(2^{b-\ell}(c_1+1) - 1),$$

$$F(2^{b-\ell}c_2) \leq k \leq F(2^{b-\ell}(c_2+1) - 1).$$

But, note that in this case,

$$(2^{b-\ell}c_2) - (2^{b-\ell}(c_1+1) - 1) = 2^{b-\ell}(c_2 - c_1 - 1) + 1 \geq 2^{b-\ell} + 1 \geq 1 \geq (d/2^b)$$

and so $F(2^{b-\ell}c_2) > F(2^{b-\ell}(c_1+1) - 1)$, a contradiction.

Also, we claim that for each value of c , there is at most one value k such that $G_k = \{c, c+1\}$. For, if not, then there would be values $k_1 < k_2$ with

$$F(2^{b-\ell}(c+1)) \leq k_1 \leq F(2^{b-\ell}(c+2) - 1),$$

$$F(2^{b-\ell}c) \leq k_2 \leq F(2^{b-\ell}(c+1) - 1).$$

But note then that

$$F(2^{b-\ell}(c+1) - 1) \geq k_2 > k_1 \geq F(2^{b-\ell}(c+1)) > F(2^{b-\ell}(c+1)) \geq F(2^{b-\ell}(c+1) - 1),$$

a contradiction.

Thus, for each $c \in \{0, \dots, 2^\ell - 1\}$, let us define W_c to be the set of values $k \in \{0, \dots, d-1\}$ such that $G_k = \{c, c+1\}$. We have just showed that $|W_c| \leq 1$.

Now consider the random experiment of assigning independent Bernoulli-1/2 values to the low-order $b - \ell$ bit-levels of u . Define the random variable Z_c to be the number of vertices in S_c assigned a value $k \in W_c$. (If $W_c = \emptyset$, then $Z_c = 0$ necessarily.) In order for e to be rainbow, every $c \in \mathcal{M}_\ell$ must have $Z_c \in \{0, 1\}$.

For any integers $0 \leq c_0 < c_1 \leq 2^\ell$ and values $z_0, z_1 \in \{0, 1\}$, let us thus define the function g by

$$g(c_0, c_1, z_0, z_1) =$$

$$P(\text{the vertices in } S_{c_0}, S_{c_0+1}, \dots, S_{c_1-1}, S_{c_1} \text{ receive distinct colors and } Z_{c_1} = z_1 \mid Z_{c_0-1} = z_0).$$

The overall probability that the random experiment results in a rainbow coloring of e is given by $g(0, 2^\ell, 0, 0)$. With a little thought, one can see that g satisfies the recurrence:

$$g(c_0, c_1, z_0, z_1) = g(c_0, c_2 - 1, z_0, 0)g(c_2, c_1, 0, z_1) + g(c_0, c_2 - 1, z_0, 1)g(c_2, c_1, 1, z_1)$$

for $c_2 = (c_0 + c_1)/2$.

If $S_c = \emptyset$, then the value c is not relevant to this calculation; thus, during this calculation, we can skip all such entries. As there are at most d values of c with $S_c \neq \emptyset$, we can recursively compute $g(0, 2^\ell, 0, 0)$ using $\text{poly}(db)$ processors and using $\tilde{O}(\log mn)$ time. (The base cases can be computed using simple arithmetic as functions of $|S_c|$.)

We next discuss how to modify this to graded PEO. Here, the top $\ell - 1$ bits of each y_v are completely known, while the lowest-order bit is in $\{0, 1, ?\}$. Now suppose we want to calculate $g(0, 1, z_0, z_1)$; in this case, some vertices are known to correspond to the sets S_0, S_1 and some vertices

(for which bit at level ℓ is unspecified) have a $1/2$ probability of going into S_0 and a $1/2$ probability of going into S_1 . We can integrate over the sizes of S_0 and S_1 (which are now binomial random variables), and use the above formulas to calculate $g(0, 1, z_0, z_1)$. $\square$

B FULL PROOF OF THEOREM 5.2

We assume here that $\Delta \leq C \log n$ for some sufficiently large constant C , and that k is larger than any needed constant. The case of large Δ has already been discussed.

We will build the coloring gradually over stages $i = 0, \dots, r+1$; at stage i , the vertices have a t_i -coloring, in which every vertex has at most Δ_i neighbors of its own color class. At stage i , and for any integer $\ell \in \{1, \dots, t_i\}$, let us define $G_{i,\ell}$ to be the subgraph of G induced on vertices with color ℓ . So $G_{i,\ell}$ has maximum degree Δ_i . We will apply Lemma 5.1 with parameter j_i to each $G_{i,\ell}$; this is done in parallel across all values of ℓ . This gives $t_{i+1} \leq 2^{j_i} t_i$, and at the end of this process, we thereby obtain a Δ_{r+1} -defective coloring with t_{r+1} colors.

We need to define the sequence of values j_i, Δ_i which will be valid for the degree splitting procedure. We do so recursively by setting $\Delta_0 = \Delta$, and

$$\Delta_{i+1} = (\Delta_i/2^{j_i})(1 + K\sqrt{2} \log^{-1/2} \Delta_i) \quad j_i = \begin{cases} \lceil \log_2(\Delta_i / \log^2 \Delta_i) \rceil & i < r \\ \lceil \log_2(\Delta_i/k) \rceil & i = r, \end{cases}$$

where r is a parameter to be determined.

Let us verify that these parameters satisfy the conditions of Lemma 5.1; specifically, we show by induction on i that G_i has maximum degree Δ_i and $2^{j_i} \leq \frac{\Delta_i}{k \log \Delta_i}$ for $i = 0, \dots, r$.

In order to carry out this analysis, let us define a sequence of real numbers by

$$b_0 = \Delta, \quad b_{i+1} = \frac{1}{2} \log^2 b_i.$$

Let r be the largest integer with $b_r \geq k$; we stop this procedure at stage $r+1$. We easily see that $r \leq O(\log^* \Delta)$. We claim that for $i = 0, \dots, r$ we have $k \leq b_i \leq \Delta_i \leq 4b_i$, and we show this by induction on i . The bound $b_i \geq k$ is immediate from the definition of r . The bound on Δ_0 is immediate. For $i < r$, the lower bound is shown by

$$\Delta_{i+1} \geq \Delta_i/2^{j_i} \geq \Delta_i/2^{\lceil \log_2(\Delta_i / \log^2 \Delta_i) \rceil} \geq \Delta_i/2^{1+\log_2(\Delta_i / \log^2 \Delta_i)} = \frac{1}{2} \log^2 \Delta_i \geq \frac{1}{2} \log^2 b_i = b_{i+1}.$$

For the upper bound, we have for $b_i \geq k$ and k sufficiently large,

$$\begin{aligned} \Delta_{i+1} &= (\Delta_i/2^{j_i})(1 + K\sqrt{2} \log^{-1/2} \Delta_i) \leq (1 + K\sqrt{2} \log^{-1/2} \Delta_i) \log^2 \Delta_i \\ &\leq (1 + K\sqrt{2} \log^{-1/2} b_i) \log^2(4b_i) \leq 1.01 \log^2(b_i) \leq 4b_{i+1}. \end{aligned}$$

We can now show that $2^{j_i} \leq \frac{\Delta_i}{K \log \Delta_i}$ holds. For $i < r$, we have

$$2^{j_i} \leq \frac{2\Delta_i}{\log^2 \Delta_i} \leq \frac{\Delta_i}{\log \Delta_i} \times \frac{2}{\log \Delta_i} \leq \frac{\Delta_i}{\log \Delta_i} \times \frac{2}{k} \leq \frac{\Delta_i}{K \log \Delta_i}.$$

For $i = r$, we note that $b_{r+1} \leq k$ and so

$$2^{j_r} \leq 2\Delta_r/k \leq (8b_r)/k \leq 16b_r/\log^2 b_r \leq 16\Delta_r/\log^2 \Delta_r \leq \Delta_r/(K \log \Delta_r).$$

So we can apply Lemma 5.1. For $i < r$, the definition of j_i gives

$$2^{j_i}/\Delta_i \leq \frac{2}{\log^2 \Delta_i}$$

and therefore Lemma 5.1 shows that the graph $G_{i+1,\ell}$ has maximum degree at most

$$(\Delta_i/2^{j_i}) \left(1 + K\sqrt{(2^{j_i}/\Delta_i) \log \Delta_i} \right) \leq (\Delta_i/2^{j_i}) \left(1 + K\sqrt{\frac{2}{\log \Delta_i}} \right) = \Delta_{i+1}.$$

Similarly, for $i = r$, Lemma 5.1 ensures that $G_{r+1,\ell}$ have maximum degree

$$\Delta_{r+1} \leq (\Delta_r/2^{j_r})(1 + O(\log^{-1/2} \Delta_r)) \leq k(1 + O(\log^{-1/2} \Delta_r)) \leq O(k).$$

Thus, the overall coloring we obtain is indeed $O(k)$ -defective. Our next task is to count the number of colors used. Let us define $a_i = \Delta_i t_i$. We want to show that $t_{r+1} \leq O(\Delta/k)$. As $\Delta_{r+1} \geq k/2$, it suffices to show that $a_{r+1} \leq O(\Delta)$. The recursive formulas for Δ_i and t_i give

$$a_{i+1} \leq (\Delta_i/2^{j_i})(1 + K\sqrt{2} \log^{-1/2} \Delta_i) \times (t_i 2^{j_i}) \leq a_i(1 + K\sqrt{2} \log^{-1/2} \Delta_i) \leq a_i(1 + K\sqrt{2} \log^{-1/2} b_i).$$

Therefore,

$$a_{r+1} \leq a_0 \prod_{i=0}^r (1 + K\sqrt{2} \log^{-1/2} b_i) \leq \Delta e^{O(\sum_{i=0}^r \log^{-1/2} b_i)} \leq O(\Delta),$$

where the last inequality follows by noting that the sequence $\log b_i$ is decreasing super-exponentially.

We finish by calculating the complexity of the algorithm. Each iteration i requires $(\Delta_i \log n + \log^2 n, \text{poly}(n))$ quasi-complexity. We see easily that $\Delta_i \leq O(\Delta) \leq O(\log n)$, so this is $\tilde{O}(\log^2 n)$ time per iteration. There are $r \leq O(\log^* \Delta)$ iterations, giving a total runtime again of $\tilde{O}(\log^2 n)$.

ACKNOWLEDGMENTS

The authors thank Aravind Srinivasan, Vance Faber, and anonymous referees for helpful comments and discussion.

REFERENCES

- [1] N. Alon and A. Srinivasan. 1996. Improved parallel approximation of a class of integer programming problems. In *Proceedings of the 23rd International Colloquium on Automata, Languages, and Programming (ICALP'96)*, 562–573.
- [2] N. Alon, L. Babai, and A. Itai. 1986. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms* 7 (1986), 567–583.
- [3] B. Berger and J. Rompel. 1991. Simulating $(\log^c n)$ -wise independence in NC. *Journal of the ACM* 38, 4 (1991), 1026–1046.
- [4] A. Borodin and S. Cook. 1982. A time-space tradeoff for sorting on a general sequential model of computation. *SIAM Journal on Computing* 11, 2 (1982), 287–297.
- [5] K. Chandrasekaran, N. Goyal, and B. Haeupler. 2013. Deterministic algorithms for the Lovász local lemma. *SIAM Journal on Computing* 42, 6 (2013), 2132–2155.
- [6] S. Chari, P. Rohatgi, and A. Srinivasan. 2000. Improved algorithms via approximations of probability distributions. *Journal of Computer and System Sciences* 61, 1 (2000), 81–107.
- [7] M. Fischer and M. Ghaffari. 2017. Sublogarithmic distributed algorithms for Lovász local lemma, and the complexity hierarchy. In *Proceedings of the 31st International Symposium on Distributed Computing (DISC'17)*, 18:1–18:16.
- [8] U. Feige, M. M. Halldórsson, G. Kortsarz, and A. Srinivasan. 2002. Approximating the domatic number. *SIAM Journal on Computing* 32, 1 (2002), 172–195.
- [9] B. Haeupler and D. G. Harris. 2017. Parallel algorithms and concentration bounds for the Lovász Local Lemma via witness-DAGs. *ACM Transactions on Algorithms* 13, 4 (2017), Article 53.
- [10] D. Harris. 2017. Deterministic parallel algorithms for bilinear objective functions. arxiv:1711-08494.
- [11] M. Luby. 1993. Removing randomness in parallel computation without a processor penalty. *Journal of Computer and System Sciences* 47, 2 (1993), 250–286.
- [12] R. Moser and G. Tardos. 2010. A constructive proof of the general Lovász Local Lemma. *Journal of the ACM* 57, 2 (2010), 11:1–11:15.
- [13] K. Mulmuley. 1987. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. *Combinatorica* 7, 1 (1987), 101–104.

- [14] J. Naor and M. Naor. 1993. Small-bias probability spaces: efficient construction and applications, *SIAM Journal of Computing* 22, 4 (1993), 835–856.
- [15] L. Schulman. 1992. Sample spaces uniform on neighborhoods. In *Proceedings of the 24th ACM Symposium on Theory of Computing (STOC'92)*, 17–25.
- [16] D. Sivakumar. 2002. Algorithmic derandomization via complexity theory. In *Proceedings of the 34th ACM Symposium on Theory of Computing (STOC'02)*, 619–626.

Received June 2017; revised May 2018; accepted May 2018