

Beyond the Lovász Local Lemma: Point to Set Correlations and Their Algorithmic Applications *

Dimitris Achlioptas [†]
University of California Santa Cruz
optas@cs.ucsc.edu

Fotis Iliopoulos [‡]
University of California Berkeley
fotis.iliopoulos@berkeley.edu

Alistair Sinclair [§]
University of California Berkeley
sinclair@berkeley.edu

Abstract

Following the groundbreaking algorithm of Moser and Tardos for the Lovász Local Lemma (LLL), there has been a plethora of results analyzing local search algorithms for various constraint satisfaction problems. The algorithms considered fall into two broad categories: *resampling* algorithms, analyzed via different algorithmic LLL conditions; and *backtracking* algorithms, analyzed via entropy compression arguments. This paper introduces a new convergence condition that seamlessly handles resampling, backtracking, and *hybrid* algorithms, i.e., algorithms that perform both resampling and backtracking steps. Unlike previous work on the LLL, our condition replaces the notion of a dependency or causality graph by quantifying *point-to-set correlations* between bad events. As a result, our condition simultaneously: (i) captures the most general algorithmic LLL condition known as a special case; (ii) significantly simplifies the analysis of entropy compression applications; (iii) relates backtracking algorithms, which are conceptually very different from resampling algorithms, to the LLL; and most importantly (iv) allows for the analysis of hybrid algorithms, which were outside the scope of previous techniques. We give several applications of our condition, including a new hybrid vertex coloring algorithm that extends the recent breakthrough result of Molloy for coloring triangle-free graphs to arbitrary graphs.

From a technical perspective, our main new insight is that LLL-inspired convergence arguments for local search algorithms can be viewed as methods for bounding the *spectral radius* of an associated transition matrix.

* A preliminary version of this paper appeared in the *60th Annual IEEE Symposium on Foundations of Computer Science*.

[†]Research supported by NSF grant CCF-1514434

[‡]Research supported by NSF grant CCF-1514434 and the Onassis Foundation

[§]Research supported by NSF grants CCF-1514434 and CCF-1815328

1 Introduction

Numerous problems in computer science and combinatorics can be formulated as searching for objects that lack certain bad properties, or “flaws”. For example, constraint satisfaction problems such as satisfiability and graph coloring can be seen as searching for objects (truth assignments, colorings) that are *flawless*, in the sense that they do not violate any constraint.

The *Lovász Local Lemma (LLL)* [25] is a powerful tool for proving the *existence* of flawless objects that has had far-reaching consequences in computer science and combinatorics [9, 53]. Roughly speaking, it asserts that, given a collection of a bad events in a probability space, if all of them are individually not too likely, and independent of most other bad events, then the probability that none of them occurs is strictly positive; hence a flawless object (i.e., an elementary event that does not belong in any bad event) exists. For example, the LLL implies that every k -CNF formula in which each clause shares variables with fewer than $2^k/e$ other clauses is satisfiable. Remarkably, this is tight [28].

In seminal work, Moser [54] and Moser-Tardos [55] showed that a simple local search algorithm can be used to make the LLL constructive for product probability spaces, i.e., spaces where elementary events correspond to sequences of independent coin-flips. For example, the Moser-Tardos algorithm for satisfiability starts at a uniformly random truth assignment and, as long as violated clauses exist, selects any such clause and resamples the values of all its variables uniformly at random. Following this work, a large amount of effort has been devoted to making different variants of the LLL constructive [17, 47, 48, 58], and to analyzing sophisticated *resampling* algorithms that extend the Moser-Tardos techniques to non-product probability spaces [2, 3, 4, 38, 40, 51, 43]. Indeed, intimate connections have been established between resampling algorithms and the so-called “lopsided” versions of the LLL: see [3, 40, 43] for more details.

In earlier groundbreaking work aimed at making the LLL algorithmic for k -SAT, Moser [54] introduced the *entropy compression* method. This method has since been used to analyze *backtracking* algorithms, e.g., [24, 27, 30, 31, 32, 57]. These natural and potentially powerful local search algorithms have a very different flavor from resampling algorithms: instead of maintaining a complete value assignment (object) and repeatedly modifying it until it becomes satisfying, they operate on *partial non-violating* assignments, starting with the empty assignment, and try to extend to a complete, satisfying assignment. To do this, at each step they assign a (random) value to a currently unassigned variable; if this leads to the violation of one or more constraints, they backtrack to a partial non-violating assignment by unassigning some set of variables (typically including the last assigned variable).

While there have been efforts to treat certain classes of backtracking algorithms systematically [27, 31], the analysis of such algorithms in general still requires *ad hoc* technical machinery. Moreover, there is no known connection between backtracking algorithms and any known LLL condition, either existential or algorithmic. The main reason for this is that backtracking steps induce non-trivial correlations among bad events, which typically result in very dense dependency graphs that are not amenable to currently known LLL conditions.

The main contribution of this paper is to introduce a new technique for analyzing *hybrid* algorithms, i.e., algorithms that (potentially) use both resampling and backtracking steps. Such algorithms combine the advantages of both approaches by using resampling to explore the state space, while detecting and backing away from unfavorable regions using backtracking steps. To analyze these algorithms, we prove a new algorithmic LLL condition which, unlike previous versions, replaces the notion of a dependency or causality graph by quantifying *point-to-set correlations* between bad events. Notably, our new condition captures the most general algorithmic LLL condition known so far and, moreover, unifies the analysis of all entropy compression applications, connecting backtracking algorithms to the LLL in the same fashion that existing analyses connect resampling algorithms to the LLL.

Our main technical insight is that essentially all existing LLL-inspired convergence arguments for local search algorithms can be viewed as methods for bounding the *spectral radius* of an associated transition

matrix by decomposing it into a sum of sparse matrices, one for each flaw. Our new condition hinges on a more refined decomposition that allows us to take advantage of higher-dimensional information.

A new coloring algorithm. Our main application is a new vertex coloring algorithm inspired by the recent breakthrough result of Molloy [51], who proved that any triangle-free graph of maximum degree Δ can be colored using $(1 + o(1))\Delta / \ln \Delta$ colors; this improves the celebrated result of Johansson [44] while, at the same time, dramatically simplifying its analysis. We generalize Molloy’s result by establishing a bound on the chromatic number of *arbitrary* graphs, as a function of the maximum number of triangles in the neighborhood of a vertex, and giving an algorithm that produces such a coloring.

Theorem 1.1 (Informal Statement). *Let G be any graph with maximum degree Δ in which the neighbors of every vertex span at most $T \geq 0$ edges between them. For every $\epsilon > 0$, if $\Delta \geq \Delta_\epsilon$ and $T \lesssim \Delta^{2\epsilon}$ then*

$$\chi(G) \leq (1 + \epsilon) \frac{\Delta}{\ln \Delta - \frac{1}{2} \ln(T + 1)}, \tag{1}$$

and such a vertex coloring can be found efficiently. (Here \lesssim hides logarithmic factors.) Moreover, the theorem holds for any $T \geq 0$ if the leading constant $(1 + \epsilon)$ is replaced by $(2 + \epsilon)$.

Importantly, as explained in Section 2.2, the bound (1) *matches* the algorithmic barrier for random graphs [1]. This implies that any improvement on the guarantee of our algorithm for $T \lesssim \Delta^{2\epsilon}$ would amount to an unexpected breakthrough in random graph theory. (Random graphs are only informative in the regime $T \lesssim \Delta^{2\epsilon}$.) For arbitrary graphs our bound is within a factor of 4 of the chromatic number, improving upon a classical result of Alon, Krivelevich and Sudakov [8] who showed (1) with an unspecified (large) leading constant.

At the heart of Theorem 1.1 is a hybrid local search algorithm which we analyze using the techniques introduced in this paper. Molloy’s result (and resampling algorithms based on it) breaks down immediately in the presence of triangles; the key to our algorithm is to allow backtracking steps in order to avoid undesirable portions of the search space. We discuss our coloring result and its optimality further in Section 2.2.

Applications to backtracking algorithms. Besides our coloring application, we give three representative applications of our techniques applied to pure backtracking algorithms.

Recently, Bissacot and Doin [15] showed that backtracking algorithms can make LLL applications in the variable setting constructive, using the entropy compression method. However, their result applies only to the uniform measure and their algorithms are relatively complicated. Our new algorithmic condition makes applications of the LLL in the variable setting [55], with any measure, constructive via a single, simple backtracking algorithm, i.e., an algorithm of very different flavor from the Moser-Tardos resampling algorithm. For example, in the case of k -SAT the algorithm takes the following form:

Randomized DPLL with single-clause backtracking

while unassigned variables exist **do**

Assign to the lowest indexed unassigned variable x a value $v \in \{0, 1\}$ with probability p_x^v

if one or more clauses become violated **then**

Unassign all k variables in the lowest indexed violated clause

We then show that applying our condition to the algorithm of Esperet and Parreau [27] for *acyclic edge coloring* recovers, in a simple, black-box fashion, the same bound of 4Δ as their highly non-trivial, problem-specific analysis via entropy compression, while guaranteeing an improved running time bound.

Finally, we make constructive in an effortless manner an existential result of Bernshteyn [12] showing improved bounds for the acyclic chromatic index of graphs that do not contain an arbitrary bipartite graph H .

We present our results on backtracking algorithms in Section 5.

1.1 A new algorithmic LLL condition

To state our new algorithmic LLL condition, we need some standard terminology. Let Ω be a finite set and let $F = \{f_1, f_2, \dots, f_m\}$ be a collection of subsets of Ω , each of which will be referred to as a “flaw.” Let $\bigcup_{i \in [m]} f_i = \Omega^*$. For example, for a given CNF formula on n variables with clauses c_1, \dots, c_m , we take $\Omega = \{0, 1\}^n$ to be the set of all possible variable assignments, and f_i the set of assignments that fail to satisfy clause c_i . Our goal is to find an assignment in $\Omega \setminus \Omega^*$, i.e., a satisfying (“flawless”) assignment. Since we will be interested in algorithms which traverse the set Ω we will also refer to its elements as *states*.

We consider algorithms which, in each flawed state $\sigma \in \Omega^*$, choose a flaw f_i present in σ , i.e., $f_i \ni \sigma$, and attempt to leave (or “fix”) f_i by moving to a new state τ selected with probability $\rho_i(\sigma, \tau)$; we refer to such an attempt as *addressing* f_i . We make minimal assumptions about how the algorithm chooses which flaw to address at each step; e.g., it will be enough for the algorithm to choose the flaw with lowest index according to some fixed permutation. We say that a transition $\sigma \rightarrow \tau$, made to address flaw f_i , *introduces* flaw f_j if $\tau \in f_j$, and either $\sigma \notin f_j$ or $j = i$.

For an arbitrary state $\tau \in \Omega$, flaw f_i , and set of flaws S , let

$$\text{In}_i^S(\tau) := \{\sigma \in f_i : \text{the set of flaws introduced by the transition } \sigma \rightarrow \tau \text{ includes } S\}. \quad (2)$$

For any fixed probability distribution $\mu > 0$ on Ω (either inherited from an application of the probabilistic method, as in the classical LLL, or introduced by the algorithm designer), we define the *charge* of the pair (i, S) with respect to μ to be

$$\gamma_i^S := \max_{\tau \in \Omega} \left\{ \frac{1}{\mu(\tau)} \sum_{\sigma \in \text{In}_i^S(\tau)} \mu(\sigma) \rho_i(\sigma, \tau) \right\}. \quad (3)$$

That is, the charge γ_i^S is an upper bound on the ratio between the ergodic flow into a state via transitions that introduce every flaw in S (and perhaps more), and the probability of the state under μ .

Our condition may now be stated informally as follows:

Theorem 1.2. *If there exist positive real numbers $\{\psi_i\}_{i=1}^m$ such that, for all $i \in [m]$,*

$$\frac{1}{\psi_i} \sum_{S \subseteq [m]} \gamma_i^S \prod_{j \in S} \psi_j < 1, \quad (4)$$

then a local search algorithm as above reaches a flawless object quickly with high probability.

The phrase “quickly with high probability” essentially means that the running time has expectation linear in the number of flaws and an exponential tail; we spell this out more formally in Section 2.

A key feature of Theorem 1.2 is the absence of a causality/dependency graph, present in all previous LLL conditions. This is because considering *point-to-set correlations*, i.e., how each flaw interacts with every other set of flaws, frees us from the traditional view of dependencies between individual events. In our condition, every flaw may interact with every other flaw, as long as the interactions are sufficiently weak. Notably, this is achieved without any loss when specialized to the traditional setting of a causality/dependency graph. To see this, note that if S contains any flaw that is never introduced by addressing flaw f_i , then $\gamma_i^S = 0$. Thus, in the presence of a causality/dependency graph, the only terms contributing to the summation in (4) are those that correspond to subsets of the graph neighborhood of flaw f_i , recovering the traditional setting.

Besides relaxing the traditional notion of dependence, our condition is also quantitatively more powerful, even in the traditional setting. To get a feeling for this, we observe that the previously most powerful

algorithmic LLL condition, due to Achlioptas, Iliopoulos and Kolmogorov [3], can be derived from (4) by replacing γ_i^S by γ_i^\emptyset for every S (and restricting S to subsets of the neighborhood of flaw f_i , as discussed in the previous paragraph). Since the charges γ_i^S are decreasing in S , replacing γ_i^\emptyset with γ_i^S can lead to a dramatic improvement. For example, if the flaws in S are never introduced simultaneously when addressing flaw f_i , then $\gamma_i^S = 0$ and S does not contribute to the sum in (4); in contrast, S contributes γ_i^\emptyset , i.e., the maximum possible charge, to the corresponding sum in [3]. For a more detailed discussion of how our new condition subsumes existing versions of the LLL see Appendix C.

A natural question is whether Theorem 1.2 can be improved by replacing the word “includes” with the word “equals” in (2), thus shrinking the sets $\text{In}_i^S(\tau)$. The short answer is “No,” i.e., such a change invalidates the theorem. The reason for this is that in resampling algorithms we must allow for the possibility that flaws introduced when addressing f_i may later be fixed “collaterally,” i.e., as the result of addressing other flaws rather than by being specifically addressed by the algorithm. While it may seem that such collateral fixes cannot possibly be detrimental, they are problematic from an analysis perspective as they can potentially increase the intensity of correlations between flaw f_i and S . Perhaps more convincingly, tracking collateral fixes and taking them into account also appears to be a bad idea in practice [60, 61, 10, 11]: for example, local search satisfiability algorithms that select which variable to flip (among those in the targeted violated clause) based *only* on which clauses will become violated, fare much better than algorithms that weigh this damage against the benefit of the collaterally fixed clauses.

Motivated by the above considerations, we introduce the notion of *primary* flaws. These are flaws which, once present, can only be eradicated by being addressed by the algorithm, i.e., they cannot be fixed collaterally. Primary flaws allow us to change the definition of the sets $\text{In}_i^S(\tau)$ in the desired direction. Specifically, say that a set of flaws T *covers* a set of flaws S if:

1. the set of primary flaws in T *equals* the set of primary flaws in S ; and
2. the set of non-primary flaws in T *includes* the set of non-primary flaws in S .

In other words, we demand equality at least for the primary flaws.

Theorem 1.3. *Theorem 1.2 continues to hold if $\text{In}_i^S(\tau)$ is redefined by replacing “includes” by “covers” in equation (2).*

The notion of primary flaws is one of our main conceptual contributions. Crucially for our applications, backtracking steps always introduce only primary flaws and thus, for such steps, we achieve an ideal level of control. The full version of our new algorithmic LLL condition, incorporating primary flaws, is spelled out formally in Theorem 2.4 in Section 2.

1.2 Technical overview: the Lovász Local Lemma as a spectral condition

We conclude this introduction by sketching the techniques we use to prove our convergence criterion. As mentioned earlier, our main insight is to interpret LLL-inspired convergence arguments for local search algorithms as methods for bounding the *spectral radius* of an associated matrix.

As above, let Ω be a (large) finite set of states and let $\Omega^* \subseteq \Omega$ be the “bad” part of Ω , comprising the flawed states. Imagine a particle trying to escape Ω^* by following a Markov chain on Ω with transition matrix P . Our task is to develop conditions under which the particle eventually escapes, thus establishing in particular that $\Omega^* \neq \Omega$. Letting A be the $|\Omega^*| \times |\Omega^*|$ submatrix of P that corresponds to transitions from Ω^* to Ω^* , and B the submatrix that corresponds to transitions from Ω^* to $\Omega \setminus \Omega^*$, we may, after a suitable permutation of its rows and columns, write P as:

$$P = \left[\begin{array}{c|c} A & B \\ \hline 0 & I \end{array} \right].$$

Here I is the identity matrix, since we assume that the particle stops after reaching a flawless state.

Let $\theta = [\theta_1 \mid \theta_2]$ be the row vector corresponding to the probability distribution of the starting state, where θ_1 and θ_2 are the vectors that correspond to states in Ω^* and $\Omega \setminus \Omega^*$, respectively. Then, the probability that after t steps the particle is still inside Ω^* is exactly $\|\theta_1 A^t\|_1$. Therefore, for any initial distribution θ , the particle escapes Ω^* if and only if the spectral radius, $\rho(A)$, of A is strictly less than 1. Moreover, the rate of convergence is dictated by $1 - \rho(A)$. Unfortunately, since A is huge and defined implicitly by an algorithm, the magnitude of its largest eigenvalue, $\rho(A)$, is not readily available.

To sidestep the inaccessibility of the spectral radius $\rho(A)$, one can instead bound some *operator norm* $\|\cdot\|$ of A and appeal to the fact that $\rho(A) \leq \|A\|$ for any such norm. Moreover, instead of bounding an operator norm of A itself, one often first performs a “change of basis” $A' = MAM^{-1}$ for some invertible matrix M and then bounds $\|A'\|$, justified by the fact that $\rho(A) = \rho(A') \leq \|A'\|$. The purpose of the change of basis is to cast A “in a good light” in the eyes of the chosen operator norm, in the hope of minimizing the cost of replacing the spectral norm with the operator norm.

As we explain in Appendix B, essentially all known analyses of LLL-inspired local search algorithms can be recast in the above framework of matrix norms. More significantly, this viewpoint allows the extension of such analyses to a wider class of algorithms. In Section 3, we will use this linear-algebraic machinery to prove our new convergence condition. The role of the measure μ will be reflected in the change of basis M , while the charges γ_i^S will correspond to norms $\|MA_i^S M^{-1}\|_1$, where the A_i^S are submatrices of the transition matrix A .

1.3 Follow-up work

Following the appearance of the conference version of this paper, Davies et al. [21] used our new algorithmic LLL condition in their general framework for coloring locally sparse graphs. Among other applications, they generalize our Theorem 2.5 so that its conclusion holds under an analogous bound for the local k -cycle-density of the input graph, for any fixed $k \geq 3$. (Thus, when $k = 3$ their result reduces to Theorem 2.5.) In a companion paper [22], the same authors also managed to improve the bound in our Theorem 2.6 by a factor of 2, but this result only implies an algorithm whose running time is exponential in the maximum degree of the graph. An interesting feature of these papers is that they apply the LLL in sophisticated probability spaces that are inspired by so-called “hardcore distributions” on independent sets.

A number of papers in the algorithmic LLL literature [33, 36, 39, 42, 49] analyze properties of local search algorithms besides fast convergence (such as parallelization, bounds on the entropy of the output distribution, etc.). It is known that such properties cannot be established for general local search algorithms [40], but it has been observed [49, 42] that they do hold for a large class of so-called *commutative* algorithms. Very recently, and inspired by our work in this paper, Harris, Iliopoulos and Kolmogorov [20] introduced a very natural and more general notion of commutativity (essentially matrix commutativity) that allows them to show a number of new refined properties of LLL-inspired local search algorithms with significantly simpler proofs.

Finally, Kolmogorov [50] extended our new algorithmic LLL condition, showing that it can be combined with well-known improvements of the standard Lovász Local Lemma for locally dense dependency graphs [16, 3]. However, we do not currently know a natural setting where this new condition applies, and Kolmogorov mentions that he views the significance of his contribution as mainly theoretical. (See also Appendix C for further discussion.)

2 Statement of results

2.1 A new algorithmic LLL condition

Below we state our main result, which is a formal version of Theorem 1.3 discussed in Section 1.1.

Recall that Ω is a finite set, $F = \{f_1, f_2, \dots, f_m\}$ is a collection of subsets of Ω which we refer to as *flaws*, and $\bigcup_{i \in [m]} f_i = \Omega^*$. Our goal is to find a flawless object, i.e., an object in $\Omega \setminus \Omega^*$. For a state σ , we denote by $U(\sigma) = \{j \in [m] : f_j \ni \sigma\}$ the set of (indices of) flaws present in σ . (Here and elsewhere, we shall blur the distinction between flaws and their indices.) We consider algorithms that start in a state sampled from a probability distribution θ and, in each flawed state $\sigma \in \Omega^*$, choose a flaw $f_i \in U(\sigma)$, and attempt to leave (“fix”) f_i by moving to a new state τ selected with probability $\rho_i(\sigma, \tau)$. We refer to an attempt to fix a flaw, successful or not, as *addressing* it. We say that a transition $\sigma \rightarrow \tau$, made to address flaw f_i , *introduces* flaw $f_j \in U(\tau)$ if $f_j \notin U(\sigma)$ or if $j = i$. (Thus, a flaw (re)introduces itself when a transition fails to address it.)

Recall that θ denotes the probability distribution of the starting state. We denote by $\text{Span}(\theta)$ the set of flaw indices that may be present in the initial state, i.e., $\text{Span}(\theta) = \bigcup_{\sigma \in \Omega: \theta(\sigma) > 0} U(\sigma)$.

Let π be an arbitrary permutation over $[m]$. We say that an algorithm follows the π -*strategy* if at each step the flaw it chooses to address is the one corresponding to the element of $U(\sigma)$ of lowest index according to π .

We now formalize the definitions of primary flaws and charges introduced informally in the introduction.

Definition 2.1. A flaw f_i is *primary* if for every $\sigma \in f_i$ and every $j \neq i$, addressing f_j at σ always results in some $\tau \in f_i$, i.e., f_i is never eradicated collaterally. For a given set $S \subseteq [m]$, we write S^P and S^N to denote the indices that correspond to primary and non-primary flaws in S , respectively.

Definition 2.2. We say that a set of flaws T covers a set of flaws S if $T^P = S^P$ and $T^N \supseteq S^N$.

Definition 2.3. For a state $\tau \in \Omega$, flaw f_i , and set of flaws S , let

$$\text{In}_i^S(\tau) = \{\sigma \in f_i : \text{the set of flaws introduced by the transition } \sigma \rightarrow \tau \text{ covers } S\}.$$

Let $\mu > 0$ be an arbitrary measure on Ω . For every $i \in [m]$ and $S \subseteq [m]$, the *charge* of (i, S) with respect to μ is,

$$\gamma_i^S = \max_{\tau \in \Omega} \left\{ \frac{1}{\mu(\tau)} \sum_{\sigma \in \text{In}_i^S(\tau)} \mu(\sigma) \rho_i(\sigma, \tau) \right\}. \quad (5)$$

We now state the formal version of our main result, Theorem 1.2 of the introduction.

Theorem 2.4 (Main Result). *If there exist positive real numbers $\{\psi_i\}_{i \in [m]}$ such that, for every $i \in [m]$,*

$$\zeta_i := \frac{1}{\psi_i} \sum_{S \subseteq [m]} \gamma_i^S \prod_{j \in S} \psi_j < 1, \quad (6)$$

then, for every permutation π over $[m]$, the probability that an algorithm following the π -strategy fails to reach a flawless state within $(T_0 + s)/\delta$ steps is 2^{-s} , where $\delta = 1 - \max_{i \in [m]} \zeta_i$, and

$$T_0 = \log_2 \mu_{\min}^{-1} + m \log_2 \left(\frac{1 + \psi_{\max}}{\psi_{\min}} \right),$$

with $\mu_{\min} = \min_{\sigma \in \Omega} \mu(\sigma)$, $\psi_{\max} = \max_{i \in [m]} \psi_i$ and $\psi_{\min} = \min_{i \in [m]} \psi_i$.

Remark 2.1. In typical applications, μ and $\{\psi_i\}_{i \in [m]}$ are such that $T_0 = O(\log |\Omega| + m)$ and the sum in (6) is easily computable, as $\gamma_i^S = 0$ for the vast majority of subsets S .

Remark 2.2. For any fixed permutation π , the charges γ_i^S can be reduced by removing from $\text{In}_i^S(\tau)$ every state for which i is not the lowest indexed element of $U(\sigma)$ according to π .

Remark 2.3. Theorem 2.4 also holds for algorithms using flaw choice strategies other than π -strategies. We discuss some such strategies in Section 3.4. However, there is good reason to expect that it does not hold for arbitrary flaw choice strategies (see [49]).

Finally, we state a refinement of our running time bound that will be important in order to get the best convergence guarantees in the applications of pure backtracking algorithms in Section 5.

Remark 2.4. The upper bound on T_0 in Theorem 2.4 can be replaced by the more refined bound:

$$T_0 = \log_2 \left(\max_{\sigma \in \Omega} \frac{\theta(\sigma)}{\mu(\sigma)} \right) + \log_2 \left(\sum_{S \subseteq \text{Span}(\theta)} \prod_{j \in S} \psi_j \right) + \log_2 \left(\max_{S \subseteq [m]} \frac{1}{\prod_{j \in S} \psi_j} \right).$$

Moreover, if (as in pure backtracking algorithms) every flaw is primary, and (as is typical in pure backtracking algorithms) every flaw is present in the initial state, and if $\psi_i \in (0, 1]$ for all i , then $T_0 = \log_2 \mu_{\min}^{-1}$.

2.2 Application to graph coloring

In graph coloring one is given a graph $G = (V, E)$ and the goal is to find a mapping of V to a set of q colors so that no edge in E is monochromatic. The *chromatic number*, $\chi(G)$, of G is the smallest integer q for which this is possible. Given a set \mathcal{L}_v of colors for each vertex v (called a *list*), a list-coloring maps each $v \in V$ to a color in \mathcal{L}_v so that no edge in E is monochromatic. A graph is q -list-colorable if it has a list-coloring no matter how one assigns a list of q colors to each vertex. The *list chromatic number*, $\chi_\ell(G)$, is the smallest q for which G is q -list-colorable. Clearly $\chi_\ell(G) \geq \chi(G)$. A celebrated result of Johansson [44] established that there exists a large constant $C > 0$ such that every *triangle-free* graph with maximum degree $\Delta \geq \Delta_0$ can be list-colored using $C\Delta / \ln \Delta$ colors. Very recently, using the entropy compression method, Molloy [51] improved Johansson's result, replacing C with $(1 + \epsilon)$ for any $\epsilon > 0$ and all $\Delta \geq \Delta_\epsilon$. (Soon thereafter, Bernshteyn [13] established the same bound for the list chromatic number, non-constructively, via the Lopsided LLL, and Iliopoulos [42] showed that the algorithm of Molloy can be analyzed using the algorithmic LLL condition of [3], avoiding the need for a problem-specific entropy compression argument.)

Our first result related to graph coloring is a generalization of Molloy's result, bounding the list-chromatic number as a function of the number of triangles in each neighborhood. Specifically, in Section 4 we establish the following theorem which is a key ingredient in the proof of Theorem 1.1. Note that, in order to comply with the standard notation used in results in the area, we express the bound on the number of triangles as Δ^2/f ; the triangle-free case then corresponds to $f = \Delta^2 + 1$. We stress that Molloy's proof breaks in the presence of even a single triangle per vertex.

Theorem 2.5. Let G be any graph with maximum degree Δ in which the neighbors of every vertex span at most Δ^2/f edges. For all $\epsilon > 0$, there exists Δ_ϵ such that if $\Delta \geq \Delta_\epsilon$ and $f \in [\Delta^{\frac{2+2\epsilon}{1+2\epsilon}} (\ln \Delta)^2, \Delta^2 + 1]$, then

$$\chi_\ell(G) \leq (1 + \epsilon)\Delta / \ln \sqrt{f}.$$

Furthermore, if G is a graph on n vertices then, for every $c > 0$, there exists an algorithm that constructs such a coloring in polynomial time with probability at least $1 - \frac{1}{n^c}$.

Theorem 2.5 is interesting for several reasons. First, random graphs suggest that it is sharp, i.e., that no efficient algorithm can color graphs satisfying the conditions of the theorem with $(1 - \epsilon)\Delta / \ln \sqrt{f}$ colors. More precisely, Proposition 2.1 below, proved in Appendix F, implies that any such algorithm would entail coloring random graphs using fewer than twice as many colors as their chromatic number. This would be a major (and unexpected) breakthrough in random graph theory, where beating this factor of two has been an elusive goal for over 30 years. Besides the lack of progress, further evidence for the optimality of this factor of two is that it corresponds precisely to a phase transition in the geometry of the set of colorings [1], known as the *shattering threshold*. Second, Theorem 2.5 establishes the existence of an algorithm that is robust enough to apply to worst-case graphs, while at the same time matching the performance of the best known (and highly tuned) algorithms for random graphs:

Proposition 2.1. *For every $\epsilon > 0$ and $d \in (d_\epsilon \ln n, (n \ln n)^{\frac{1}{3}})$, there exist $\Delta = \Delta(d, \epsilon)$ and $f = f(d, \epsilon)$ such that with probability tending to 1 as $n \rightarrow \infty$, a random graph $G = G(n, d/n)$ satisfies the conditions of Theorem 2.5 and $\chi(G) \geq (\frac{1}{2} - \epsilon)\Delta / \ln \sqrt{f}$.*

Third, armed with Theorem 2.5, we are able to prove the following result concerning the chromatic number of *arbitrary* graphs, as a function of the maximum degree and the maximum number of triangles in any neighborhood:

Theorem 2.6. *Let G be a graph with maximum degree Δ in which the neighbors of every vertex span at most Δ^2/f edges. For all $\epsilon > 0$, there exist $\Delta_\epsilon, f_\epsilon$ such that if $\Delta \geq \Delta_\epsilon$ and $f \in [f_\epsilon, \Delta^2 + 1]$, then*

$$\chi(G) \leq (2 + \epsilon)\Delta / \ln \sqrt{f}. \quad (7)$$

Furthermore, if G is a graph on n vertices then, for every $c > 0$, there exists an algorithm that constructs such a coloring in polynomial time with probability at least $1 - \frac{1}{n^c}$.

Theorem 2.6 improves a classical result of Alon, Krivelevich and Sudakov [8] which established (7) with an unspecified (large) constant in place of $2 + \epsilon$. The main idea in their analysis is to break down the input graph into triangle-free subgraphs, and color each one of them separately using distinct sets of colors by applying the result of Johansson [44]. Note that even if one used Molloy's recent result [51] in place of Johansson's in this scheme, the corresponding constant would still be in the thousands. Instead, we break down the graph into subgraphs with *few* triangles per neighborhood, and use Theorem 2.5 to color the pieces. The proof of Theorem 2.6 can be found in Appendix E.

As final remark, we note that Vu [63] proved the analogue of the main result of [8] (again with a large constant) for the list chromatic number. While we don't currently see how to sharpen Vu's result to an analogue of Theorem 2.6 for the list chromatic number using our techniques, we note that our Theorem 2.5 improves over [63] for all $f \geq \Delta^{\frac{2+2\epsilon}{1+2\epsilon}} (\ln \Delta)^2$.

3 Proof of main theorem

In Sections 3.2 and 3.3 we present the proof of our main result, Theorem 2.4. In Section 3.4 we show how to extend the theorem to allow flaw choice strategies other than following a fixed permutation over flaws.

Throughout this section we use standard facts about operator norms, summarized briefly in Appendix A.

3.1 Charges as norms of transition matrices

We will first show how charges can be seen as the norms of certain transition matrices. For more concrete examples of this connection, see Appendix B.

Recall that for any $S \subseteq [m]$, we denote by S^P and S^N the subsets of S that correspond to primary and non-primary flaws, respectively.

Definition 3.1. For every $i \in [m]$ and every set of flaw indices $S \subseteq [m]$, let A_i^S be the $|\Omega| \times |\Omega|$ matrix where $A_i^S[\sigma, \tau] = \rho_i(\sigma, \tau)$ if the set of flaws introduced by $\sigma \rightarrow \tau$ covers S , i.e., the set of primary flaws introduced by the transition $\sigma \rightarrow \tau$ equals S^P and the set of non-primary flaws introduced by $\sigma \rightarrow \tau$ contains S^N ; otherwise $A_i^S[\sigma, \tau] = 0$.

Let $\|\cdot\|_1$ denote the matrix norm induced by the L^1 -vector-norm, and recall that it is equal to the max column sum. Let also $M = \text{diag}(\mu(\sigma))$ denote the $|\Omega| \times |\Omega|$ diagonal matrix whose entries correspond to the probability measure μ . Our key observation is that the charges γ_i^S introduced in (5) can be expressed as

$$\gamma_i^S = \|MA_i^SM^{-1}\|_1. \quad (8)$$

The reader is encouraged to verify this equivalence, which is an immediate consequence of the definitions.

Remark 3.1. Although we are specializing here to the $\|\cdot\|_1$ norm and matrix $M = \text{diag}(\mu(\sigma))$, Theorem 2.4 holds for any choice of matrix norm and invertible matrix M . It is an interesting research direction whether using other norms can be useful in applications.

3.2 Tracking the set of current flaws

We say that a trajectory $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_{t+1})$ followed by the algorithm is a *bad t -trajectory* if every state σ_i , $i \in [t+1]$, is flawed. Thus, our goal is to bound the probability that the algorithm follows a bad t -trajectory.

Given a bad trajectory, we will track the flaws introduced into the state in each step, where a flaw is said to “introduce itself” whenever addressing it fails to remove it. Of the flaws introduced at each step, we disregard those that later get eradicated collaterally, i.e., by an action addressing some other flaw. The rest form the “witness sequence” of the trajectory, i.e., a sequence of sets of flaws.

Fix any permutation π on $[m]$. For any $S \subseteq [m]$, let $\pi(S) = \min_{j \in S} \pi(j)$, i.e., the lowest index in S according to π . Recalling that $U(\sigma)$ is the set of indices of flaws present in σ , in the following we assume that the index of the flaw addressed in state σ is $\pi(U(\sigma))$, which we sometimes abbreviate as $\pi(\sigma)$. Also, to lighten notation, we will denote $A \setminus \{\pi(B)\}$ by $A - \pi(B)$.

Definition 3.2. Let $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_{t+1})$ be any bad t -trajectory. Let $B_0 = U(\sigma_1)$. For $1 \leq i \leq t$, let

$$B_i = U(\sigma_{i+1}) \setminus [U(\sigma_i) - \pi(\sigma_i)],$$

i.e., B_i comprises the indices of the flaws introduced in the i -th step. For $0 \leq i \leq t$, let

$$C_i = \{k \in B_i \mid \exists j \in [i+1, t] : k \notin U(\sigma_{j+1}) \wedge \forall \ell \in [i+1, j] : k \neq \pi(\sigma_\ell)\},$$

i.e., C_i comprises the indices of the flaws introduced in the i -th step that get eradicated collaterally. The witness sequence of Σ is the sequence of sets

$$w(\Sigma) = (B_0 \setminus C_0, B_1 \setminus C_1, \dots, B_t \setminus C_t).$$

A crucial feature of witness sequences is that they allow us to recover the sequence of flaws addressed.

Definition 3.3. Given an arbitrary sequence S_0, \dots, S_t of subsets of $[m]$, let $S_1^* = S_0$, while for $1 \leq i \leq t$, let

$$S_{i+1}^* = \begin{cases} [S_i^* - \pi(S_i^*)] \cup S_i & \text{if } S_i^* \neq \emptyset; \\ \emptyset & \text{otherwise.} \end{cases}$$

If $S_i^* \neq \emptyset$ for all $1 \leq i \leq t$, then we say that $(S_i)_{i=0}^t$ is plausible and write $\pi(S_i^*) = (i)$.

Lemma 3.4. *If $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_{t+1})$ is any bad t -trajectory, then $w(\Sigma) = (S_0, \dots, S_t)$ is plausible, $\pi(\sigma_i) = \pi(S_i^*) = (i)$ for all $1 \leq i \leq t$, and for every flaw index $z \in [m]$, the number of times z occurs in the multiset $\bigcup_{i=0}^t S_i$ minus the number of times it occurs in the multiset $\bigcup_{i=1}^t (i)$ equals $\mathbf{1}_{z \in S_{t+1}^*}$.*

Proof. Recall that $S_i = B_i \setminus C_i$. For $1 \leq i \leq t+1$, let L_i comprise the elements of $U(\sigma_i)$ eradicated collaterally during the i -th step, and let H_i comprise the elements of $U(\sigma_i)$ eradicated collaterally during any step $j \geq i$. Observe that $H_{i+1} = (H_i \setminus L_i) \cup C_i$. We will prove, by induction, that for all $1 \leq i \leq t+1$,

$$S_i^* \subseteq U(\sigma_i); \quad (9)$$

$$U(\sigma_i) \setminus S_i^* = H_i. \quad (10)$$

Note that if (9) and (10) hold for a given i , then $\pi(\sigma_i) = \pi(S_i^*)$, since $\pi(\sigma_i) \notin H_i$ by the definition of H_i , and $\pi(A) = \pi(A \setminus B)$ whenever $\pi(A) \notin B$. Moreover, $S_i^* \neq \emptyset$, because otherwise $U(\sigma_i) = H_i$, an impossibility. To complete the proof it suffices to note that for any $z \in [m]$, the difference in question equals $\mathbf{1}_{z \in U(\sigma_{t+1})}$ and that $U(\sigma_{t+1}) = S_{t+1}^*$ since, by definition, $H_{t+1} = \emptyset$. The inductive proof is as follows.

For $i = 1$, (9) and (10) hold since $S_1^* = B_0 \setminus C_0$, while $U(\sigma_1) = B_0$. If (9) and (10) hold for some $i \geq 1$, then $S_{i+1}^* = [S_i^* - \pi(\sigma_i)] \cup S_i$ while, by definition, $U(\sigma_{i+1}) = [(U(\sigma_i) - \pi(\sigma_i)) \setminus L_i] \cup B_i$. Thus, the fact that $S_i^* \subseteq U(\sigma_i)$ trivially implies $S_{i+1}^* \subseteq U(\sigma_{i+1})$, while

$$U(\sigma_{i+1}) \setminus S_{i+1}^* = ((U(\sigma_i) \setminus S_i^*) \setminus L_i) \cup (B_i \setminus S_i) = (H_i \setminus L_i) \cup C_i = H_{i+1}. \quad \square$$

The first step in our proof of Theorem 2.4 is to give an upper bound on the probability that a given witness sequence occurs in terms of the charges γ_i^S . In particular, and in order to justify Remark 3.1, we will use an arbitrary norm $\|\cdot\|$ and invertible matrix M .

Recall that $\|\cdot\|_*$ denotes the *dual* of norm $\|\cdot\|$ and let $\theta^\top \in [0, 1]^{|\Omega|}$ denote the row vector corresponding to the probability distribution of the initial state σ_1 . Moreover, for a state σ , let e_σ denote the indicator vector of σ , i.e., $e_\sigma[\sigma] = 1$ and $e_\sigma[\tau] = 0$ for all $\tau \in \Omega \setminus \{\sigma\}$.

Lemma 3.5. *Fix any integer $t \geq 0$ and let Σ be the random variable $(\sigma_1, \dots, \sigma_{t+1})$. Fix any arbitrary invertible matrix M and operator norm $\|\cdot\|$, and let $\lambda_i^S = \|MA_i^S M^{-1}\|$. For any plausible sequence $\phi = (S_0, \dots, S_t)$,*

$$\Pr[w(\Sigma) = \phi] \leq \|\theta^\top M^{-1}\|_* \left(\sum_{\tau \in \Omega} \|M e_\tau\| \right) \prod_{i=1}^t \lambda_{(i)}^{S_i}. \quad (11)$$

Proof. Recall that for any $S \subseteq [m]$, we denote by S^P and S^N the subsets of S that correspond to primary and non-primary flaws, respectively. By Definition 3.2 and Lemma 3.4, a necessary condition for $w(\Sigma) = \phi$ to occur is that $(i) \in U(\sigma_i)$ and $S_i \subseteq B_i$, for every $1 \leq i \leq t$. Moreover, since primary flaws are never eradicated collaterally, i.e., $C_i^P = \emptyset$ always, it must also be that $S_i^P = B_i^P$ for $1 \leq i \leq t$. Fix any state $\tau \in \Omega$. The probability that $(1) \in U(\sigma_1) \wedge S_1^P = B_1^P(\Sigma) \wedge S_1^N \subseteq B_1^N(\Sigma) \wedge \sigma_2 = \tau$ equals the τ -column (coordinate) of the row-vector $\theta^\top A_{(1)}^{S_1^P}$. More generally, we see that for any $t \geq 1$,

$$\Pr \left[\bigwedge_{i=1}^t ((i) \in U(\sigma_i)) \bigwedge_{i=1}^t (S_i^P = B_i^P) \bigwedge_{i=1}^t (S_i^N \subseteq B_i^N) \bigwedge \sigma_{t+1} = \tau \right] = \theta^\top \prod_{i=1}^t A_{(i)}^{S_i^P} e_\tau. \quad (12)$$

Consider now any vector norm $\|\cdot\|$ and the corresponding operator norm. By (29),

$$\theta^\top \prod_{i=1}^t A_{(i)}^{S_i^P} e_\tau = \theta^\top M^{-1} \left(\prod_{i=1}^t M A_{(i)}^{S_i^P} M^{-1} \right) M e_\tau \leq \left\| \theta^\top M^{-1} \left(\prod_{i=1}^t M A_{(i)}^{S_i^P} M^{-1} \right) \right\|_* \|M e_\tau\|. \quad (13)$$

Summing (13) over all $\tau \in \Omega$ we conclude that

$$\Pr[w(\Sigma) = \phi] = \sum_{\tau \in \Omega} \Pr[w(\Sigma) = \phi \wedge \sigma_{t+1} = \tau] \leq \left\| \theta^\top M^{-1} \prod_{i=1}^t M A_{(i)}^{S_i} M^{-1} \right\|_* \sum_{\tau \in \Omega} \|M e_\tau\|. \quad (14)$$

Applying (31) and then (30) to (14), and recalling the definition of $\lambda_{(i)}^{S_i}$, we conclude that

$$\Pr[w(\Sigma) = \phi] \leq \|\theta^\top M^{-1}\|_* \left(\sum_{\tau \in \Omega} \|M e_\tau\| \right) \prod_{i=1}^t \|M A_{(i)}^{S_i} M^{-1}\| = \|\theta^\top M^{-1}\|_* \left(\sum_{\tau \in \Omega} \|M e_\tau\| \right) \prod_{i=1}^t \lambda_{(i)}^{S_i},$$

as claimed. \square

Now define

$$\mathcal{F}_t = \{w(\Sigma) : \Sigma \text{ is a bad } t\text{-trajectory of the algorithm}\}.$$

Since \mathcal{F}_t contains only plausible sequences, an immediate corollary of Lemma 3.5 is a bound on the probability that the algorithm fails in t steps.

Corollary 3.6. *The probability that the algorithm fails to reach a flawless state within t steps is at most*

$$\left(\max_{\sigma \in \Omega} \frac{\theta(\sigma)}{\mu(\sigma)} \right) \cdot \sum_{\phi \in \mathcal{F}_t} \prod_{i=1}^t \gamma_{(i)}^{S_i}. \quad (15)$$

Proof. We apply Lemma 3.5 with $M = \text{diag}(\mu(\sigma))$ and the $\|\cdot\|_1$ -norm. Since the dual norm of $\|\cdot\|_1$ is $\|\cdot\|_\infty$, we have $\|\theta^\top M^{-1}\|_\infty = \max_{\sigma \in \Omega} \frac{\theta(\sigma)}{\mu(\sigma)}$. Combining this with the fact that $\sum_{\tau \in \Omega} \|M e_\tau\|_1 = 1$ concludes the proof. \square

Thus, to complete the proof of Theorem 2.4 we are left with the task of bounding the sum in (15).

3.3 Bounding the sum

Given $\psi_1, \dots, \psi_m > 0$ and $S \subseteq [m]$, let $\Psi(S) = \prod_{j \in S} \psi_j$, with $\Psi(\emptyset) = 1$. For each $i \in [m]$, let

$$\zeta_i = \frac{1}{\psi_i} \sum_{S \subseteq [m]} \gamma_i^S \Psi(S).$$

Finally, for each $i \in [m]$ consider the distribution on $2^{[m]}$ that assigns to each $S \subseteq [m]$ the probability

$$p(i, S) := \frac{\gamma_i^S \Psi(S)}{\sum_{S \subseteq [m]} \gamma_i^S \Psi(S)} = \frac{\gamma_i^S \Psi(S)}{\zeta_i \psi_i}.$$

For any $S_0 \subseteq [m]$, let $\mathcal{F}_t(S_0)$ comprise the witness sequences in \mathcal{F}_t whose first set is S_0 . Consider the probability distribution on sequences of subsets of $[m]$ generated as follows: $R_1 = S_0$; for $i \geq 1$, if $R_i \neq \emptyset$, then $R_{i+1} = (R_i - \pi(R_i)) \cup S_i$, where $\Pr[S_i = S] = p(\pi(R_i), S)$ for any $S \subseteq [m]$. Under this distribution, by Lemma 3.4, each $\phi = (S_0, \dots, S_t) \in \mathcal{F}_t(S_0)$ receives probability $p_\phi = \prod_{i=1}^t p((i), S_i)$, while $\sum_{\phi \in \mathcal{F}_t(S_0)} p_\phi \leq 1$. At the same time, by the last claim in Lemma 3.4,

$$p_\phi = \prod_{i=1}^t p((i), S_i) = \left(\prod_{i=1}^t p((i), S_i) \frac{\psi_{(i)}}{\Psi(S_i)} \right) \frac{\Psi(S_{t+1}^*)}{\Psi(S_0)} = \frac{\Psi(S_{t+1}^*)}{\Psi(S_0)} \prod_{i=1}^t \frac{\gamma_{(i)}^{S_i}}{\zeta_{(i)}}. \quad (16)$$

Combining (16) with the fact $\sum_{\phi \in \mathcal{F}_t(S_0)} p_\phi \leq 1$, we get

$$\sum_{\phi \in \mathcal{F}_t(S_0)} \prod_{i=1}^t \frac{\gamma_{(i)}^{S_i}}{\zeta_{(i)}} \leq \max_{S \subseteq [m]} \frac{\Psi(S_0)}{\Psi(S)}. \quad (17)$$

Let $\zeta = \max_{i \in [m]} \zeta_i$. Then, summing equation (17) over all possible sets S_0 yields

$$\sum_{\phi \in \mathcal{F}_t} \prod_{i=1}^t \gamma_{(i)}^{S_i} = \sum_{S_0 \subseteq \text{Span}(\theta)} \sum_{\phi \in \mathcal{F}_t(S_0)} \prod_{i=1}^t \gamma_{(i)}^{S_i} \leq \zeta^t \sum_{S_0 \subseteq \text{Span}(\theta)} \sum_{\phi \in \mathcal{F}_t(S_0)} \prod_{i=1}^t \frac{\gamma_{(i)}^{S_i}}{\zeta_{(i)}} \leq \max_{S \subseteq [m]} \sum_{S_0 \subseteq \text{Span}(\theta)} \frac{\Psi(S_0)}{\Psi(S)}. \quad (18)$$

Proofs of Theorem 2.4 and Remark 2.4. Combining (18) with Corollary 3.6, we see that the binary logarithm of the probability that the algorithm does not encounter a flawless state within t steps is at most $t \log_2 \zeta + T_0$, where

$$T_0 = \log_2 \left(\max_{\sigma \in \Omega} \frac{\theta(\sigma)}{\mu(\sigma)} \right) + \log_2 \left(\sum_{S \subseteq \text{Span}(\theta)} \Psi(S) \right) + \log_2 \left(\max_{S \subseteq [m]} \frac{1}{\Psi(S)} \right).$$

Therefore, if $t = (T_0 + s)/\log_2(1/\zeta) \leq (T_0 + s)/\delta$, the probability that the algorithm does not reach a flawless state within t steps is at most 2^{-s} . This concludes the proofs of the first part of Remark 2.4 and Theorem 2.4, since $\max_{\sigma \in \Omega} \theta(\sigma) \leq 1$ and

$$\log_2 \left(\sum_{S \subseteq \text{Span}(\theta)} \Psi(S) \right) + \log_2 \left(\max_{S \subseteq [m]} \frac{1}{\Psi(S)} \right) \leq \log_2 \frac{\prod_{i=1}^m (1 + \psi_i)}{(\psi_{\min})^m} \leq m \log_2 \left(\frac{1 + \psi_{\max}}{\psi_{\min}} \right).$$

To see the second part of Remark 2.4, let $\mathcal{I}(\theta)$ denote the set comprising the sets of flaw-indices that may be present in a state selected according to θ . Recall now that when every flaw is primary, the only equivalence classes of \mathcal{F}_t that contribute to the sum in (18) are those for which $S_0 \in \mathcal{I}(\theta)$. Thus, for backtracking algorithms, the sum over $S \subseteq \text{Span}(\theta)$ in the definition of T_0 can be restricted to $S \in \mathcal{I}(\theta)$. Finally, if every flaw is always present in the initial state and $\psi_i \in (0, 1]$ for every $i \in [m]$, then $\mathcal{I}(\theta) = \{F\}$ and $\log \left(\frac{1}{\max_{S \subseteq [m]} \prod_{j \in S} \psi_j} \right) = -\log_2 \prod_{j \in [m]} \psi_j$. This implies that the second and third term in the expression for T_0 in Remark 2.4 cancel out, concluding its proof. \square

3.4 Other flaw choice strategies

The only place where we used the fact that the flaw choice is based on a fixed permutation was to assert, in Lemma 3.4, that the witness sequence of a trajectory determines the sequence of addressed flaws. Thus, our analysis is in fact valid for every flaw choice strategy that shares this property.

One example of such a strategy is “pick a random occurring flaw and address it”. To implement this, we can fix a priori an infinite sequence of uniformly random permutations π_1, π_2, \dots and at the i -th step address the lowest indexed flaw present according to π_i . It is straightforward to see that Lemma 3.4 still holds if we replace π with π_i therein and in Definition 3.3.

As a second example, consider the following recursive way to chose which flaw to address at each step (which makes the algorithm non-Markovian). The algorithm now maintains a stack. The flaws present in σ_1 , ordered according to some permutation π , comprise the initial stack contents. The algorithm starts by

addressing the flaw at the top of the stack, i.e., $\pi(\sigma_1)$, as before. Now, though, any flaws introduced in the i -th step, i.e., the elements of B_i , go on the top of the stack (ordered by π), while all eradicated flaws are removed from the stack. The algorithm terminates when the stack empties. It is not hard to see that, by taking S_0 to be the initial stack contents, popping the flaw at the top of the stack at each step, and adding S_i to the top of the stack (ordered by π), the sequence of popped flaws is the sequence of addressed flaws.

4 Graph coloring proofs

4.1 The algorithm

To prove Theorem 2.5 we will generalize the algorithm of Molloy [51] for coloring triangle-free graphs. The main issue we have to address is that in the presence of triangles, the natural generalization of Molloy's algorithm introduces monochromatic edges when the neighborhood of a vertex is recolored. As a result, the existing analysis fails completely even if each vertex participates in just one triangle. To get around this problem, we introduce backtracking steps into the algorithm, whose analysis is enabled by our new convergence condition, Theorem 2.4.

For each vertex $v \in V$, let N_v denote the neighbors of v and let $E_v = \{\{u_1, u_2\} : u_1, u_2 \in N_v\}$ denote the edges spanned by them. Recall that the color-list of v is denoted by \mathcal{L}_v . It will be convenient to treat Blank also as a color. Indeed, the initial distribution θ of our algorithm assigns all its probability mass to the state where every vertex is colored Blank. Whenever assigning a color to a vertex creates monochromatic edges, the algorithm will immediately uncolor enough vertices so that no monochromatic edge remains. Edges with two Blank endpoints are not considered monochromatic. To uncolor a vertex v , the algorithm picks a monochromatic edge e incident to v and assigns e to v instead of a color, thus also creating a record of the reason for the uncoloring. Thus,

$$\Omega \subseteq \prod_{v \in V} \{\mathcal{L}_v \cup \{\text{Blank}\} \cup E_v\}.$$

Let $L = (1 + \epsilon) \frac{\Delta}{\ln f} f^{-\frac{1}{2+2\epsilon}}$ and assume Δ is sufficiently large so that $L \geq 10$.

4.1.1 The flaws

We let $L_v(\sigma) \subseteq (\mathcal{L}_v \cup \{\text{Blank}\})$ be the set of colors we can assign to v in state σ without creating any monochromatic edge. We call these the *available colors for v in σ* and note that Blank is always available. For each $v \in V$, we define a flaw expressing the fact that there are “too few available colors for v ,” namely

$$B_v = \{\sigma \in \Omega : |L_v(\sigma)| < L\}.$$

For each color c other than Blank, let $T_{v,c}(\sigma)$ be the set of Blank neighbors of v for which c is available in σ , i.e., the vertices that may “compete” with v for color c . For each $v \in V$, we define a flaw expressing the fact that there is “too much competition for v 's available (real) colors,” namely

$$Z_v = \left\{ \sigma \in \Omega : \sum_{c \in L_v(\sigma) \setminus \text{Blank}} |T_{v,c}(\sigma)| > \frac{L}{10} |L_v(\sigma)| \right\}.$$

Finally, for each $v \in V$ and $e \in E$ we define a flaw for the fact that v is uncolored (because of e), namely

$$f_v^e = \{\sigma \in \Omega : \sigma(v) = e\}.$$

Let $F_v = B_v \cup Z_v \cup \bigcup_{e \in E} f_v^e$ and let $\Omega^+ = \Omega - \bigcup_{v \in V} F_v$; thus Ω^+ denotes the partial colorings that do not suffer from any of the above flaws.

Lemma 4.1 ([51]). *Given $\sigma \in \Omega^+$, a complete list-coloring of G can be found efficiently.*

The proof of Lemma 4.1 is a fairly standard application of the (algorithmic) LLL, showing that σ can be extended to a complete list-coloring by coloring all Blank vertices with actual colors. Thus, the heart of the matter is reaching a state in Ω^+ (i.e., a partial coloring avoiding all the above flaws).

4.1.2 The flaw choice

The algorithm can use any π -strategy in which every B -flaw has priority over every f -flaw.

4.1.3 The actions

To address f_v^e at σ , i.e., to color v , the algorithm simply chooses a color from $L_v(\sigma)$ uniformly at random and assigns it to v . The fact that B -flaws have higher priority than f -flaws implies that there are always at least L such choices.

Addressing B - and Z -flaws is significantly more sophisticated. For each vertex v , for each vertex $u \in N_v$, let $R_u^v(\sigma) \supseteq L_u(\sigma)$ comprise those colors having the property that assigning them to u in state σ creates no monochromatic edge except, perhaps, in E_v . To address either B_v or Z_v in σ , the algorithm selects an action according to the following procedure:

-
- 1: **procedure** RECOLOR(v, σ)
 - 2: Assign to each colored vertex u in N_v a uniformly random color from $R_u^v(\sigma)$
 - 3: **while** monochromatic edges exist **do**
 - 4: Let u be the lowest indexed vertex participating in a monochromatic edge
 - 5: Let e be the lowest indexed monochromatic edge with u as an endpoint
 - 6: Uncolor u by assigning e to u
-

Lemma 4.2. *Let $S'(v, \sigma)$ be the set of colorings that can be reached at the end of Step 2 of RECOLOR(v, σ) and let $S''(v, \sigma)$ be the set of possible final colorings. Then $|S'(v, \sigma)| = |S''(v, \sigma)|$.*

Proof. Since Steps 4–6 are deterministic, $|S''(v, \sigma)| \leq |S'(v, \sigma)|$. To prove that $|S''(v, \sigma)| \geq |S'(v, \sigma)|$, we will prove that if $u \in N_v$ has distinct colors in $\sigma'_1, \sigma'_2 \in S'$, then there exists $z \in V$ such that $\sigma''_1(z) \neq \sigma''_2(z)$. Imagine that in Step 6 we also oriented e to point away from u . Then, in the resulting partial orientation, every vertex would have outdegree at most 1 and there would be no directed cycles. Consider the (potentially empty) oriented paths starting at u in σ''_1 and σ''_2 , and let z be their last common vertex. If z is uncolored, then $\sigma''_1(z) = e_1$ and $\sigma''_2(z) = e_2$, where $e_1 \neq e_2$; if z is colored, then $\sigma''_i(z) = \sigma'_i(u)$. \square

4.2 Proving termination

Let D_v be the set of vertices at distance 1, 2 or 3 from v and let

$$S_v = \{B_u\}_{u \in D_v} \cup \{Z_u\}_{u \in D_v} \cup \{f_u^{\{u,w\}}\}_{u,w \in N_v}.$$

To lighten notation, in the following we write $\gamma^S(f)$ instead of γ_f^S . Let $q = (1 + \epsilon) \frac{\Delta}{\ln \sqrt{f}} \geq 1$.

Lemma 4.3. *For every vertex $v \in V$ and edge $e \in E$:*

- (a) *if $S \not\subseteq S_v$, then $\gamma^S(B_v) = \gamma^S(Z_v) = \gamma^S(f_v^e) = 0$;*
- (b) *if $S \supseteq \{f_{u_1}^{\{u_1, u_2\}}, f_{u_2}^{\{u_1, u_2\}}\}$, then $\gamma^S(B_v) = \gamma^S(Z_v) = \gamma^S(f_v^e) = 0$;*

- (c) $\max_{S \subseteq F} \gamma^S(f_v^e) \leq \frac{1}{L} =: \gamma(f_v^e)$;
- (d) $\max_{S \subseteq F} \gamma^S(B_v) \leq 2e^{-\frac{L}{6}} =: \gamma(B_v)$;
- (e) $\max_{S \subseteq F} \gamma^S(Z_v) \leq 3qe^{-\frac{L}{60}} =: \gamma(Z_v)$,

where the charges are computed with respect to the uniform measure over Ω .

We note that, while we give uniform bounds on the charges corresponding to each flaw, the analysis of our algorithm cannot be captured by the result of [3]. This is because we will crucially exploit the existence of primary flaws.

Before giving the proof of Lemma 4.3, we first use it to derive Theorem 2.5.

Proof of Theorem 2.5. For every flaw $f \in F$, we will take $\psi_f = \gamma(f)\psi$, where $\psi > 0$ will be chosen later.

For any vertex $v \in V$, flaw $f \in \{B_v, Z_v, f_v^e\}$, and set of flaws $S \subseteq F$, Lemma 4.3 implies that $\gamma^S(f) = 0$ unless all B - and Z -flaws in S correspond to vertices in D_v , per part (a), and every edge $e \in E_v$ contributes at most one flaw to S , per part (b). Therefore, for $f \in \{B_v, Z_v, f_v^e\}$,

$$\frac{1}{\psi_f} \sum_{S \subseteq F} \gamma^S(f) \prod_{g \in S} \psi_g \leq \frac{1}{\psi} \prod_{u \in D_v} (1 + \gamma(B_u)\psi)(1 + \gamma(Z_u)\psi) \prod_{e=\{u_1, u_2\} \in E_v} (1 + \gamma(f_{u_1}^e)\psi + \gamma(f_{u_2}^e)\psi). \quad (19)$$

To bound the right hand side of (19) we use parts (c)–(e) of Lemma 4.3 along with the facts $|D_v| \leq \Delta^3 + 1$ and $|E_v| \leq \Delta^2/f$ to derive (20) below. To derive (21), we use the facts that $2e^{-\frac{L}{6}} \leq 3qe^{-\frac{L}{60}}$, since $q \geq 1$, and that $1 + x \leq e^x$ for all x . Thus, for $f \in \{B_v, Z_v, f_v^e\}$, we conclude

$$\frac{1}{\psi_f} \sum_{S \subseteq F} \gamma^S(f) \prod_{g \in S} \psi_g \leq \frac{1}{\psi} \left(1 + 2e^{-\frac{L}{6}}\psi\right)^{\Delta^3+1} \left(1 + 3qe^{-\frac{L}{60}}\psi\right)^{\Delta^3+1} \left(1 + \frac{2\psi}{L}\right)^{\frac{\Delta^2}{f}} \quad (20)$$

$$\leq \frac{1}{\psi} \exp\left(\frac{2\psi\Delta^2}{fL} + 6qe^{-\frac{L}{60}}\psi(\Delta^3 + 1)\right) =: \frac{1}{\psi} \exp(Q). \quad (21)$$

Setting $\psi = (1 + \epsilon)$, we see that the right hand side of (21) is strictly less than 1 for all $\Delta \geq \Delta_\epsilon$, since $Q \xrightarrow{\Delta \rightarrow \infty} 0$ for all $f \in [\Delta^{\frac{2+2\epsilon}{1+2\epsilon}}(\ln \Delta)^2, \Delta^2 + 1]$. To see this last claim, recall that $L = (1 + \epsilon)\frac{\Delta}{\ln f}f^{-\frac{1}{2+2\epsilon}}$ and $q = (1 + \epsilon)\frac{\Delta}{\ln \sqrt{f}}$, and note that $\ln f < 3 \ln \Delta$ and $f^{\frac{1+2\epsilon}{2+2\epsilon}} \geq \Delta(\ln \Delta)^{\frac{2+4\epsilon}{2+2\epsilon}}$. Thus,

$$\frac{2\psi\Delta^2}{fL} = \frac{2\Delta^2}{f\frac{\Delta}{\ln f}f^{-\frac{1}{2+2\epsilon}}} = \frac{2\Delta \ln f}{f^{\frac{1+2\epsilon}{2+2\epsilon}}} \leq \frac{2 \ln f}{(\ln \Delta)^{\frac{2+4\epsilon}{2+2\epsilon}}} \leq \frac{6 \ln \Delta}{(\ln \Delta)^{\frac{2+4\epsilon}{2+2\epsilon}}} = \frac{6}{(\ln \Delta)^{\frac{\epsilon}{1+\epsilon}}} \xrightarrow{\Delta \rightarrow \infty} 0, \quad (22)$$

while the facts $L = \Omega(\Delta^{\frac{\epsilon}{1+2\epsilon}})$ and $q \leq (1 + \epsilon)\Delta$ imply that $6qe^{-\frac{L}{60}}\psi(\Delta^3 + 1) \xrightarrow{\Delta \rightarrow \infty} 0$. □

4.2.1 Proof of Lemma 4.3

Proof of part (a). Addressing B_v or Z_v by executing $\text{RECOLOR}(v, \cdot)$ only changes the color of vertices in N_v , with any resulting uncolorings being due to edges in E_v . Thus, only flaws in S_v may be introduced. Addressing f_v^e by coloring v trivially can only introduce flaws B_u, Z_u , where $u \in N_v$. □

Proof of part (b). Since addressing an f -flaw never introduces another f -flaw, we only need to discuss procedure RECOLOR. Therein, vertices are uncolored serially in time, so that any time a vertex w is uncolored there exists, at the time of w 's uncoloring, a monochromatic edge $e = \{w, u\}$. Therefore, an edge $e = \{u_1, u_2\}$ can never be the reason for the uncoloring of both its endpoints, i.e., $f_{u_1}^e \cap f_{u_2}^e = \emptyset$. \square

Proof of part (c). If addressing f_v^e results in τ , then the previous state σ must be the mutation of τ that results from assigning e to v . Since $\pi(\sigma) = f_v^e$ implies $\sigma \notin B_v$, it follows that $|L_v(\sigma)| \geq L$. Since colors are chosen uniformly from $L_v(\sigma)$, it follows that $\gamma(f_v^e) \leq 1/L$. \square

Proof of parts (d) and (e). Observe that every flaw corresponding to an uncolored vertex is primary, since procedure RECOLOR never colors an uncolored vertex and addressing f_v^e only colors v . Thus, when computing $\gamma^S(f)$, for $f \in \{B_v, Z_v\}$ and $S \subseteq F$, we can restrict to pairs (σ, τ) such that the set of uncolored vertices in τ is exactly the union of the set of uncolored vertices in σ and the set $\{u \in N_v : f_u^e \in S\}$. Fixing $f \in \{B_v, Z_v\}$, $S \subseteq F$, and τ , let us denote by $\text{In}_f^S(\tau)$ the candidate set of originating states, and by $\mathcal{U}_f^S(\tau)$ their common set of uncolored vertices. Then, for any $f \in \{B_v, Z_v\}$ and any $S \subseteq F$,

$$\gamma^S(f) = \max_{\tau \in \Omega} \sum_{\sigma \in \text{In}_f^S(\tau)} \rho_f(\sigma, \tau). \quad (23)$$

To bound $\rho_f(\sigma, \tau)$ in (23), we recall that RECOLOR assigns to each colored vertex $u \in N_v$ a random color from $R_u^v(\sigma)$ and invoke Lemma 4.2 to derive the first equality in (24). For the second equality we observe that for every $u \in N_v$, the set R_u^v is determined by the colors of the vertices in $V \setminus N_v$. Since RECOLOR only changes the color of vertices in N_v , it follows that $R_u^v(\sigma) = R_u^v(\tau)$, yielding

$$\rho_f(\sigma, \tau) = \frac{1}{\prod_{u \in N_v \setminus \mathcal{U}_f^S(\tau)} |R_u^v(\sigma)|} = \frac{1}{\prod_{u \in N_v \setminus \mathcal{U}_f^S(\tau)} |R_u^v(\tau)|} := \frac{1}{\Lambda_f^S(\tau)}. \quad (24)$$

Next we bound $|\text{In}_f^S(\tau)|$ as follows. First we observe that if $\sigma \in \text{In}_f^S(\tau)$, then $\sigma(u) \neq \tau(u)$ implies $u \in N_v \setminus \mathcal{U}_f^S(\tau)$ and, therefore, $\sigma(u) \in R_u^v(\tau)$ since $\sigma(u) \in L_u(\sigma) \subseteq R_u^v(\sigma) = R_u^v(\tau)$. Thus, the set of τ -mutations that result from recoloring each vertex in $N_v \setminus \mathcal{U}_f^S(\tau)$ with a color from $R_u^v(\tau)$ so that the resulting state belongs to f is a superset of $\text{In}_f^S(\tau)$. Denoting this last set by $\text{Viol}(f, \tau)$, we conclude that

$$\gamma^S(f) = \max_{\tau \in \Omega} \frac{|\text{In}_f^S(\tau)|}{\Lambda_f^S(\tau)} \leq \max_{\tau \in \Omega} \frac{|\text{Viol}(f, \tau)|}{\Lambda_f^S(\tau)} = \max_{\tau \in \Omega} \Pr[\text{RECOLOR}(v, \tau) \in f], \quad (25)$$

where for the last equality we use the definition of $\Lambda_f^S(\tau)$.

Remark 4.1. We note that expressing the sum of the transition probabilities into a state in terms of a random experiment, as we do in (25), was the key technical idea of [51] in order to apply the entropy compression method. It is also the one that breaks down if we allow our algorithm to go through improper colorings.

To conclude the proof of Lemma 4.3 we prove the following bounds in Appendix D via fairly routine calculations.

Lemma 4.4. For each vertex v and $\sigma \in \Omega$:

- (a) $\Pr[\text{RECOLOR}(v, \sigma) \in B_v] \leq 2e^{-\frac{L}{6}}$.
- (b) $\Pr[\text{RECOLOR}(v, \sigma) \in Z_v] \leq 3qe^{-\frac{L}{60}}$.

\square

5 Applications to backtracking algorithms

An important class of algorithms naturally devoid of “collateral fixes” are *backtracking* algorithms. In particular, consider a Constraint Satisfaction Problem (CSP) over a set of variables $V = \{v_1, v_2, \dots, v_n\}$, each variable v_i taking values in a domain \mathcal{D}_i , with a set of constraints $\mathcal{C} = \{c_1, c_2, \dots, c_m\}$ over these variables. The backtracking algorithms we consider operate as follows. (Note that in Step 1, we can always take θ to be the distribution under which all variables are unassigned; this does not affect the convergence condition (6) but may have a mild effect on the running time.)

Generic Backtracking Algorithm

- 1: Sample a partial non-violating assignment σ_0 according to a distribution θ and set $i = 0$
 - 2: **while** unassigned variables exist **do**
 - 3: Let v be the lowest indexed unassigned variable in σ_i
 - 4: Choose a new value for v according to a state-dependent probability distribution
 - 5: **if** one or more constraints are violated **then**
 - 6: Remove the values from enough variables so that no constraint is violated
 - 7: Let σ_{i+1} be the resulting assignment
 - 8: $i \leftarrow i + 1$
-

Let Ω be the set of partial assignments to V that do not violate any constraint in \mathcal{C} . For each variable $v_i \in V$, let flow $f_i \subseteq \Omega$ comprise the partial assignments in which v_i is unassigned. Clearly, each flow f_i can only be removed by addressing it, as addressing any other flow can only unassign v_i . Thus, every flow is primary and a flawless state is a complete satisfying assignment.

In this section we present applications of our main theorem to analyze backtracking search algorithms. First, we prove a useful corollary of Theorem 2.4 that holds in the so-called *variable setting*. Second, we analyze a backtracking algorithm of Esperet and Parreau [27] for acyclic edge coloring that lies outside the variable setting. In particular, we recover their 4Δ bound on the acyclic chromatic index and, further, we show how it can make constructive an existential result of Bernshteyn [12]. We emphasize that all these analyses follow very easily from our framework.

5.1 The variable setting

In this section we show how we can use Theorem 2.4 to employ backtracking algorithms in order to capture applications in the variable setting, i.e., the setting considered by Moser and Tardos. In particular, we consider a product measure over variables V and define a bad event for each constraint $c \in \mathcal{C}$ being violated. We show the following corollary of Theorem 2.4.

Theorem 5.1. *Let P be any product measure over a set of variables V and let A_c be the event that constraint c is violated. If there exist positive real numbers $\{\psi_v\}_{v \in V}$ such that for every variable $v \in V$,*

$$\frac{1}{\psi_v} \left(1 + \sum_{c \ni v} P(A_c) \prod_{u \in c} \psi_u \right) < 1, \quad (26)$$

then there exists a backtracking algorithm that finds a satisfying assignment after an expected number of $O\left(\log(P_{\min}^{-1}) + |V| \log_2 \left(\frac{1 + \psi_{\max}}{\psi_{\min}}\right)\right)$ steps.

Before proving Theorem 5.1, we first use it to capture a well-known application of the Lovász Local Lemma to sparse k -SAT formulas when P is the uniform measure. For a k -SAT formula Φ , we denote its maximum degree by $\Delta \equiv \Delta(\Phi)$, i.e., each variable of Φ is contained in at most Δ clauses.

Theorem 5.2. Every k -SAT formula Φ with maximum degree $\Delta < \frac{2^k}{ek}$ is satisfiable. Moreover, there exists a backtracking algorithm that finds a satisfying assignment of Φ efficiently.

Proof. Setting $\psi_v = \psi = 2\alpha > 0$ we see that it suffices to find a value $\alpha > 0$ such that

$$\frac{1}{\psi} + \frac{1}{2^k} \Delta \psi^{k-1} = \frac{1}{2\alpha} + \frac{1}{2} \Delta \alpha^{k-1} < 1,$$

which is feasible whenever

$$\Delta < \max_{\alpha > 0} \frac{2\alpha - 1}{\alpha^k} = \frac{2^k}{k} \cdot \left(1 - \frac{1}{k}\right)^{k-1} \leq \frac{2^k}{ek}.$$

□

Remark 5.1. In [28] it is shown that using a non-uniform product measure P one can improve the bound of Theorem 5.2 to $\Delta < \frac{2^{k+1}}{e(k+1)}$ and that this is asymptotically tight. We note that we can achieve the same bound using Theorem 5.1 with the same P , but since this is a rather involved LLL application we will not explicitly present it here.

5.1.1 Proof of Theorem 5.1

We consider the following very simple backtracking algorithm. Start with each variable unassigned. Then, in each state σ , choose the lowest indexed unassigned variable v and sample a value for it according to the product measure P . If one or more constraints become violated, remove the value from every variable in the lowest indexed violated constraint.

Let Ω be the set of partial non-violating assignments. Let $\mu : \Omega \rightarrow \mathbb{R}$ be the probability measure that assigns to each state $\sigma \in \Omega$ the value $\mu(\sigma) \propto \prod_{\substack{v \in V \\ v \notin U(\sigma)}} P(\sigma(v))$, where for brevity we abuse notation by letting $P(\sigma(v))$ denote the event that variable v is assigned value $\sigma(v)$.

Theorem 5.1 will follow immediately from the following lemma. (For brevity, we will index flaws with variables instead of integers.)

Lemma 5.3. For each vertex v and set of variables $S \neq \emptyset$,

$$\gamma_v^S = \begin{cases} 1 & \text{if } S = \emptyset; \\ P(A_c) & \text{if } S = c, \text{ where } c \text{ is a constraint containing } v; \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Notice that the actions related to flaw f_v can only remove the value from sets of variables that correspond to constraints that contain v . Thus, $\gamma_v^S = 0$ for every set $S \neq \emptyset$ that does not correspond to a constraint containing v . Recalling the definition of charges and $U(\sigma)$, we have

$$\gamma_v^S = \max_{\tau \in \Omega} \sum_{\substack{\sigma \in f_v \\ S = U(\tau) \setminus (U(\sigma) \setminus \{v\})}} \frac{\mu(\sigma)}{\mu(\tau)} \rho_v(\sigma, \tau). \quad (27)$$

To see the claim for the case of the empty set, notice that, given a state τ , there exists at most one state σ such that $\rho_v(\sigma, \tau) > 0$ and that $U(\tau) \setminus (U(\sigma) \setminus \{v\}) = \emptyset$. This is because we can uniquely reconstruct σ from τ by removing the value from v at τ . Then we have

$$\frac{\mu(\sigma)}{\mu(\tau)} \rho_v(\sigma, \tau) = \frac{\prod_{u \in V \setminus U(\sigma)} P(\sigma(u))}{\prod_{u \in V \setminus U(\tau)} P(\tau(u))} P(\tau(v)) = \frac{1}{P(\tau(v))} P(\tau(v)) = 1.$$

To see the claim for the case where $S = c$, consider the set $\text{viol}(c)$ consisting of the set of assignments of the variables of c that violate c . Notice now that for every state $\tau \in \Omega$ there is an injection from the set of states σ such that $\rho_v(\sigma, \tau) > 0$ and $S = U(\tau) \setminus (U(\sigma) \setminus \{v\})$ to $\text{viol}(c)$. This is because c should be violated at each such state σ , and hence each state σ should be of the form $\sigma = \tau_\alpha$ for $\alpha \in \text{viol}(c)$, where τ_α is the state induced by τ when assigning α to the variables of c . Observe further that, for every state of the form $\tau_\alpha, \alpha \in \text{viol}(c)$, we have that

$$\frac{\mu(\tau_\alpha)}{\mu(\tau)} \rho_v(\tau_\alpha, \tau) = \left(\prod_{u \in c \setminus \{v\}} P(X_u = \tau_\alpha(u)) \right) P(X_v = \tau(v)) = P(A_c^\alpha), \quad (28)$$

where $P(A_c^\alpha)$ is the probability of the event that the variables of c receive assignment α . Combining (28) with (27) and the fact that $P(A_c) = \sum_{\alpha \in \text{viol}(c)} P(A_c^\alpha)$ concludes the proof of Lemma 5.3. \square

Finally, plugging Lemma 5.3 into Theorem 2.4 concludes the proof of Theorem 5.1.

5.2 Acyclic edge coloring

An edge-coloring of a graph is *proper* if all edges incident to each vertex have distinct colors. A proper edge coloring is *acyclic* if it has no bichromatic cycles, i.e., no cycle receives exactly two (alternating) colors. The smallest number of colors for which a graph G has an acyclic edge-coloring is denoted by $\chi'_a(G)$.

Acyclic Edge Coloring was originally motivated by the work of Coleman et al. [18, 19] on the efficient computation of Hessians and, since then, there has been a series of papers [7, 27, 33, 47, 52, 56] that upper bound $\chi'_a(G)$ for graphs with bounded degree. The current best result was given recently by Giotis *et al.* [29], who showed that $\chi'_a(G) \leq 3.74\Delta$ in graphs with maximum degree Δ .

5.2.1 A simple backtracking algorithm

We show how one can apply Theorem 2.4 to recover the main application of the framework of [27] with a much simpler proof. We chose the result of [27] because, although it is not the sharpest known bound for the acyclic chromatic index, it is a canonical example of how the entropy compression method is applied to analyze backtracking algorithms and has inspired many other results in the area. (Indeed, the authors in [27] already give several applications of their techniques to other problems besides acyclic edge coloring.)

Let G be a graph with m edges $E = \{e_1, \dots, e_m\}$ and suppose we have q available colors.

Definition 5.4. *Given a graph $G = (V, E)$ and a (possibly partial) edge-coloring of G , say that color c is 4-forbidden for $e \in E$ if assigning c to e would result in either a violation of proper edge-coloration, or a bichromatic 4-cycle containing e . Say that c is 4-available if it is not 4-forbidden.*

Lemma 5.5 ([27]). *In any proper edge-coloring of G , at most $2(\Delta - 1)$ colors are 4-forbidden for any $e \in E$.*

Proof. The 4-forbidden colors for $e = \{u, v\}$ can be enumerated as: (i) the colors on edges adjacent to u ; and (ii) for each edge e_v adjacent to v , either the color of e_v (if no edge with that color is adjacent to u), or the color of some edge e' which together with e, e_v and an edge adjacent to u form a cycle of length 4. \square

Consider the following backtracking algorithm for Acyclic Edge Coloring with $q = 2(\Delta - 1) + Q$ colors. At each step, choose the lowest indexed uncolored edge e and attempt to color it choosing uniformly at random among the 4-available colors for e . If one or more bichromatic cycles are created, then choose the lowest indexed one of them, say $C = \{e_{i_1}, e_{i_2}, \dots, e_{i_{2\ell}} = e\}$, and remove the colors from all its edges except e_{i_1} and e_{i_2} .

The main result of [27] states that every graph G admits an acyclic edge coloring with $q > 4(\Delta - 1)$ colors. Moreover, such a coloring can be found in $O(|E||V|\Delta^2 \ln \Delta)$ time with high probability.

We prove the following theorem, which achieves the same bound on q and improves the running time bound when the graph is dense. More significantly, our analysis is much simpler.

Theorem 5.6. *Every graph G admits an acyclic edge coloring with $q > 4(\Delta - 1)$ colors. Moreover, such a coloring can be found in $O(|E||V|\Delta)$ time with high probability.*

Proof. Let Ω be the set of partial acyclic edge colorings of G . For each edge e , let f_e be the subset (flaw) of Ω that contains the partial acyclic edge colorings of G in which e is uncolored. We will apply Theorem 2.4 using the $\|\cdot\|_1$ norm and $M = \text{diag}(\frac{1}{|\Omega|})$.

We first compute the charges γ_e^S for each edge e and set of edges S . Notice that for γ_e^S to be non-zero, it should either be that $S = \emptyset$, or that S contains e and there exists a cycle $C = \{e_{i_1}, e_{i_2}\} \cup S$ so that, when a recoloring of e makes C bichromatic, the backtracking step uncolors precisely the edges in S . With that in mind, for each edge e and each set S that contains e , let $\mathcal{C}_e(S)$ denote the set of cycles with the latter property.

Lemma 5.7. *For each edge e , let*

$$\gamma_e^S = \begin{cases} \frac{1}{Q} & \text{if } S = \emptyset \\ \frac{|\mathcal{C}_e(S)|}{Q} & \text{if } e \in S \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Notice that

$$\gamma_e^S = \max_{\tau \in \Omega} \sum_{\substack{\sigma \in f_e \\ S = U(\tau) \setminus (U(\sigma) \setminus \{e\})}} \rho_e(\sigma, \tau) \leq \max_{\tau \in \Omega} \sum_{\substack{\sigma \in f_e \\ S = U(\tau) \setminus (U(\sigma) \setminus \{e\})}} \frac{1}{Q},$$

since, according to Lemma 5.5, $\rho_e(\sigma, \tau) \leq \frac{1}{Q}$ for each pair $(\sigma, \tau) \in f_e \times \Omega$. The proof follows by observing that, for each state τ :

- If $S = \emptyset$, then there exists at most one state σ such that $\rho_e(\sigma, \tau) > 0$ and $U(\tau) \setminus (U(\sigma) \setminus \{e\}) = \emptyset$ (we can reconstruct σ from τ by uncoloring e).
- If $S \ni e$ and $|S| = 2\ell - 2$, then there exist at most $|\mathcal{C}_e(S)|$ states such that $\rho_e(\sigma, \tau) > 0$ and $S = U(\tau) \setminus (U(\sigma) \setminus \{e\})$. Given a cycle $C = S \cup \{e_{i_1}, e_{i_2}\}$ we reconstruct σ from τ by finding the colors of edges in $S \setminus \{e\}$ from $\tau(e_{i_1}), \tau(e_{i_2})$, exploiting the facts that the backtracking step corresponds to an uncoloring of a bichromatic cycle, edge e is uncolored, and every other edge has the same color as in τ .
- For all other S there exists no state σ such that $\rho_e(\sigma, \tau) > 0$ and $S = U(\tau) \setminus (U(\sigma) \setminus \{e\})$.

□

Observe that there are at most $(\Delta - 1)^{2\ell - 2}$ cycles of length 2ℓ containing a specific edge e . In other words, there exist at most $(\Delta - 1)^{2\ell - 3}$ sets of edges S of size $2\ell - 2$ that contain e and such that $\gamma_e^S > 0$ and, in addition, note that we always have $|\mathcal{C}_e(S)| \leq \Delta - 1$.

Thus, if $Q = c(\Delta - 1)$ for some constant c , setting $\psi_e = \alpha\gamma_\emptyset^e = \frac{\alpha}{Q}$, where $\alpha \in (1, c)$ is a constant, Lemma 5.7 implies

$$\begin{aligned} \frac{1}{\psi_e} \left(\sum_{S \subseteq E} \gamma_e^S \prod_{e \in S} \psi_j \right) &\leq \min_{\alpha \in (1, c)} \left(\frac{1}{\alpha} + \sum_{i=3}^{\infty} \left(\frac{\Delta - 1}{Q} \right)^{2i-2} \alpha^{2i-3} \right) \\ &\leq \min_{\alpha \in (1, c)} \left(\frac{1}{\alpha} + \frac{1}{c} \sum_{i=3}^{\infty} \left(\frac{\alpha}{c} \right)^{2i-3} \right) \\ &= \min_{\alpha \in (1, c)} \left(\frac{1}{\alpha} + \frac{\alpha^3}{c^2(c^2 - \alpha^2)} \right) = \frac{2}{c} \end{aligned}$$

for $\alpha^* = c \left(\frac{\sqrt{5}-1}{2} \right)$. Thus, if $c > 2$ the probability that the algorithm fails to find an acyclic edge coloring within $\frac{T_0+s}{\delta}$ steps is 2^{-s} , where $\delta = 1 - \frac{2}{c}$, and, according to Remark 2.4,

$$T_0 = \log_2 |\Omega| = O(|E|).$$

The proof is concluded by observing that each step can be performed in time $O(|V|\Delta)$ (the time it takes to find a 2-colored cycle containing a given edge, if such a cycle exists, in a graph with a proper edge-coloring). \square

5.2.2 An application of the local cut lemma

Bernshteyn [14] introduced a non-constructive generalized LLL condition, called the ‘‘Local Cut Lemma’’, with the aim of drawing connections between the LLL and the entropy compression method. He later applied it in [12] to the problem of Acyclic Edge Coloring, giving improved bounds assuming further constraints on the graph besides sparsity. In particular, he proved the following:

Theorem 5.8 ([12]). *Let G be a graph with maximum degree Δ and let H be a fixed bipartite graph. If G does not contain H as a subgraph, then there exists an acyclic edge coloring of G using at most $3(\Delta + o(1))$ colors.*

We now show how to use our framework to give a constructive proof of Theorem 5.8. This will follow immediately from the following structural lemma in [12].

Lemma 5.9 ([12]). *There exist positive constants γ, δ such that the following holds. Let G be a graph with maximum degree Δ that does not contain H as a subgraph. Then, for any edge $e \in E(G)$ and for any integer $k \geq 4$, the number of cycles of length k in G that contain e is at most $\gamma\Delta^{k-2-\delta}$.*

Constructive Proof of Theorem 5.8. Notice that in this case, making almost identical calculations to those in the proof of Theorem 5.6, invoking Lemma 5.9 to upper bound the number of cycles that contain e and setting $\alpha = \frac{c}{\beta}$, we obtain

$$\frac{1}{\psi_e} \left(\sum_{S \subseteq E} \gamma_e^S \prod_{h \in S} \psi_h \right) \leq \min_{\alpha \in (1, c)} \left(\frac{1}{\alpha} + \frac{(\alpha)^3 \gamma \Delta^{-\delta}}{c^2(c^2 - \alpha^2)} \right) = \frac{1}{c} \min_{\beta > 1} \left(\beta + \frac{\beta \gamma \Delta^{-\delta}}{\beta(\beta^2 - 1)} \right).$$

Thus, as Δ grows, the value of c required for the algorithm to terminate approaches 1, concluding the proof. \square

6 Acknowledgements

We are grateful to Paris Syminelakis for various insightful remarks and comments. We also thank David Harris for helpful discussions.

References

- [1] Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 793–802. IEEE Computer Society, 2008.
- [2] Dimitris Achlioptas and Fotis Iliopoulos. Random walks that find perfect objects and the Lovász local lemma. *J. ACM*, 63(3):22:1–22:29, 2016.
- [3] Dimitris Achlioptas, Fotis Iliopoulos, and Vladimir Kolmogorov. A local lemma for focused stochastic algorithms. *SIAM Journal on Computing*, 48:1583–1602, 2019.
- [4] Dimitris Achlioptas, Fotis Iliopoulos, and Nikos Vlassis. Stochastic control via entropy compression. In *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 83:1–83:13. Springer, 2017.
- [5] Dimitris Achlioptas and Assaf Naor. The two possible values of the chromatic number of a random graph. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 587–593. ACM, 2004.
- [6] Michael Albert, Alan Frieze, and Bruce Reed. Multicoloured Hamilton cycles. *The Electronic Journal of Combinatorics*, 2(1):R10, 1995.
- [7] Noga Alon. A parallel algorithmic version of the local lemma. *Random Structures & Algorithms*, 2(4):367–378, 1991.
- [8] Noga Alon, Michael Krivelevich, and Benny Sudakov. Coloring graphs with sparse neighborhoods. *Journal of Combinatorial Theory, Series B*, 77(1):73–82, 1999.
- [9] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley Publishing, 4th edition, 2016.
- [10] Adrian Balint and Uwe Schöning. Choosing probability distributions for stochastic local search and the role of make versus break. In *Proceedings of the International Conference on Theory and Applications of Satisfiability Testing*, pages 16–29. Springer, 2012.
- [11] Adrian Balint and Uwe Schöning. Engineering a lightweight and efficient local search SAT solver. In Lasse Kliemann and Peter Sanders, editors, *Algorithm Engineering: Selected Results and Surveys*, volume 9220 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2016.
- [12] Anton Bernshteyn. New bounds for the acyclic chromatic index. *Discrete Mathematics*, 339(10):2543–2552, 2016.
- [13] Anton Bernshteyn. The Johansson–Molloy theorem for DP-coloring. *arXiv preprint arXiv:1708.03843*, 2017.
- [14] Anton Bernshteyn. The local cut lemma. *European Journal of Combinatorics*, 63:95–114, 2017.

- [15] Rodrigo Bissacot and Luís Doin. Entropy compression method and legitimate colorings in projective planes. *arXiv preprint arXiv:1710.06981*, 2017.
- [16] Rodrigo Bissacot, Roberto Fernández, Aldo Procacci, and Benedetto Scoppola. An improvement of the Lovász local lemma via cluster expansion. *Combinatorics, Probability & Computing*, 20(5):709–719, 2011.
- [17] Karthekeyan Chandrasekaran, Navin Goyal, and Bernhard Haeupler. Deterministic algorithms for the Lovász local lemma. *SIAM Journal on Computing*, 42(6):2132–2155, 2013.
- [18] Thomas F Coleman and Jin Yi Cai. The cyclic coloring problem and estimation of sparse Hessian matrices. *SIAM Journal on Algebraic Discrete Methods*, 7(2):221–235, 1986.
- [19] Thomas F. Coleman and Jorge J. Moré. Estimation of sparse Hessian matrices and graph coloring problems. *Mathematical Programming*, 28(3):243–270, 1984.
- [20] Harris David G., Fotis Iliopoulos, and Vladimir Kolmogorov. A new notion of commutativity for the algorithmic Lovász local lemma. *arXiv preprint arXiv:2008.05569*, 2020.
- [21] Ewan Davies, Ross J Kang, François Pirot, and Jean-Sébastien Sereni. An algorithmic framework for colouring locally sparse graphs. *arXiv preprint arXiv:2004.07151*, 2020.
- [22] Ewan Davies, Ross J Kang, François Pirot, and Jean-Sébastien Sereni. Graph structure via local occupancy. *arXiv preprint arXiv:2003.14361*, 2020.
- [23] Andrzej Dudek, Alan Frieze, and Andrzej Ruciński. Rainbow Hamilton cycles in uniform hypergraphs. *The Electronic Journal of Combinatorics*, 19(1):46, 2012.
- [24] Vida Dujmović, Gwenaél Joret, Jakub Kozik, and David R Wood. Nonrepetitive colouring via entropy compression. *Combinatorica*, 36(6):661–686, 2016.
- [25] Paul Erdős and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and Finite Sets (Colloq. Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vol. II, pages 609–627. Colloq. Math. Soc. János Bolyai, Vol. 10. North-Holland, Amsterdam, 1975.
- [26] Paul Erdős and Joel Spencer. Lopsided Lovász local lemma and Latin transversals. *Discrete Applied Mathematics*, 30(2-3):151–154, 1991.
- [27] Louis Esperet and Aline Parreau. Acyclic edge-coloring using entropy compression. *European Journal of Combinatorics*, 34(6):1019–1027, 2013.
- [28] Heidi Gebauer, Tibor Szabó, and Gábor Tardos. The local lemma is tight for SAT. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 664–674. Society for Industrial and Applied Mathematics, 2011.
- [29] Ioannis Giotis, Lefteris M. Kirousis, Kostas I. Psaromiligkos, and Dimitrios M. Thilikos. Acyclic edge coloring through the Lovász local lemma. *Theoretical Computer Science*, 665:40–50, 2017.
- [30] Adam Gkagol, Gwenaél Joret, Jakub Kozik, and Piotr Micek. Pathwidth and nonrepetitive list coloring. *arXiv preprint arXiv:1601.01886*, 2016.
- [31] Daniel Gonçalves, Mickaël Montassier, and Alexandre Pinlou. Entropy compression method applied to graph colorings. *arXiv preprint arXiv:1406.4380*, 2014.

- [32] Jarosław Grytczuk, Jakub Kozik, and Piotr Micek. New approach to nonrepetitive sequences. *Random Structures & Algorithms*, 42(2):214–225, 2013.
- [33] Bernhard Haeupler, Barna Saha, and Aravind Srinivasan. New constructive aspects of the Lovász local lemma. *J. ACM*, 58(6):28:1–28:28, 2011.
- [34] David G. Harris. Lopsidedependency in the Moser-Tardos framework: Beyond the lopsided Lovász local lemma. *ACM Transactions on Algorithms*, 13(1):17:1–17:26, 2016.
- [35] David G. Harris. Comparison of two convergence criteria for the variable-assignment lopsided Lovasz local lemma. *arXiv preprint arXiv:1610.01926*, 2018.
- [36] David G Harris. New bounds for the Moser-Tardos distribution. *Random Structures & Algorithms*, 51:97–131, 2020.
- [37] David G. Harris and Aravind Srinivasan. The Moser-Tardos framework with partial resampling. In *54th Annual IEEE Symposium on Foundations of Computer Science*, pages 469–478. IEEE Computer Society, 2013.
- [38] David G. Harris and Aravind Srinivasan. A constructive algorithm for the Lovász local lemma on permutations. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 907–925. Society for Industrial and Applied Mathematics, 2014.
- [39] David G. Harris and Aravind Srinivasan. Algorithmic and enumerative aspects of the Moser-Tardos distribution. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2004–2023. Society for Industrial and Applied Mathematics, 2016.
- [40] Nicholas J. A. Harvey and Jan Vondrák. An algorithmic proof of the Lovász local lemma via resampling oracles. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, pages 1327–1346. IEEE Computer Society, 2015.
- [41] Kun He, Liang Li, Xingwu Liu, Yuyi Wang, and Mingji Xia. Variable-version Lovász local lemma: Beyond Shearer’s bound. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, pages 451–462. IEEE Computer Society, 2017.
- [42] Fotis Iliopoulos. Commutative algorithms approximate the LLL-distribution. In *Proceedings of APPROX/RANDOM*, pages 44:1–44:20. Schloß Dagstuhl–Leibniz-Zentrum für Informatik, 2018.
- [43] Fotis Iliopoulos and Alistair Sinclair. Efficiently list-edge coloring multigraphs asymptotically optimally. *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2319–2336, 2020.
- [44] A. Johansson. Asymptotic choice number for triangle free graphs. Unpublished manuscript. 1996.
- [45] Jeff Kahn. Asymptotics of the chromatic index for multigraphs. *Journal of Combinatorial Theory, Series B*, 68(2):233–254, 1996.
- [46] Jeff Kahn. Asymptotics of the list-chromatic index for multigraphs. *Random Structures & Algorithms*, 17(2):117–156, 2000.
- [47] Kashyap Kolipaka, Mario Szegedy, and Yixin Xu. A sharper local lemma with improved applications. In *Proceedings of APPROX/RANDOM*, pages 603–614. Springer Berlin Heidelberg, 2012.

- [48] Kashyap Babu Rao Kolipaka and Mario Szegedy. Moser and Tardos meet Lovász. In *Proceedings of the 43rd Annual ACM Symposium on the Theory of Computing*, pages 235–244. ACM, 2011.
- [49] Vladimir Kolmogorov. Commutativity in the algorithmic Lovász local lemma. *SIAM Journal on Computing*, 47(6):2029–2056, 2018.
- [50] Vladimir Kolmogorov. Extensions of the algorithmic Lovász local lemma. *arXiv preprint arXiv:1910.01552*, 2019.
- [51] Michael Molloy. The list chromatic number of graphs with small clique number. *Journal of Combinatorial Theory, Series B*, 134:264–284, 2019.
- [52] Michael Molloy and Bruce Reed. Further algorithmic aspects of the local lemma. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*, pages 524–529. ACM, 1998.
- [53] Michael Molloy and Bruce Reed. *Graph Colouring and the Probabilistic Method*. Springer-Verlag, Berlin, 2002.
- [54] Robin A. Moser. A constructive proof of the Lovász local lemma. In *Proceedings of the 41st Annual ACM Symposium on the Theory of Computing*, pages 343–350. ACM, 2009.
- [55] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2):Art. 11, 15, 2010.
- [56] Sokol Ndreca, Aldo Procacci, and Benedetto Scoppola. Improved bounds on coloring of graphs. *Eur. J. Comb.*, 33(4):592–609, 2012.
- [57] Pascal Ochem and Alexandre Pinlou. Application of entropy compression in pattern avoidance. *arXiv preprint arXiv:1301.1873*, 2013.
- [58] Wesley Pegden. An extension of the Moser-Tardos algorithmic local lemma. *SIAM Journal on Discrete Mathematics*, 28(2):911–917, 2014.
- [59] Alexander D Scott and Alan D Sokal. The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma. *Journal of Statistical Physics*, 118(5-6):1151–1261, 2005.
- [60] Bart Selman, Henry A. Kautz, and Bram Cohen. Local search strategies for satisfiability testing. In *Proceedings of DIMACS Workshop on Cliques, Coloring and Satisfiability*, pages 521–532. DIMACS/AMS, 1993.
- [61] Bart Selman, Henry A. Kautz, and Bram Cohen. Noise strategies for improving local search. In *Proceedings of the 12th National Conference on Artificial Intelligence, Volume 1*, pages 337–343. AAAI Press/MIT Press, 1994.
- [62] J.B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3):241–245, 1985.
- [63] Van H Vu. A general upper bound on the list chromatic number of locally sparse graphs. *Combinatorics, Probability and Computing*, 11(1):103–111, 2002.

A Matrices and norms

Let $\|\cdot\|$ be any norm over vectors in \mathbb{R}^n . The *dual* norm, also over vectors in \mathbb{R}^n , is defined as

$$\|z\|_* = \sup_{\|x\|=1} |z^\top x|.$$

For example, the dual norm of $\|\cdot\|_\infty$ is $\|\cdot\|_1$. It can be seen that $\|\cdot\|_{**} = \|\cdot\|$ and that for any vectors x, z ,

$$z^\top x = \|x\| \left(\frac{z^\top x}{\|x\|} \right) \leq \|z\|_* \|x\|. \quad (29)$$

The corresponding *operator norm*, over $n \times n$ real matrices, is defined as

$$\|A\| \equiv \sup_{\|x\|=1} \|Ax\|.$$

For example, if A is a matrix with non-negative entries then $\|A\|_\infty$ and $\|A\|_1$ can be seen to be the maximum *row* and *column* sum of A , respectively. Operator norms are submultiplicative, i.e., for every operator norm $\|\cdot\|$ and any two $n \times n$ matrices A, B ,

$$\|AB\| \leq \|A\| \|B\|. \quad (30)$$

Finally, for any vector norm $\|\cdot\|$, any row vector x^\top and any $n \times n$ matrix A , we have

$$\|x^\top A\|_* \leq \|x^\top\|_* \|A\|. \quad (31)$$

B The matrix norms framework in action

In this appendix we illustrate how the framework of matrix norms captures a wide variety of convergence arguments for local search algorithms, both LLL-inspired ones and others. Recall from our discussion in the introduction that our goal is to bound $\rho(A)$, the spectral radius of the transition matrix between flawed states.

We begin with the classical potential function argument. Consider any function ϕ on Ω such that $\phi(\sigma) > 0$ for $\sigma \in \Omega^*$, while $\phi(\sigma) = 0$ for $\sigma \notin \Omega^*$. The potential argument asserts that eventually $\phi = 0$ (i.e., the particle escapes Ω^*) if ϕ is always reduced in expectation, i.e., if for every $\sigma \in \Omega^*$,

$$\sum_{\tau \in \Omega} P[\sigma, \tau] \phi(\tau) < \phi(\sigma). \quad (32)$$

To express this argument via matrix norms, let $A' = MAM^{-1}$ where M is the diagonal $|\Omega^*| \times |\Omega^*|$ matrix $\text{diag}(1/\phi(\sigma))$. Thus, $A'[\sigma, \tau] = A[\sigma, \tau] \phi(\tau) / \phi(\sigma)$. Recalling that $\|\cdot\|_\infty$ is the maximum row sum of a matrix, we see that condition (32) for the potential function is nothing other than $\|A'\|_\infty < 1$, implying $\rho(A) = \rho(A') \leq \|A'\|_\infty < 1$.

Next we show how the same approach, using the dual matrix norm $\|\cdot\|_1$, captures the Moser-Tardos algorithm for k -SAT. Given a k -SAT formula on n variables with clauses c_i , let $\Omega = \{0, 1\}^n$ denote the set of all assignments, and Ω^* the set of non-satisfying assignments. To simplify exposition, we assume that in each step, the algorithm picks the lowest-indexed unsatisfied clause c_i and resamples all variables in c_i . Thus, the state-evolution is a Markov chain, and if we denote its transition matrix by P , we are interested in the submatrix A that is the projection of P onto Ω^* . For each clause c_i , let A_i be the $|\Omega^*| \times |\Omega^*|$ submatrix of A comprising all rows (states) where the resampled clause is c_i . (All other rows of A_i are 0.) For $t \geq 1$,

let \mathcal{W}_t contain every t -sequence of (indices of) clauses that has non-zero probability of being the first t clauses resampled by the algorithm. In other words, \mathcal{W}_t is the set of all t -sequences of indices from $[m]$ corresponding to non-vanishing t -products of matrices from $\{A_1, \dots, A_m\}$, i.e., $\mathcal{W}_t = \{W = (w_i) \in [m]^t : \prod_{i=1}^t A_{w_i} \neq 0\}$. For every operator norm $\|\cdot\|$ we get:

$$\rho(A)^t = \rho(A^t) \leq \|A^t\| = \left\| \left(\sum_{i \in [m]} A_i \right)^t \right\| = \left\| \sum_{W \in \mathcal{W}_t} \prod_{i=1}^t A_{w_i} \right\| \leq \sum_{W \in \mathcal{W}_t} \left\| \prod_{i=1}^t A_{w_i} \right\| \leq \sum_{W \in \mathcal{W}_t} \prod_{i=1}^t \|A_{w_i}\|. \quad (33)$$

The first inequality here follows from the fact that $\rho(A) \leq \|A\|$ for any operator norm $\|\cdot\|$, the second inequality is the triangle inequality, and the third follows by submultiplicativity of operator norms.

To get a favorable bound, we will apply (33) with the norm $\|\cdot\|_1$, i.e., the maximum column sum. We see that for all $j \in [m]$, every column of A_j has at most one non-zero entry, since $A_j(\sigma, \tau) > 0$ only if σ is the (unique) mutation of τ so that c_j is violated. Recalling that all non-zero entries of A equal 2^{-k} , we conclude that $\|A_j\|_1 = 2^{-k}$ for all $j \in [m]$. Therefore, $\|A^t\|_1 \leq |\mathcal{W}_t| 2^{-kt}$. To bound $|\mathcal{W}_t|$ we use a simple necessary condition for membership in \mathcal{W}_t which, by a standard counting argument, implies that if each clause shares variables with at most Δ other clauses then $|\mathcal{W}_t| \leq 2^m (e\Delta)^t$. Therefore $\rho(A)^t \leq 2^m (e\Delta 2^{-k})^t$, implying that if $\Delta < 2^k/e$ then $1 > \|A\|_1 \geq \rho(A)$ and the algorithm terminates within $O(m)$ steps with high probability. This matches exactly the Moser-Tardos condition (which is tight).

A very similar argument can be used to capture even the most general existing versions of the algorithmic LLL [2, 40, 3], which are described by arbitrary flaws and, for each flaw f_i , an arbitrary corresponding transition matrix A_i for addressing the flaw. Note that (33) is in essence a weighted counting of witness sequences, the weight of each sequence being the product of the norms $\|A_{w_i}\|_1$. Observe also that in our k -SAT example above, the only probabilistic notion was the transition matrix A and we did not make any implicit or explicit reference to a probability measure μ . To cast general algorithmic LLL arguments in this same form, any measure μ is incorporated as a *change of basis* for the transition matrix A , i.e., we bound $\|A'\|_1 = \|MAM^{-1}\|_1$ as $\sum_{W \in \mathcal{W}_t} \prod_{i=1}^t \|MA_{w_i}M^{-1}\|_1$, where M is the diagonal $|\Omega^*| \times |\Omega^*|$ matrix $\text{diag}(\mu(\sigma))$, similarly to the potential function argument. We thus see that the measure μ is nothing other than a tool for analyzing the progress of the algorithm.

C Comparison with previous LLL conditions

Here we give background on the LLL and explain how our main theorem compares with previous algorithmic LLL conditions.

C.1 Non-constructive conditions

The original statement of the LLL asserts, roughly, that, given a family of “bad” events in a probability space, if each bad event individually is not very likely and, in addition, is independent of all but a small number of other bad events, then the probability of avoiding all bad events is strictly positive. Erdős and Spencer [26] later noted that independence in the LLL can be replaced by positive correlation, yielding the original version of what is known as the *Lopsided* LLL, more sophisticated versions of which have also been established in [6, 23]. Below we state the Lopsided LLL in its most powerful form that holds for arbitrary probability spaces and families of bad events (see e.g., [53, p.228]).

General LLL. Let (Ω, μ) be a probability space and let $\mathcal{B} = \{B_1, B_2, \dots, B_m\}$ be a set of m (bad) events. For each $i \in [m]$, let $L(i) \subseteq [m] \setminus \{i\}$ be such that $\mu(B_i \mid \bigcap_{j \in S} \overline{B_j}) \leq b_i$ for every $S \subseteq [m] \setminus (L(i) \cup \{i\})$. If there exist positive real numbers $\{\psi_i\}_{i=1}^m$ such that, for all $i \in [m]$,

$$\frac{b_i}{\psi_i} \sum_{S \subseteq L(i) \cup \{i\}} \prod_{j \in S} \psi_j \leq 1, \quad (34)$$

then the probability that none of the events in \mathcal{B} occurs is at least $\prod_{i=1}^m 1/(1 + \psi_i) > 0$.

Writing $x_i = \psi_i/(1 + \psi_i) \in [0, 1)$, condition (34) takes the more familiar form $b_i \leq x_i \prod_{j \in L(i)} (1 - x_j)$. The form in (34), though, is more amenable to refinement (under additional assumptions) and comparison.

The above form of the LLL is motivated by the fact that, in complex applications, small but non-vanishing correlations tend to travel arbitrarily far in the space Ω . To isolate these dependencies so that they can be treated locally, it can be crucial [6, 23, 45, 46] to allow mild negative correlations between each bad event B_i and the events outside its “special” set $L(i)$, achieved by allowing $b_i \geq \mu(B_i)$.

C.1.1 Extensions

In the setting of the General LLL above, we often define the *lopsidependency* graph, which is the directed graph on $[m]$ in which each vertex i points to the vertices in $L(i)$. In many applications the lopsidependency graph is in fact symmetric, and even when it is not it is useful to consider an undirected version in which the neighborhood of each vertex $i \in [m]$ includes $L(i)$ and possibly some other vertices. (One can trivially achieve this by simply ignoring the direction of the edges in the lopsidependency graph.) Given such an undirected graph G , one can restrict the sum in (34) to $S \subseteq L(i)$ that are *independent* in G . This yields the so-called *cluster expansion* condition [16]. It should be clear that, if the lopsidependency graph is already symmetric (so that the neighborhood of i is exactly $L(i)$), then this leads to an improvement over (34).

In fact, for a given graph G as above, the *exact* condition for avoiding all bad events as a function of the bounds $\{b_i\}_{i=1}^m$ is known, and is due to Shearer [62]. Unlike the cluster-expansion condition, though, Shearer’s condition involves a separate condition for every independent set in G . Moreover, when μ is a product measure and the dependencies between events can be expressed in terms of variable sharing (that is, under the assumptions of the variable setting [55]), several works [35, 41, 48] have shown that Shearer’s condition can be improved, i.e., that more permissive conditions exist.

Finally, Scott and Sokal [59] introduced a so-called *soft-core* LLL condition, in an effort to quantify interaction strengths between bad events (whereas in all other works two bad events either interact or they don’t). Unlike our present work, which quantifies general point-to-set interactions, the condition in [59] only quantifies *pairwise* (i.e., point-to-point) interactions. Finding combinatorial applications for the soft-core LLL condition was left as an open question in [59]. To the best of our knowledge, it remains open.

C.2 Constructive conditions

Using the framework of Section 1.1 for local search algorithms, we say that flaw f_i *causes* flaw f_j , and write $i \rightarrow j$, if there exist $\sigma \in f_i, \tau \in f_j$ such that $\rho_i(\sigma, \tau) > 0$ and the transition $\sigma \rightarrow \tau$ introduces flaw f_j . (Thus, causality is to be interpreted as *potential* causality.) Let $\Gamma(i) = \{j : i \rightarrow j\}$ be the set of flaws caused by f_i . We call the digraph over $[m]$ in which each vertex i points to the vertices in $\Gamma(i)$ the *causality* digraph.

Let μ be an arbitrary probability measure on Ω . For $i \in [m]$ and $\tau \in \Omega$, let $\nu_i(\tau)$ be the probability of ending up in state τ after sampling a state $\sigma \in f_i$ according to μ , and then addressing f_i at σ . The *distortion* associated with f_i is defined as

$$d_i := \max_{\tau \in \Omega} \frac{\nu_i(\tau)}{\mu(\tau)} \geq 1.$$

Thus, the distortion of μ associated with flaw f_i is the greatest inflation of a state probability induced by sampling $\sigma \in f_i$ according to μ and addressing flaw f_i at σ . As shown in [43], if we set $b_i = \mu(f_i)d_i$ for all $i \in [m]$, then a causality digraph is a lopsidedependency graph for μ , thus making the following algorithmic LLL condition, established in [3], the algorithmic counterpart of the General LLL (34):

Algorithmic LLL. Let $\gamma_i := \mu(f_i)d_i$. If there exist positive real numbers $\{\psi_i\}_{i=1}^m$ such that, for all $i \in [m]$,

$$\frac{\gamma_i}{\psi_i} \sum_{S \subseteq \Gamma(i)} \prod_{j \in S} \psi_j < 1, \quad (35)$$

then a local search algorithm as above reaches a flawless object quickly with high probability.

If $d_i = 1$, i.e., $\nu_i(\tau) = \mu(\tau)$ for all $\tau \in \Omega$, we say that the algorithm is a *resampling oracle* [40] for f_i . When one has resampling oracles for all flaws, i.e., $d_i = 1$ for all $i \in [m]$, then condition (34) is the algorithmic counterpart of the Lopsided LLL, as established by Harvey and Vondrák [40]. Condition (35) also subsumes the flaws/actions condition of [2]: in that setting $\rho_i(\sigma, \cdot)$ is uniform over the set of possible next states, while the analysis does not reference a measure μ . Taking μ to be uniform and applying condition (35) in fact sharpens the condition of [2].

Observe that a resampling oracle perfectly removes the conditioning on the old state σ belonging to f_i , since the new state τ is distributed according to μ . For example, if μ is a product measure, this is precisely what is achieved by the resampling algorithm of Moser and Tardos [55]. Requiring an algorithm to be a resampling oracle for every flaw may be impossible to achieve by local exploration within Ω , i.e., by “local search.” (Note that restricting to local search is crucial since longer-range resampling, even if it were possible, would tend to rapidly densify the causality digraph.) Allowing distortion greater than 1 frees the algorithm from the strict resampling requirement of perfect deconditioning. Optimizing the tradeoff between distortion and the density of the causality digraph has recently led to strong algorithmic results [3, 42, 43, 51].

C.2.1 Our new condition

Our Theorem 1.2 is a strict generalization of the above algorithmic LLL condition (35). To see this, observe that $\gamma_i^S \leq \gamma_i^\emptyset = \gamma_i$ for every set $S \subseteq [m]$, as

$$\gamma_i^S \leq \max_{\tau \in \Omega} \left\{ \frac{1}{\mu(\tau)} \sum_{\sigma \in f_i} \mu(\sigma) \rho_i(\sigma, \tau) \right\} = \max_{\tau \in \Omega} \frac{\nu(\tau)}{\mu(\tau)} \mu(f_i) = d_i \mu(f_i) = \gamma_i.$$

Hence, since $\gamma_i^S = 0$ for $S \not\subseteq \Gamma(i)$, the l.h.s. of our condition (4) is never larger than the l.h.s. of (35).

C.2.2 Extensions

Condition (35) can be improved under additional assumptions. Specifically, let G be an *undirected* graph on $[m]$ such that $\Gamma(i)$ is a subset of the neighbors of i , for every $i \in [m]$. (As discussed in Section C.1.1, one can again trivially get such a G by ignoring the direction of edges in the causality graph.) In [3] it was shown that condition (35) can be replaced by the cluster expansion condition [16] on G , while in [49] it was shown that condition (35) can be replaced by Shearer’s condition [62]. Both of these conditions benefit by restricting consideration to independent sets of G , as discussed in Section C.1.1. Finally, Kolmogorov [50] showed a new condition that interpolates between the cluster expansion condition and our new condition.

Harris and Srinivasan [34, 37] have developed improved conditions for the convergence of algorithms operating in the so-called variable setting [55], based on refinements of the notion of dependency between bad events. These improvements are incomparable to condition (35), as they do not apply to general local search algorithms (for instance, all algorithms in the variable setting are commutative [49]).

D Proof of Lemma 4.4

Our computations are similar to the ones in [51]. The following version of Chernoff Bounds will be useful:

Lemma D.1. *Suppose $\{X_i\}_{i=1}^m \in \{0, 1\}$ are boolean variables, and set $Y_i = 1 - X_i$, $X = \sum_{i=1}^m X_i$. If $\{Y_i\}_{i=1}^m$ are negatively correlated, then for any $0 < t \leq \mathbb{E}[X]$*

$$\Pr[|X - \mathbb{E}[X]| > t] < 2 \exp\left(-\frac{t^2}{3\mathbb{E}[X]}\right).$$

Proof of part (a) of Lemma 4.4. Let $v \in V$ and $\sigma \in \Omega$ be arbitrary and let $\tau \in \Omega$ be the (random) output (state) of $\text{RECOLOR}(v, \sigma)$. For each color $c \in \mathcal{L}_v$, let $P_v^c = \{u \in N_v : c \in R_u^v(\sigma)\}$ and define

$$\rho(c) = \sum_{u \in P_v^c} \frac{1}{|R_u^v(\sigma)| - 1}.$$

Since $c \in R_u^v(\sigma)$ implies $|R_u^v(\sigma)| \geq 2$, and since $1 - 1/x > \exp(-1/(x-1))$ for $x \geq 2$, we see that

$$\mathbb{E}[|L_v(\tau)|] = 1 + \sum_{c \in \mathcal{L}_v} \prod_{u \in P_v^c} \left(1 - \frac{1}{|R_u^v(\sigma)|}\right) > \sum_{c \in \mathcal{L}_v} \prod_{u \in P_v^c} \exp\left(-\frac{1}{|R_u^v(\sigma)| - 1}\right) = \sum_{c \in \mathcal{L}_v} e^{-\rho(c)}. \quad (36)$$

Also, since each $R_u^v(\sigma)$ has $|R_u^v(\sigma)| - 1$ non-Blank colors, we see that

$$Z_v := \sum_{c \in \mathcal{L}_v} \rho(c) \leq \sum_{u \in N_v} \sum_{c \in R_u^v(\sigma) \setminus \text{Blank}} \frac{1}{|R_u^v(\sigma)| - 1} \leq \Delta. \quad (37)$$

The fact that e^{-x} is convex implies that the right hand side of (36) is at least $|\mathcal{L}_v| \exp(-Z_v/|\mathcal{L}_v|)$. Recalling that $|\mathcal{L}_v| = q = (1 + \epsilon) \frac{\Delta}{\ln \sqrt{f}}$ and combining (36) with (37) yields

$$\mathbb{E}[|L_v(\tau)|] > q e^{-Z_v/q} \geq (1 + \epsilon) \frac{\Delta}{\ln \sqrt{f}} e^{-\Delta/q} = 2(1 + \epsilon) \frac{\Delta}{\ln f} f^{-\frac{1}{2(1+\epsilon)}} = 2L.$$

Let X_c be the indicator variable that $c \in L_v(\tau)$ so that $|L_v(\tau)| = 1 + \sum_{c \in \mathcal{L}_v(\tau)} X_c$. It is not hard to see that the variables $Y_c = 1 - X_c$ are negatively correlated, so that applying Lemma D.1 with $t = \frac{1}{2} \mathbb{E}[|L_v(\tau)|] > L$ yields

$$\Pr[|L_v(\tau)| < \frac{1}{2} \mathbb{E}[|L_v(\tau)|]] \leq 2e^{-\mathbb{E}[|L_v(\tau)|]/12} < 2e^{-L/6}.$$

This concludes the proof. \square

Proof of part (b) of Lemma 4.4. Let $\Psi = \{c \in \mathcal{L}_v(\sigma) : \rho(c) \geq L/20\} \setminus \text{Blank}$. The probability that $L_v(\tau)$ contains at least one color from Ψ is at most

$$\mathbb{E}[|L_v(\tau) \cap \Psi|] = \sum_{c \in \Psi} \prod_{u \in P_v^c} \left(1 - \frac{1}{|R_u^v(\sigma)|}\right) < \sum_{c \in \Psi} \prod_{u \in P_v^c} \exp\left(-\frac{1}{2(|R_u^v(\sigma)| - 1)}\right) < \sum_{c \in \Psi} e^{-\rho(c)/2},$$

where we used that $c \in R_u^v(\sigma)$ implies $|R_u^v(\sigma)| \geq 2$, and that $1 - 1/x < \exp(-1/(2(x-1)))$ for $x \geq 2$. Finally note that $\sum_{c \in \Psi} e^{-\rho(c)/2} \leq q e^{-L/40}$ by the definition of the set Ψ .

Recall that $T_{v,c}(\tau) = \{u \in N_v : \tau(u) = \text{Blank and } c \in L_u(\tau)\}$. Since $L_u(\tau) \subseteq R_u(\tau) = R_u^v(\sigma)$, it follows that $T_{v,c}(\tau) \subseteq P_v^c$ and, therefore, $\mathbb{E}[|T_{v,c}(\tau)|] \leq \sum_{u \in P_v^c} 1/|R_u^v(\sigma)| \leq \rho(c)$. Since the vertices in P_v^c are colored (and thus become Blank) independently and since $\rho(c) < L/20$ for $c \notin \Psi$, applying

Lemma D.1 with $t = L/20$ yields $\Pr [|T_{v,c}(\tau)| > \mathbb{E} [|T_{v,c}(\tau)|] + L/20] < 2e^{-L/60}$. Applying the union bound over all q colors, we see that the probability there is at least one $c \notin \Psi$ for which $|T_{v,c}(\tau)| > L/10$ is at most $2qe^{-L/60}$. Thus, with probability at least $1 - 3qe^{-L/60}$,

$$\sum_{c \in L_v(\tau) \setminus \text{Blank}} |T_{v,c}(\tau)| = \sum_{c \in L_v(\tau) \setminus (\Psi \cup \text{Blank})} |T_{v,c}(\tau)| < \frac{L}{10} |L_v(\tau)|.$$

□

E Proof of Theorem 2.6

We will follow closely the approach adopted by the authors in [8]. Throughout the proof we assume that $\epsilon \in (0, \epsilon_0)$, where ϵ_0 is sufficiently small, and that $f_\epsilon > 0$ and $\Delta_\epsilon > 0$ are sufficiently large.

We distinguish two cases, depending on whether $f \geq \Delta^{(2+\epsilon^2)\epsilon}$ or not. To prove Theorem 2.6 for the case $f \geq \Delta^{(2+\epsilon^2)\epsilon}$ we will prove the following.

Theorem E.1. *For every $\theta, \zeta \in (0, 1)$, there exists $\Delta_{\theta, \zeta} > 0$ such that every graph G with maximum degree $\Delta \geq \Delta_{\theta, \zeta}$ in which the neighbors of every vertex span at most $\Delta^{2-(2+\zeta)\theta}$ edges, has chromatic number $\chi(G) \leq (1 + \zeta)(1 + \theta^{-1}) \frac{\Delta}{\ln \Delta}$.*

Proof of Theorem 2.6 for $f \geq \Delta^{(2+\epsilon^2)\epsilon}$. We apply Theorem E.1 with $\zeta = \epsilon^2$ and $\theta = \frac{\ln f}{(2+\epsilon^2) \ln \Delta} \geq \epsilon$, so that $\Delta^2/f = \Delta^{2-(2+\zeta)\theta}$. Since $\zeta, \theta < 1$, we obtain

$$\begin{aligned} \chi(G) &\leq (1 + \zeta) \left(1 + \frac{(2 + \zeta) \ln \Delta}{\ln f} \right) \frac{\Delta}{\ln \Delta} \\ &= (1 + \zeta) \frac{\Delta}{\ln \Delta} + (1 + \zeta)(1 + \zeta/2) \frac{\Delta}{\ln \sqrt{f}} \\ &\leq (1 + 2\zeta) \frac{\Delta}{\ln \sqrt{f}} + (1 + \zeta)(1 + \zeta/2) \frac{\Delta}{\ln \sqrt{f}} \\ &= \left(2 + \frac{7\zeta}{2} + \frac{\zeta^2}{2} \right) \frac{\Delta}{\ln \sqrt{f}} \\ &\leq (2 + \epsilon) \frac{\Delta}{\ln \sqrt{f}}. \end{aligned}$$

□

Theorem E.1 follows immediately from the following lemma, whose proof is similar to Lemma 2.3 in [8] and can be found in Section E.1. The proof of Lemma E.2 uses the standard Local Lemma and Theorem 2.5, so it can be made constructive using the Moser-Tardos algorithm and the algorithm in Theorem 2.5.

Lemma E.2. *For every $\theta, \zeta \in (0, 1)$ there exists $\Delta_{\theta, \zeta} > 0$ such that for every graph $G = (V, E)$ with maximum degree $\Delta \geq \Delta_{\theta, \zeta}$ in which the neighbors of every vertex span at most $\Delta^{2-(2+\zeta)\theta}$ edges, there exists a partition of the vertex set $V = V_1 \cup \dots \cup V_k$ with $k = \Delta^{1-\theta}$, such that for every $1 \leq i \leq k$,*

$$\chi(G[V_i]) \leq (1 + \zeta)(1 + \theta^{-1}) \frac{\Delta^\theta}{\ln \Delta}.$$

Proof of Theorem E.1. If $V_1, V_2, \dots, V_k, k = \Delta^{1-\theta}$ is the partition promised by Lemma E.2 then

$$\chi(G) \leq \sum_{i=1}^{\Delta^{1-\theta}} \chi(G[V_i]) \leq (1 + \zeta)(1 + \theta^{-1}) \frac{\Delta}{\ln \Delta}.$$

□

To prove Theorem 2.6 for $f \in [f_\epsilon, \Delta^{(2+\epsilon^2)\epsilon}]$, we will perform a sequence of random halving steps, as in [8], to partition the graph into subgraphs satisfying the condition of Theorem 2.6 with $f \geq \Delta^{(2+\epsilon^2)\epsilon}$ and color these subgraphs using disjoint sets of colors. To perform the partition we use the following lemma from [8]. As it is proven via the standard LLL, it can be made constructive using the Moser-Tardos algorithm.

Lemma E.3 ([8]). *Let $G(V, E)$ be a graph with maximum degree $\Delta \geq 2$ in which the neighbors of every vertex span at most s edges. There exists a partition $V = V_1 \cup V_2$ such that the induced subgraph $G[V_i]$, $i = 1, 2$, has maximum degree at most $\Delta/2 + 2\sqrt{\Delta \ln \Delta}$ and the neighborhoods of every vertex in $G[V_i]$, $i = 1, 2$, span at most $s/4 + 2\Delta^{\frac{3}{2}}\sqrt{\ln \Delta}$ edges.*

We will also use the following lemma whose proof, presented in Section E.2, is almost identical to a similar statement in the proof of Theorem 1.1 of [8].

Lemma E.4. *Given Δ, f sufficiently large, let the sequences Δ_t and s_t be defined as follows. $\Delta_0 = \Delta, s_0 = \Delta^2/f$ and*

$$\Delta_{t+1} = \Delta_t/2 + 2\sqrt{\Delta_t \ln \Delta_t}, \quad s_{t+1} = s_t/4 + 2\Delta_t^{\frac{3}{2}}\sqrt{\ln \Delta_t}.$$

For any $\delta \in (0, 1/100)$ and $\zeta > 0$ such that $\zeta(2 + \delta) < 1/10$, let j be the smallest integer for which $f > \left(\frac{(1+\delta)\Delta}{2^j}\right)^{(2+\delta)\zeta}$. Then $\Delta_j \leq (1 + \delta)\Delta/2^j$ and $s_j \leq ((1 + \delta)\Delta/2^j)^2/f$.

Proof of Theorem 2.6 for $f \in [f_\epsilon, \Delta^{(2+\epsilon^2)\epsilon}]$. Let $\epsilon_0 = 1/11$. For $\epsilon \in (0, \epsilon_0]$, let $\delta = \zeta = \epsilon^2$. Since $\zeta(2 + \delta) < 1/10$, apply Lemma E.4 and let $j = j(\Delta, f, \delta, \zeta)$ be the integer described therein. Let S be the process which, given a graph G , does nothing if $\Delta(G) < 2$, and otherwise partitions G as described in Lemma E.3. Apply S to G to get subgraphs $G[V_1], G[V_2]$. Apply S to $G[V_1], G[V_2]$ to get $G[V_{1,1}], G[V_{1,2}], G[V_{2,1}], G[V_{2,2}]$. Repeating these steps j times, we obtain a partition of G into at most 2^j induced subgraphs. Observe that for each such subgraph H , either $\Delta(G) < 2$ and hence $\chi(H) \leq 2$, or, by Lemma E.4, $\Delta(H) \leq (1 + \delta)\Delta/2^j =: \Delta_*$ and the neighbors of every vertex in H span at most Δ_*^2/f edges, where $f \geq \Delta_*^{(2+\delta)\zeta} = \Delta_*^{(2+\zeta)\zeta} \geq \Delta_*^{(2+\zeta^2)\zeta}$. Therefore, by the already established case of Theorem 2.6, either $\chi(H) \leq 2$ or $\chi(H) \leq (2 + \zeta)\Delta_*/\ln \sqrt{f}$. Thus,

$$\chi(G) \leq 2^j \max \left\{ 2, (2 + \zeta) \frac{(1 + \delta)\Delta/2^j}{\ln \sqrt{f}} \right\} \leq \max \left\{ 2^{j+1}, (2 + \zeta) \frac{(1 + \delta)\Delta}{\ln \sqrt{f}} \right\}.$$

To bound 2^{j+1} from above we first observe that, for all f sufficiently large, i.e., for all $f \geq f_\epsilon$,

$$\left(\frac{(1 + \delta)\Delta}{2^j} \right)^{(2+\delta)\zeta} = \left(2(1 + \delta) \ln \sqrt{f} \right)^{(2+\delta)\zeta} < f. \quad (38)$$

Now, since j was defined as the smallest integer for which $\left(\frac{(1+\delta)\Delta}{2^j}\right)^{(2+\delta)\zeta} < f$, we see that (38) implies $2^j \leq \frac{\Delta}{2 \ln \sqrt{f}}$ and, therefore, $2^{j+1} \leq \frac{\Delta}{\ln \sqrt{f}}$. Finally, we observe that $(2 + \zeta)(1 + \delta) = (2 + \epsilon^2)(1 + \epsilon^2) < 2 + \epsilon$ for all $\epsilon \in (0, \epsilon_0]$. Therefore, as claimed,

$$\chi(G) \leq (2 + \epsilon) \frac{\Delta}{\ln \sqrt{f}}.$$

□

E.1 Proof of Lemma E.2

We follow an approach similar to that of Lemma 2.3 in [8], making appropriate modifications as needed. First we partition the vertices of G into $\Delta^{1-\theta}$ parts by coloring them randomly and independently with $\Delta^{1-\theta}$ colors. For a vertex v , call a neighbor u of v a *bad neighbor* if u and v have at least $\Delta^{1-(1+\zeta/2)\theta}$ common neighbors. Otherwise, say that u is a *good neighbor*. Since the neighbors of every vertex span at most $\Delta^{2-(2+\zeta)\theta}$ edges, there are at most $2\Delta^{1-(1+\zeta/2)\theta}$ bad neighbors for any vertex in G .

For any vertex v , define three types of bad event with respect to the random partitioning experiment:

- A_v : more than $(1 + \theta)\Delta^\theta$ neighbors of v receive the same color as v .
- B_v : more than $\frac{10}{\theta\zeta}$ bad neighbors of v receive the same color as v .
- C_v : the good neighbors of v that receive the same color as v span more than $\frac{100}{(\theta\zeta)^2}$ edges.

We will use the symmetric version of the Local Lemma [25] to show that we can find a coloring of the graph that avoids all bad events. First, note that each of the bad events A_v, B_v, C_v is independent of all but at most Δ^2 others, as it is independent of all events A_u, B_u, C_u corresponding to vertices u whose distance from v is more than 2. Since the degree of any vertex in its color class is binomially distributed with mean at most Δ^θ , standard Chernoff estimates imply that the probability that v has more than $(1 + \theta)\Delta^\theta$ neighbors with the same color as v is at most $e^{-\Omega(\Delta^\theta)}$, which means that $\Pr[A_v] < \Delta^{-3}$ for large enough Δ . Moreover, we also have

$$\Pr[B_v] \leq \binom{2\Delta^{1-(1+\zeta/2)\theta}}{\frac{10}{\theta\zeta}} \left(\frac{1}{\Delta^{1-\theta}} \right)^{\frac{10}{\theta\zeta}} \leq \left(\frac{2}{\Delta^{\frac{\theta\zeta}{2}}} \right)^{\frac{10}{\theta\zeta}} \leq \Delta^{-3},$$

for large enough Δ . Finally, to bound the probability of C_v we make the following observation. If a graph has at least e^2 edges, then either it has a vertex of degree at least e , or every vertex has degree strictly less than e , implying that the graph can be edge-colored with $\lceil e \rceil$ colors, in which case the largest color class must contain at least $e^2/\lceil e \rceil \geq e - 1$ edges. Thus, a graph with more than $100/(\theta\zeta)^2$ edges either has a vertex of degree at least $10/(\theta\zeta) \geq 9/(\theta\zeta)$ or a matching with at least $10/(\theta\zeta) - 1 \geq 9/(\theta\zeta)$ edges, where the inequality follows from the fact that $\theta, \zeta < 1$. Thus, C_v can happen only if there is a good neighbor u of v such that u and v have at least $9/(\theta\zeta)$ common neighbors with the same color as v , or if there is a matching of size at least $9/(\theta\zeta)$ on the good neighbors of v that have the same color as v . The probabilities of the first and second of these events are bounded, respectively, by

$$\Delta \binom{\Delta^{1-(1+\zeta/2)\theta}}{\frac{9}{\theta\zeta}} \left(\frac{1}{\Delta^{1-\theta}} \right)^{\frac{9}{\theta\zeta}} \leq \left(\frac{1}{\Delta^{\frac{\theta\zeta}{2}}} \right)^{\frac{9}{\theta\zeta}} \leq \frac{1}{2}\Delta^{-3},$$

$$\binom{\Delta^{2-(2+\zeta)\theta}}{\frac{9}{\theta\zeta}} \left(\left(\frac{1}{\Delta^{1-\theta}} \right)^2 \right)^{\frac{9}{\theta\zeta}} \leq \left(\frac{1}{\Delta^{\theta\zeta}} \right)^{\frac{9}{\theta\zeta}} \leq \frac{1}{2}\Delta^{-3}.$$

Therefore the probability of C_v is at most Δ^{-3} . Thus, the Local Lemma applies since each bad event has probability at most Δ^{-3} and is independent of all but at most Δ^2 other bad events. This means that we can find a partition $V = V_1, \dots, V_k$, where $k = \Delta^{1-\theta}$, so that in each induced subgraph $G[V_i]$, every vertex has degree at most $(1 + \theta)\Delta^\theta$, has at most $\frac{10}{\theta\zeta}$ bad neighbors, and is contained in at most $\frac{100}{(\theta\zeta)^2}$ triangles in which both other vertices are good. We will show that, given such a partition, each $G[V_i]$ can be colored with at most $\frac{(1+\zeta)(1+\theta^{-1})\Delta^\theta}{\ln \Delta}$ colors, assuming Δ is large enough.

To see this, consider the partition $B_i, V_i \setminus B_i$ of V_i , where B_i is the set of vertices $u \in V_i$ for which there exists a vertex $v \in V_i$, such that u is a bad neighbor of v . We claim that $\chi(G[B_i]) \leq \frac{20}{\theta\zeta} + 1$ and $\chi(G[V_i \setminus B_i]) \leq \frac{(1+\zeta/2)(1+\theta^{-1})\Delta^\theta}{\ln \Delta}$. Assuming this claim, observe that

$$\chi(G[V_i]) \leq \frac{20}{\theta\zeta} + 1 + \frac{(1+\zeta/2)(1+\theta^{-1})\Delta^\theta}{\ln \Delta} \leq \frac{(1+\zeta)(1+\theta^{-1})\Delta^\theta}{\ln \Delta},$$

where the last inequality holds for all $\Delta \geq \Delta_{\theta,\zeta}$.

To see the first part of the claim, note that it is well-known (and easy to see) that if a graph has an orientation with maximum outdegree d , then it is $(2d+1)$ -colorable. Consider the orientation of the graph on B_i that results when every vertex points to its bad neighbors in B_i . Clearly, the maximum outdegree is at most $\frac{10}{\theta\zeta}$ and, thus, $\chi(G[B_i]) \leq \frac{20}{\theta\zeta} + 1$.

To see the second part of the claim, observe that each vertex of $G[V_i \setminus B_i]$ is contained in at most $\frac{100}{(\theta\zeta)^2}$ triangles. Let $\Delta_* = (1+\theta)\Delta^\theta$ and

$$f = \frac{((1+\theta)\Delta^\theta)^2}{100/(\theta\zeta)^2} = \frac{(\theta\zeta)^2 \Delta_*^2}{100} \geq \Delta_*^{\frac{2+\frac{2\zeta}{3}}{1+\frac{2\zeta}{3}}} (\ln \Delta_*)^2,$$

where the last inequality holds for $\Delta \geq \Delta_{\theta,\zeta}$. Applying Theorem 2.5 to $G[V_i \setminus B_i]$ (plugging in $\zeta/3$ for the ϵ in that theorem) we get that, for all $\Delta \geq \Delta_{\theta,\zeta}$,

$$\begin{aligned} \chi_\ell(G[V_i \setminus B_i]) &\leq (1+\zeta/3) \frac{\Delta_*}{\ln \sqrt{f}} \\ &= (1+\zeta/3) \frac{(1+\theta)\Delta^\theta}{\ln \frac{(1+\theta)\theta\zeta\Delta^\theta}{10}} \\ &= (1+\zeta/3) \frac{(1+\theta)\Delta^\theta}{\theta \ln \Delta + \ln \frac{(1+\theta)\theta\zeta}{10}} \\ &\leq (1+\zeta/2) \frac{(1+\theta)}{\theta} \frac{\Delta^\theta}{\ln \Delta} \\ &= (1+\zeta/2) \frac{(1+\theta^{-1})\Delta^\theta}{\ln \Delta}, \end{aligned}$$

as claimed.

E.2 Proof of Lemma E.4

Let $\epsilon' := \zeta(2+\delta)$ and recall that $\epsilon' < \frac{1}{10}$ by hypothesis. By the definition of j , for every $t < j$,

$$\Delta_t \geq \Delta/(2^t) > \frac{f^{\frac{1}{\epsilon'}}}{1+\delta},$$

and $f^{\frac{1}{\epsilon'}}/(1+\delta)$ can be made arbitrarily large by taking f sufficiently large. Hence, we can assume that Δ_t is sufficiently large so that $\Delta_{t+1} \leq \frac{\Delta_t}{2} + \Delta_t^{\frac{2}{3}} \leq \frac{1}{2}(\Delta_t^{\frac{1}{3}} + 1)^3$. Taking cube roots and subtracting $\frac{1}{2^{\frac{1}{3}-1}}$ from both sides we get

$$\Delta_{t+1}^{\frac{1}{3}} - \frac{1}{2^{\frac{1}{3}-1}} \leq \frac{1}{2^{\frac{1}{3}}}(\Delta_t^{\frac{1}{3}} + 1) - \frac{1}{2^{\frac{1}{3}-1}} = \frac{1}{2^{\frac{1}{3}}} \left(\Delta_t^{\frac{1}{3}} - \frac{1}{2^{\frac{1}{3}-1}} \right).$$

Therefore,

$$\Delta_j^{\frac{1}{3}} - \frac{1}{2^{\frac{1}{3}} - 1} \leq \frac{1}{2^{j/3}} \left(\Delta_0^{\frac{1}{3}} - \frac{1}{2^{\frac{1}{3}} - 1} \right). \quad (39)$$

Since $\Delta_0 = \Delta$, $2^{\frac{1}{3}} - 1 > \frac{1}{4}$ and $\Delta/2^{j-1} > \frac{f\epsilon^{\frac{1}{j}}}{1+\delta}$ is large enough, (39) implies that

$$\Delta_j^{\frac{1}{3}} \leq \frac{\Delta^{\frac{1}{3}}}{2^{j/3}} + 4 \leq (1+\delta)^{\frac{1}{3}} \frac{\Delta^{\frac{1}{3}}}{2^{j/3}}.$$

Therefore, we have shown that $\Delta_j \leq (1+\delta)\frac{\Delta}{2^j}$. Note also that the same proof shows that for every $t \leq j$ we have $\Delta_t \leq (1+\delta)\frac{\Delta}{2^t}$.

We turn now to the claim regarding s_j . For all $t < j$, we have by definition

$$s_t \geq \frac{s_0}{4^t} = \frac{\Delta^2}{4^t f} = \frac{1}{(1+\delta)^2} \frac{\left(\frac{(1+\delta)\Delta}{2^t}\right)^2}{f} \geq \frac{1}{(1+\delta)^2} \left(\frac{(1+\delta)\Delta}{2^t}\right)^{2-\epsilon'} \geq \frac{1}{(1+\delta)^2} \Delta_t^{2-\epsilon'}, \quad (40)$$

where in the last inequality we used the fact that $\Delta_t \leq (1+\delta)\frac{\Delta}{2^t}$ for all $t \leq j$. Using (40) to bound Δ_t in the expression that defines s_t , we get

$$s_{t+1} \leq \frac{s_t}{4} + 2((1+\delta)^2 s_t)^{\frac{3}{2} \frac{1}{(2-2\epsilon')}}. \quad (41)$$

To bound the r.h.s. of (41) we recall that $\epsilon' < \frac{1}{10}$ implying $\frac{3}{2(2-2\epsilon')} < \frac{3}{2(2-1/5)} = 5/6$. Assuming that f (and hence also Δ_t) is large enough, we obtain

$$s_{t+1} \leq \frac{s_t}{4} + 3s_t^{5/6} \leq \frac{1}{4}(s_t^{1/6} + 2)^6 = \frac{1}{4}(s_t + 12s_t^{5/6} + 60s_t^{2/3} + \dots). \quad (42)$$

Taking 6-th roots and subtracting $\frac{5}{6^{1/6}-1}$ from both sides, we obtain

$$s_{t+1}^{1/6} - \frac{5}{6^{1/6}-1} \leq \frac{1}{4^{1/6}}(s_t^{1/6} + 2) - \frac{5}{6^{1/6}-1} \leq \frac{1}{4^{1/6}} \left(s_t^{1/6} - \frac{5}{6^{1/6}-1} \right).$$

Therefore,

$$s_j^{1/6} - \frac{5}{6^{1/6}-1} \leq \frac{1}{4^{j/6}} \left(s_0^{1/6} - \frac{5}{6^{1/6}-1} \right),$$

and, since $s_0 = \Delta^2/f$,

$$s_j^{1/6} \leq \frac{(\Delta^2)^{1/6}}{4^{j/6} f^{1/6}} + \frac{5}{6^{1/6}-1} \leq \left(\frac{\Delta^2}{4^j f} \right)^{1/6} + 15.$$

Since

$$\frac{\Delta^2}{4^j f} = \left(\frac{\Delta}{2^j} \right)^2 \frac{1}{f} = \left(\frac{\Delta}{2^{j-1}} \right)^2 \frac{1}{4f} > \left(\frac{f^{1/\epsilon'}}{1+\delta} \right)^2 \frac{1}{4f} = \frac{f^{\frac{2}{\epsilon'}-1}}{4(1+\delta)^2}$$

can be made arbitrarily large by taking f sufficiently large, we see that

$$\left(\frac{\Delta^2}{4^j f} \right)^{1/6} + 15 \leq (1+\delta)^{1/3} \left(\frac{\Delta^2}{4^j f} \right)^{1/6}.$$

Thus $s_j \leq ((1+\delta)\Delta/2^j)^2/f$, completing the proof.

F Proof of Proposition 2.1

We use the term “with high probability” to refer to probabilities that tend to 1 as n goes to infinity. Proposition 2.1 follows in a straightforward way from the following lemma.

Lemma F.1. *For any $\delta \in (0, 1)$ there exists a constant d_0 such that, for any $d \in \left(d_0 \ln n, (n \ln n)^{\frac{1}{3}}\right)$, each vertex of the random graph $G = G(n, d/n)$ is contained in at most Δ^δ triangles with high probability, where Δ is the maximum degree of G .*

Proof of Proposition 2.1. According to [5], for a graph $G \in G(n, d/n)$ we know that with high probability

$$\chi(G) = \frac{1}{2} \frac{d}{\ln d} (1 + o(1)). \quad (43)$$

Fix $\zeta \in (0, 1)$ and $\delta \in (0, \frac{2\zeta}{1+2\zeta})$. According to Lemma F.1, there exists a constant d_0 such that for any $d \in \left(d_0 \ln n, (n \ln n)^{\frac{1}{3}}\right)$ each vertex of $G = G(n, d/n)$ is contained in at most Δ^δ triangles with probability that tends to 1 as n goes to infinity. Thus, we can apply Theorem 2.5 with parameter $\zeta > 0$ since

$$f = \frac{\Delta^2}{\Delta^\delta} > \Delta^{2 - \frac{2\zeta}{1+2\zeta}} (\ln \Delta)^2,$$

for large enough Δ . This yields an upper bound q on the chromatic number of G that is at most

$$\begin{aligned} q &= (1 + \zeta) \frac{\Delta}{\ln \sqrt{f}} \\ &\leq (1 + \zeta) \frac{\Delta}{\frac{1+\zeta}{1+2\zeta} \ln \Delta + \ln \ln \Delta} \\ &\leq (1 + 2\zeta) \frac{\Delta}{\ln \Delta}. \end{aligned} \quad (44)$$

Moreover, since the expected degree of every vertex of G is d and its distribution is binomial with parameter $\frac{d}{n}$, standard Chernoff bounds and the union bound imply that for any $\eta \in (0, 1)$, $\Delta \leq (1 + \eta)d$ with high probability, for large enough d_0 .

Combining the latter fact with (43) and (44), we deduce that we can find an arbitrarily small constant $\eta' \in (0, 1)$ such that

$$q \leq (2 + \eta') \chi(G)$$

by choosing ζ and η sufficiently small. Picking $\eta' = \frac{4\epsilon}{1-2\epsilon}$ we obtain $\chi(G) \geq \frac{q}{2+\eta'} \geq q(\frac{1}{2} - \epsilon)$, concluding the proof of Proposition 2.1. \square

Finally, we go back and prove Lemma F.1.

Proof of Lemma F.1. Let Δ_v be the random variable that equals the degree of vertex v of G . Observe that $\Delta_v \sim \text{Binom}(n-1, \frac{d}{n})$ and, therefore, using a standard Chernoff bound and the fact that $d \geq d_0 \log n$ we get that

$$\Pr \left[\Delta_v \notin \left(1 \pm \frac{1}{10}\right)d \right] \leq \frac{1}{n^2},$$

for large enough d_0 . Thus, by a union bound we get that $\Pr[\Delta \in (1 \pm \frac{1}{10})d] \leq \frac{1}{n}$. Let T_v be the number of triangles that contain vertex v and B be the event that there exists a vertex u such that $\Delta_u \notin (1 \pm \frac{1}{10})d$. Then,

$$\Pr[T_v > \Delta^\delta] \leq \Pr[T_v > \Delta^\delta \mid \overline{B}] + \Pr[B] \leq \Pr[T_v > \Delta^\delta \mid \overline{B}] + \frac{1}{n}. \quad (45)$$

To upper bound the first term in the righthand side of (45), we fix an arbitrary $\Delta_0 \in (1 \pm \frac{1}{10})d$, and upper bound $\Pr[T_v > \Delta^\delta \mid \Delta_v = \Delta_0]$. Towards that end, we observe that

$$\mathbb{E}[T_v \mid \Delta_v = \Delta_0] = \binom{\Delta_0}{2} \frac{d}{n} \leq \frac{(1 + \frac{1}{10})^2 d^3}{2n}.$$

Thus, for any fixed values of $\Delta, \Delta_0 \in (1 \pm \frac{1}{10})d$, setting $1 + \beta = \frac{\Delta^\delta}{(1 + \frac{1}{10})^2 d^3 / 2n}$ and using a standard Chernoff bound we obtain:

$$\Pr[T_v > \Delta^\delta \mid \Delta_v = \Delta_0] \leq e^{-\frac{\beta^2 1.1^2 d^3 / 2n}{3}} \leq \frac{1}{n^2}$$

since

$$\begin{aligned} \beta &\geq \frac{((1 - \frac{1}{10})d)^\delta - 1.1^2 d^3 / 2n}{1.1^2 d^3 / 2n} > 0, \\ \frac{1.1^2}{3} \beta^2 \frac{d^3}{2n} &\geq \frac{1.1^2 \left(((1 - \frac{1}{10})d)^\delta - d^3 / 2n \right)^2}{3 d^3 / 2n} \geq 2 \ln n, \end{aligned}$$

whenever $d \in [d_0 \ln n, (n \ln n)^{\frac{1}{3}}]$ and for large enough n and d_0 . Since Δ_0 was chosen arbitrarily, we obtain that $\Pr[T_v > \Delta^\delta \mid \overline{B}] \leq \frac{1}{n^2}$. Taking a union bound over v concludes the proof of the lemma. \square