

# Commutative Algorithms Approximate the LLL-distribution

Fotis Iliopoulos \*

University of California Berkeley

fotis.iliopoulos@berkeley.edu

June 8, 2019

## Abstract

Following the groundbreaking Moser-Tardos algorithm for the Lovász Local Lemma (LLL), a series of works have exploited a key ingredient of the original analysis, the witness tree lemma, in order to: derive deterministic, parallel and distributed algorithms for the LLL, to estimate the entropy of the output distribution, to partially avoid bad events, to deal with super-polynomially many bad events, and even to devise new algorithmic frameworks. Meanwhile, a parallel line of work has established tools for analyzing stochastic local search algorithms motivated by the LLL that do not fall within the Moser-Tardos framework. Unfortunately, the aforementioned results do not transfer to these more general settings. Mainly, this is because the witness tree lemma, provably, does not longer hold. Here we prove that for commutative algorithms, a class recently introduced by Kolmogorov and which captures the vast majority of LLL applications, the witness tree lemma does hold. Armed with this fact, we extend the main result of Haeupler, Saha, and Srinivasan to commutative algorithms, establishing that the output of such algorithms well-approximates the *LLL-distribution*, i.e., the distribution obtained by conditioning on all bad events being avoided, and give several new applications. For example, we show that the recent algorithm of Molloy for list-coloring number of sparse, triangle-free graphs can output exponential many list-colorings of the input graph.

---

\*Research supported by NSF grant CCF-1514434 and the Onassis Foundation.

# 1 Introduction

Many problems in combinatorics and computer science can be phrased as finding an object that lacks certain bad properties, or “flaws”. In this paper we study algorithms that take as input a flawed object and try to remove all flaws by transforming the object through repeated probabilistic action.

Concretely, let  $\Omega$  be a set of objects and let  $F = \{f_1, f_2, \dots, f_m\}$  be a collection of subsets of  $\Omega$ . We will refer to each  $f_i \in F$  as a *flaw* to express that its elements share some negative feature. For example, if a CNF formula  $F$  on  $n$  variables has clauses  $c_1, c_2, \dots, c_m$ , we can define for each clause  $c_i$  the flaw (subcube)  $f_i \subseteq \{0, 1\}^n$  whose elements violate  $c_i$ . Following linguistic rather than mathematical convention we say that  $f$  is present in  $\sigma$  if  $f \ni \sigma$  and that  $\sigma \in \Omega$  is *flawless* (perfect) if no flaw is present in  $\sigma$ .

To prove the *existence* of flawless objects we can often use the Probabilistic Method. As a matter of fact, in many interesting cases, this is the only way we know how to do so. To employ the Probabilistic Method, we introduce a probability measure  $\mu$  over  $\Omega$  and consider the collection of “bad” events corresponding to flaws. If we are able to show that the probability to avoid all bad events is strictly positive, then this implies the existence of a flawless object. A trivial example is the case where all the bad events are independent of one another and none of them has probability one. One of the most powerful tools of the Probabilistic Method is the Lovász Local Lemma [17] which weakens the latter restrictive condition of independence to a condition of limited dependence.

Making the LLL constructive was the study of intensive research for over two decades [7, 4, 33, 16, 41]. The breakthrough was made by Moser [35] who gave a very simple algorithm that finds a satisfying assignment of a  $k$ -CNF formula, under conditions that nearly match the LLL condition for satisfiability. Very shortly afterwards, Moser and Tardos [36] made the general LLL constructive for any *product* probability measure over explicitly presented variables. Specifically, they proved that whenever the general LLL condition holds, the *Resample* algorithm, which repeatedly selects *any* occurring bad event and resamples all its variables according to the measure, i.e., independently, quickly converges to a flawless object.

The first result that made the LLL constructive in a non-product probability space was due to Harris and Srinivasan in [25], who considered the space of permutations endowed with the uniform measure. Subsequent works by Achlioptas and Iliopoulos [2, 1] introducing the flaws/actions framework, and of Harvey and Vondrák [27] introducing the resampling oracles framework, made the LLL constructive in more general settings. These frameworks [2, 1, 27] provide tools for analyzing *focused* stochastic search algorithms [38], i.e., algorithms which, like *Resample*, search by repeatedly selecting a flaw of the current state and moving to a random nearby state that avoids it, in the hope that, more often than not, more flaws are removed than introduced, so that a flawless object is eventually reached. At this point, all LLL applications we are aware of have efficient algorithms analyzable in these frameworks.

Besides conditions for existence and fast convergence to perfect objects, one could ask further questions regarding properties of focused search algorithms. For instance, “are they parallelizable?”, “how many solutions can they output?”, “what is the expected “weight” of a solution?”, etc. These questions and more have been answered for the Moser-Tardos algorithm in a long series of work [36, 21, 26, 28, 10, 13, 20, 23]. As a prominent example, the result of Haeupler, Saha and Srinivasan [21], as well as follow-up works of Harris and Srinivasan [26, 22], allow one to argue about the dynamics of the MT process, resulting in several new applications such as estimating the entropy of the output distribution, partially avoiding bad events, dealing with super-polynomially many bad events, and even new frameworks [24, 11].

Unfortunately, most of these follow-up results that further enhance, or exploit, our understanding of the MT process are not transferable to the general settings of [2, 27, 1]. Mainly, this is because a key and elegant technical result of the original analysis of Moser and Tardos, *the witness tree lemma*, does not longer hold under the most general assumptions [27]. Roughly, it states that any tree of bad events growing backwards in time from a certain root bad event  $A_i$ , with the children of each node  $A_j$  being bad events that are adjacent to  $A_j$  in the dependency graph, has probability of being consistent with the trajectory of

the algorithm that is bounded by the product of the probabilities of all events in this tree. The witness tree lemma and its variations [28, 20] has been used for several other purposes besides those already mentioned, such as designing deterministic, parallel and distributed algorithms for the LLL [36, 10, 13, 20, 23].

On the other hand, Harris and Srinivasan [25] do manage to prove the witness tree lemma for their algorithm for the LLL on the space of permutations, via an analysis that is tailored specifically to this setting. Although their proof does not seem to be easily generalizable to general spaces, their success makes it natural to ask if we can impose mild assumptions in the general settings of [2, 27, 1] under which the witness tree lemma (and most of its byproducts) can be established.

The main contribution of this paper is to answer this question positively by showing that it is possible to prove the witness tree lemma in the *commutative setting*. The latter was recently introduced by Kolmogorov [30], who showed that under its assumptions one can obtain parallel algorithms, as well as the flexibility of having arbitrary flaw choice strategy in the frameworks of [2, 27, 1]. We note that the commutative setting captures the vast majority of LLL applications, including but not limited to both the variable and the permutation settings.

Subsequently to the present work, Achlioptas, Iliopoulos and Sinclair [3] gave a simpler proof of the witness tree lemma under a more general notion of commutativity (essentially matrix commutativity) at the mild cost of slightly restricting the family of flaw choice strategies (as we will see, in this paper the flaw choice strategy can be arbitrary).

Armed with the witness tree lemma, we are able to study properties of algorithms in the commutative setting and give several applications.

**Distributional Properties.** As already mentioned, one of the most important applications of the witness tree lemma is given in the paper of Haeupler, Saha and Srinivasan [21], who study properties of the *MT-distribution*, the output distribution of the MT algorithm. Their main result is that the MT-distribution well-approximates the *LLL-distribution*, i.e., the distribution obtained by conditioning on all bad events being avoided. As an example, an immediate consequence of this fact is that one can argue about the expected *weight* of the output of the MT algorithm, given a weighting function over the space  $\Omega$ . Furthermore, as shown in the same paper [21] and follow-up papers by Harris and Srinivasan [26, 22], one can lower bound the entropy of the MT distribution, go beyond the LLL conditions (if one is willing to only partially avoid bad events), and deal with applications with super-polynomially many bad events.

Here we extend the result of [21] to the commutative setting: Given a commutative algorithm that is perfectly compatible with the underlying probability measure, its output well-approximates the LLL-distribution in the same sense the MT-distribution does in the variable setting. For arbitrary commutative algorithms, the quality of the approximation additionally depends on the compatibility of the algorithm with the measure on the event(s) of interest. A simplified, and imprecise, version of our main theorem, which assumes that the initial state of the algorithm is sampled according to the underlying probability distribution  $\mu$ , is as follows. The formal statement of our main theorem can be found in Section 3.

**Theorem 1.1** (Informal and Imprecise Statement). *If algorithm  $\mathcal{A}$  is commutative and the algorithmic LLL conditions hold then, for each  $E \subseteq \Omega$ ,*

$$\Pr[E] \leq \gamma(E) \left(1 + \frac{1}{d}\right)^{D_E},$$

where  $E$  is independent of all but at most  $D_E$  flaws,  $d$  is the maximum degree of the dependency graph,  $\gamma(E) \geq \mu(E)$  is a measure of the “compatibility” between  $\mathcal{A}$  and the underlying probability distribution  $\mu$  at  $E$ , and  $\Pr[E]$  is the probability that  $\mathcal{A}$  ever reaches  $E$  during its execution.

Moreover, we quantitatively improve the bounds of [21] under the weaker assumptions of *Shearer’s condition* [40], i.e., the most general LLL condition under the assumption that the dependency graph is

undirected. This allows us to study distributional properties of commutative algorithms using criteria that lie between the General LLL and Shearer’s condition such as the Clique LLL [29]. In Section 5 we discuss the byproducts of our main theorem that we will use in our applications. Finally, in Appendix A we discuss how one can deal with settings with super-polynomially many flaws.

**Algorithmic LLL Without a Slack and Arbitrary Flaw Choice Strategy.** The works of Achlioptas, Iliopoulos and Kolmogorov [2, 1, 30] require a multiplicative slack in the generalized LLL conditions in order to establish fast convergence to a perfect object. On the other hand, Harvey and Vondrák [27] dispense with this requirement in the important case of algorithms that are perfectly compatible with the underlying measure under the mild assumption that the dependency graph is undirected.

Using the witness lemma, we are able to dispense with the multiplicative slack requirement for arbitrary algorithms in the commutative setting and also have the flexibility of arbitrary flaw choice strategy, as in the result of Kolmogorov [30].

**Improved Running Time Bounds.** We are able to improve the running time bounds of Harvey and Vondrák [27] for commutative algorithms, matching those of Kolipaka and Szegedy [28] for the MT algorithm. Whether this could be done was left as an open question in [27]. We note that while the results of Achlioptas, Iliopoulos and Kolmogorov [2, 1, 30] also manage to give improved running time bounds they require a multiplicative-slack in the LLL conditions.

**Concrete Applications.** In Section 6 we give concrete applications of commutative algorithms showing new results for the problems of *rainbow matchings*, *list-coloring* and *acyclic edge coloring*. Each application is chosen so that it demonstrates specific features of our results.

The first application is in the space of matchings of a complete graph. We use this problem as an example that allows us to show how several byproducts of approximating the LLL-distribution can be applied in a black-box manner to a setting that is not captured either by the variable or the permutation setting, and for which we know [2, 27, 30] how to design commutative algorithms that are perfectly compatible with the uniform measure over the state space.

The second, and perhaps most interesting, application is to show that the algorithm of Molloy [32] for finding proper colorings in triangle-free graphs with maximum degree  $\Delta$  using  $(1 + \epsilon) \frac{\Delta}{\ln \Delta}$  colors, can actually output exponentially many such colorings with positive probability. First, we show that Molloy’s algorithm can be analyzed in the general frameworks of the algorithmic LLL and that it is commutative, a fact that gives us access to properties of its output distribution. Then, we apply results regarding the entropy of the output of commutative algorithms. We show the following theorem.

**Theorem 1.2.** *For every  $\epsilon > 0$  there exists  $\Delta_\epsilon$  such that every triangle-free graph  $G$  with maximum degree  $\Delta \geq \Delta_\epsilon$  has list-chromatic number  $\chi_\ell(G) \leq (1 + \epsilon) \frac{\Delta}{\ln \Delta}$ . Furthermore, if  $G$  is a graph on  $n$  vertices then, for every  $\eta > 0$ , there exists an algorithm  $\mathcal{A}$  that constructs such a coloring in polynomial time with probability at least  $1 - \frac{1}{n^\eta}$ . In addition,  $\mathcal{A}$  is able to output  $e^{cn}$  distinct list-colorings with positive probability, where  $c > 0$  is a constant that depends on  $\epsilon$  and  $\Delta$ .*

We emphasize that the algorithm of Molloy is a sophisticated stochastic local search algorithm whose analysis is far from any standard LLL setting. The fact that our results allow us to state non-trivial facts about its distributional properties almost in a black-box fashion is testament to their flexibility.

In the third application we show how one can use bounds on the output distribution of commutative algorithms that are induced by the Shearer’s condition in order to analyze applications of the Clique version of the Local Lemma in the problem of acyclic edge coloring of a graph.

## 2 Background and Preliminaries

In this section we present the necessary background and definitions to describe our setting. In Subsection 2.1 we describe the Lovász Local Lemma. In Subsections 2.2 and 2.3 we formally outline the algorithmic assumptions of [2, 27, 1, 30]. In Subsection 2.4 we describe improved Lovász Local Lemma criteria formulated in our setting.

### 2.1 The Lovász Local Lemma

To prove the *existence* of flawless objects we can often use the Probabilistic Method. To do so, we introduce a probability measure  $\mu$  over  $\Omega$  and consider the collection of “bad” events corresponding to flaws. If we are able to show that the probability to avoid all bad events is strictly positive, then this implies the existence of a flawless object. One of the most powerful tools to establish the latter is the Lovász Local Lemma [17].

**General LLL.** *Let  $(\Omega, \mu)$  be a probability space and  $\mathcal{A} = \{A_1, A_2, \dots, A_m\}$  be a set of  $m$  (bad) events. For each  $i \in [m]$ , let  $D(i) \subseteq [m] \setminus \{i\}$  be such that  $\mu(A_i \mid \bigcap_{j \in S} \overline{A_j}) = \mu(A_i)$  for every  $S \subseteq [m] \setminus (D(i) \cup \{i\})$ . If there exist positive real numbers  $\{\psi_i\}_{i=1}^m$  such that for all  $i \in [m]$ ,*

$$\frac{\mu(A_i)}{\psi_i} \sum_{S \subseteq D(i) \cup \{i\}} \prod_{j \in S} \psi_j \leq 1, \quad (1)$$

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

**Remark 2.1.** *Condition (1) above is equivalent to the more well-known form  $\mu(A_i) \leq x_i \prod_{j \in D(i)} (1 - x_j)$ , where  $x_i = \psi_i / (1 + \psi_i)$ . As we will see, formulation (1) facilitates refinements.*

Let  $G$  be the digraph over the vertex set  $[m]$  with an edge from each  $i \in [m]$  to each element of  $D(i) \cup \{i\}$ . We call such a graph a *dependency graph*. Therefore, at a high level, the LLL states that if there exists a sparse dependency graph and each bad event is not too likely, then perfect objects exist.

### 2.2 Algorithmic Framework

Here we describe the class of algorithms we will consider as well as the algorithmic LLL criteria for fast convergence to a perfect object. Since we will be interested in algorithms that search for perfect objects, we sometimes refer to  $\Omega$  as a state space and to its elements as states.

For a state  $\sigma$ , we denote by  $U(\sigma) = \{j \in [m] \text{ s.t. } f_j \ni \sigma\}$  the set of indices of flaws that are present at  $\sigma$ . We consider algorithms which at each flawed state  $\sigma$  choose an element of  $U(\sigma)$  and randomly move to a nearby state in an effort to *address* the corresponding flaw. Concretely, we will assume that for every flaw  $f_i$  and every state  $\sigma \in f_i$  there is a probability distribution  $\rho_i(\sigma, \cdot)$  with a non-empty support  $A(i, \sigma) \subseteq \Omega$  such that addressing flaw  $f_i$  at state  $\sigma$  amounts to selecting the next state  $\sigma'$  from  $A(i, \sigma)$  with probability  $\rho_i(\sigma, \sigma')$ . We call  $A(i, \sigma)$  the set of *actions* for addressing flaw  $f_i$  at  $\sigma$  and note that potentially  $A(i, \sigma) \cap f_i \neq \emptyset$ , i.e., addressing a flaw does not necessarily imply removing it. The actions for flaw  $f_i$  form a digraph  $D_i$  on  $\Omega$  having an arc  $\sigma \xrightarrow{i} \sigma'$  for each pair  $(\sigma, \sigma') \in f_i \times A(i, \sigma)$ . Let  $D$  be the multi-digraph on  $\Omega$  that is the union of all  $D_i$ .

We consider algorithms that start from a state  $\sigma \in \Omega$  picked from an initial distribution  $\theta$ , and then repeatedly pick a flaw that is present in the current state and address it. The algorithm always terminates when it encounters a flawless state.

To state the algorithmic LLL criteria for fast convergence of such algorithms we need to introduce two key ingredients. The first one is a notion of *causality* among flaws that will be used to induce a graph

over  $[m]$ , which will play a role similar to the one of the dependency graph in the existential Local Lemma formulation. We note that there is a formal connection between causality graphs and dependency graphs (for more details see [27]).

**Causality.** For an arc  $\sigma \xrightarrow{i} \sigma'$  in  $D_i$  and a flaw  $f_j$  present in  $\sigma'$  we say that  $f_i$  causes  $f_j$  if  $f_i = f_j$  or  $f_j \not\exists \sigma$ . If  $D_i$  contains any arc in which  $f_i$  causes  $f_j$  we say that  $f_i$  potentially causes  $f_j$ .

**Causality Digraph.** Any digraph  $C = C(\Omega, F, D)$  on  $[m]$  where  $i \rightarrow j$  exists whenever  $f_i$  potentially causes  $f_j$  is called a causality digraph. The neighborhood of a flaw  $f_i$  in  $C$  is  $\Gamma(i) = \{j : i \rightarrow j \text{ exists in } C\}$ .

The second ingredient is a measure of *compatibility* between the actions of the algorithm for addressing each flaw  $f_i$  (that is, digraph  $D_i$ ) and the probability measure  $\mu$  over  $\Omega$  which we will use for the analysis. As was shown in [27, 1, 30] one can capture compatibility by letting

$$d_i = \max_{\sigma \in \Omega} \frac{\nu_i(\sigma)}{\mu(\sigma)} \geq 1, \quad (2)$$

where  $\nu_i(\sigma)$  is the probability of ending up at state  $\sigma$  at the end of the following experiment: sample  $\omega \in f_i$  according to  $\mu$  and address flaw  $f_i$  at  $\omega$ . An algorithm achieving perfect compatibility for flaw  $f_i$ , i.e.,  $d_i = 1$ , is a *resampling oracle* for flaw  $f_i$  (observe that the Moser-Tardos algorithm is trivially a resampling oracle for every flaw). More generally, ascribing to each flaw  $f_i$  the *charge*

$$\gamma(f_i) = d_i \cdot \mu(f_i) = \max_{\sigma' \in \Omega} \frac{1}{\mu(\sigma')} \sum_{\sigma \in f_i} \mu(\sigma) \rho_i(\sigma, \sigma'),$$

yields the following algorithmization condition. If for every flaw  $f_i \in F$ ,

$$\frac{\gamma(f_i)}{\psi_i} \sum_{S \subseteq \Gamma(i)} \prod_{j \in S} \psi_j < 1 \quad (3)$$

then there exists a flaw choice strategy under which the algorithm will reach a perfect object fast. (In most applications, that is in  $O(\log |\Omega| + m \max_{i \in [m]} \log_2(1 + \psi_i))$  steps with high probability.)

Throughout the paper we assume that we are given an undirected causality graph  $C$  (and thus the relation  $\Gamma(\cdot)$  is symmetric) and we will sometimes write  $i \sim j$  if  $j \in \Gamma(i) \leftrightarrow j \in \Gamma(j)$ . Furthermore, for a set  $S \subseteq [m]$  we define  $\Gamma(S) = \bigcup_{i \in S} \Gamma(i)$ . Finally, we denote by  $\text{Ind}(S) = \text{Ind}_C(S)$  the set of independent subsets of  $S$  with respect to  $C$ .

### 2.3 Commutativity

We will say that  $\sigma \xrightarrow{i} \sigma'$  is a *valid trajectory* if it is possible to get from state  $\sigma$  to state  $\sigma'$  by addressing flaw  $f_i$  as described in the algorithm, i.e., if two conditions hold:  $i \in U(\sigma)$  and  $\sigma' \in A(i, \sigma)$ . Kolmogorov [30] described the following *commutativity* condition. We call the setting in which Definition 2.1 holds *the commutative setting*.

**Definition 2.1** (Commutativity [30]). *A tuple  $(F, \sim, \rho)$  is called commutative if there exists a mapping  $\text{Swap}$  that sends any trajectory  $\Sigma = \sigma_1 \xrightarrow{i} \sigma_2 \xrightarrow{j} \sigma_3$  with  $i \approx j$  to another valid trajectory  $\text{Swap}(\Sigma) = \sigma_1 \xrightarrow{j} \sigma'_2 \xrightarrow{i} \sigma_3$ , and:*

1. *Swap is injective,*
2.  $\rho_i(\sigma_1, \sigma_2) \rho_j(\sigma_2, \sigma_3) = \rho_j(\sigma_1, \sigma'_2) \rho_i(\sigma'_2, \sigma_3)$  .

It is straightforward to check that the Moser Tardos algorithm satisfies the commutativity condition. Furthermore, Kolmogorov showed that the same is true for resampling oracles in the permutation [25] and perfect matchings [27] settings, and Harris [23] designed commutative resampling oracles for hamiltonian cycles.

Finally, as already mentioned, Kolmogorov showed that in the commutativity setting one may choose an arbitrary flaw choice strategy which is a function of the entire past execution history. The same will be true for our results, so we make the convention that given a tuple  $(F, \sim, \rho)$  we always fix some arbitrary flaw choice strategy to get a well-defined, commutative algorithm  $\mathcal{A} = (F, \sim, \rho)$ .

## 2.4 Improved LLL Criteria

Besides the general form of the LLL (1) there exist improved criteria that apply in the full generality of the LLL setting. The most well-known are the *cluster expansion condition* [9] and the *Shearer's condition* [40]. Both of these criteria apply when the dependency graph is undirected and have been made constructive [28, 39, 2, 27, 1, 30] in the most general algorithmic LLL settings.

**Cluster Expansion Condition.** The cluster expansion condition strictly improves upon the General LLL condition (1) by taking advantage of the local density of the dependency graph.

**Definition 2.2.** Given a sequence of positive real numbers  $\{\psi_i\}_{i=1}^m$ , we say that the cluster expansion condition is satisfied if for each  $i \in [m]$ :

$$\frac{\gamma(f_i)}{\psi_i} \sum_{S \in \text{Ind}(\Gamma(i))} \prod_{j \in S} \psi_j \leq 1 . \quad (4)$$

**Shearer's Condition.** Let  $\gamma \in \mathbb{R}^m$  be the real vector such that  $\gamma_i = \gamma(f_i)$ . Furthermore, for  $S \subseteq [m]$  define  $\gamma_S = \prod_{j \in S} \gamma_j$  and the polynomial  $q_S$ :

$$q_S = q_S(\gamma) = \sum_{\substack{I \in \text{Ind}([m]) \\ S \subseteq I}} (-1)^{|I|-|S|} \gamma_I .$$

**Definition 2.3.** We say that the Shearer's condition is satisfied if  $q_S(\gamma) \geq 0$  for all  $S \subseteq [m]$ , and  $q_\emptyset(\gamma) > 0$ .

## 3 Statement of Results

Assuming that the LLL conditions (1) hold, the *LLL-distribution*, which we denote by  $\mu_{\text{LLL}}$ , is defined as the distribution induced by the measure  $\mu$  conditional on no bad event occurring. The following proposition relates the LLL-distribution to measure  $\mu$  making it a powerful tool that can be used to argue about properties of flawless objects. The idea is that if an (not necessarily bad) event  $E$  is independent from most bad events, then its probability under the LLL-distribution is not much larger than its probability under the probability measure  $\mu$ .

**Proposition 3.1** ([21]). *If the LLL conditions (1) hold, then for any event  $E$ :*

$$\mu_{\text{LLL}}(E) \leq \mu(E) \sum_{S \subseteq D(E)} \prod_{j \in S} \psi_j , \quad (5)$$

where  $D(E) \subseteq [m]$  is such that  $\mu(E \mid \bigcap_{j \in S} \bar{A}_j) = \mu(E)$  for all  $S \subseteq [m] \setminus D(E)$ .

The main result of Haeupler, Saha and Srinivasan [21] is that the Moser-Tardos algorithm approximates well the LLL-distribution, in the sense that the left-hanside of (5) bounds the probability that it ever reaches a subspace  $E \subseteq \Omega$  during its execution. Building on this fact, [21] and followup works [26, 22] manage to show several new applications.

Here we extend the latter result to arbitrary commutative algorithms. Given an arbitrary set  $E \subseteq \Omega$  and a commutative algorithm  $\mathcal{A}$ , consider an extension,  $\mathcal{A}_E$ , of  $\mathcal{A}$  by defining an extra flaw  $f_{m+1} \equiv E$  with its own set of probability distributions  $\rho_{m+1}(\sigma, \cdot)$ ,  $\sigma \in E$ . If  $\mathcal{A}$  is commutative with respect to  $\sim$ , we will say that  $\mathcal{A}_E$  is a *commutative extension* of  $\mathcal{A}$  if  $\mathcal{A}_E = (F \cup \{m+1\}, \sim, \rho)$  is also commutative.

Commutative extensions should be interpreted as a tool to bound the probability that  $\mathcal{A}$  ever reaches a subset  $E$  of the state space. That is, they are defined only for the purposes of the analysis and, typically in applications, they are a natural extension of the algorithm. For example, in the case of the Moser-Tardos algorithm applied to  $k$ -SAT, if one would like to bound the probability that the algorithm ever reaches a state such that variables  $x_1, x_2$  of the formula are both set to true, then one could define  $f_{m+1} = \{\sigma \in \Omega \text{ s.t. } \sigma(x_1) = \sigma(x_2) = 1\}$  along with the corresponding commutative extension of the Moser-Tardos algorithm that addresses  $f_{m+1}$  by resampling variables  $x_1, x_2$  according to the product measure over the variables of the formula that the Moser-Tardos algorithm uses whenever it needs to resample a violated clause. Indeed, commutative extensions of this form are implicitly defined in the analysis of [21] for the Moser-Tardos algorithm.

We will use the notation  $\Pr[\cdot] = \Pr_{\mathcal{A}}[\cdot]$  to refer to the probability of events in the probability space induced by the execution of algorithm  $\mathcal{A}$ . For example, the probability that  $\mathcal{A}$  ever reaches a set  $E \subseteq \Omega$  of the state space during its execution will be denoted by  $\Pr[E]$ .

**Theorem 3.2.** *If  $\mathcal{A} = (F, \sim, \rho)$  is commutative and the cluster expansion condition is satisfied then:*

1. for each  $i \in [m]$ :  $\mathbb{E}[N_i] \leq \lambda_{\text{init}} \psi_i$  ;
2. for each  $E \subseteq \Omega$ :  $\Pr[E] \leq \lambda_{\text{init}} \gamma(E) \sum_{S \in \text{Ind}(\Gamma(E))} \prod_{j \in S} \psi_j$  ;

where  $N_i$  is the number of times flaw  $f_i$  is addressed during the execution of  $\mathcal{A}$ ,  $\lambda_{\text{init}} = \max_{\sigma \in \Omega} \frac{\theta(\sigma)}{\mu(\sigma)}$ , and  $\Gamma(E)$  and  $\gamma(E)$  are defined with respect to a fixed commutative extension  $\mathcal{A}_E$ .

**Corollary 3.3.** *Algorithm  $\mathcal{A}$  terminates after  $O(\lambda_{\text{init}} \sum_{i \in [m]} \psi_i)$  steps in expectation.*

**Remark 3.1.** *If the Shearer’s condition is satisfied, then one can replace  $\psi_i$  in Theorem 3.2 with  $\frac{q_{(i)}(\gamma)}{q_{\emptyset}(\gamma)}$ .*

We note that the first part of Theorem 3.2 allows us to guarantee fast convergence of  $\mathcal{A}$  to a perfect object without having to assume a “slack” in the cluster expansion and Shearer’s conditions (unlike the works of [1, 30]) and, moreover, improves upon the (roughly quadratically worse) running bound of [27], matching the one of [28]. Whether the latter could be done was left as an open question in [27].

## 4 Proof of Main Results

In this section we state and prove the witness tree lemma for our setting. We then use it to prove Theorem 3.2.

### 4.1 The Witness Tree Lemma

Given a trajectory  $\Sigma = \sigma_1 \xrightarrow{w_1} \dots \sigma_t \xrightarrow{w_t} \sigma_{t+1}$  we denote by  $W(\Sigma) = (w_1, \dots, w_t)$  the *witness sequence* of  $\Sigma$ . (Recall that according to our notation,  $w_i$  denotes the index of the flaw that was addressed at the  $i$ -th step).

To state the witness tree lemma, we will first need to recall the definition of witness trees from [36], slightly reformulating to fit our setting. A witness tree  $\tau = (T, \ell_T)$  is a finite rooted, unordered, tree  $T$  along with a labelling  $\ell_T : V(T) \rightarrow [m]$  of its vertices with indices of flaws such that the children of a vertex  $v \in V(T)$  receives labels from  $\Gamma(\ell(v))$ . To lighten the notation, we will sometimes write  $(v)$  to denote  $\ell(v)$  and  $V(\tau)$  instead of  $V(T)$ . Given a witness sequence  $W = (w_1, w_2, \dots, w_t)$  we associate with each  $i \in [t]$  a witness tree  $\tau_W(i)$  that is constructed as follows: Let  $\tau_W^i(i)$  be an isolated vertex labelled by  $w_i$ . Then, going backwards for each  $j = i - 1, i - 2, \dots, 1$ : if there is a vertex  $v \in \tau_W^{j+1}(i)$  such that  $(v) \sim w_j$  then we choose among those vertices the one having the maximum distance (breaking ties arbitrarily) from the root and attach a new child vertex  $u$  to  $v$  that we label  $w_j$  to get  $\tau_W^j(i)$ . If there is no such vertex  $v$  then  $\tau_W^{j+1}(i) = \tau_W^j(i)$ . Finally, let  $\tau_W(i) = \tau_W^1(i)$ .

We will say that a witness tree  $\tau$  occurs in a trajectory  $\Sigma$  if  $W(\Sigma) = (w_1, w_2, \dots, w_t)$  and there is  $k \in [t]$  such that  $\tau_W(k) = \tau$ .

**Theorem 4.1** (The witness tree lemma). *Assume that  $\mathcal{A} = (F, \sim, \rho)$  is commutative. Then, for every witness tree  $\tau$  we have that:*

$$\Pr[\tau] \leq \lambda_{\text{init}} \prod_{v \in V(\tau)} \gamma(f(v)) .$$

We show the proof of Theorem 4.1 in Section 4.4.

## 4.2 Witness Trees and Stable Witness Sequences

Here we prove some properties of witness trees (which are induced by witness sequences of the algorithm) that will be useful to us later. We also draw a connection between witness trees and *stable witness sequences*, which we will need in the proof of Theorem 3.2. Stable witness sequences were first introduced in [28] to make the Shearer's condition constructive in the variable setting.

### 4.2.1 Properties of Witness Trees

The following propositions capture the main properties of witness trees we will need.

**Proposition 4.2.** *For a witness tree  $\tau = (T, \ell_T)$  let  $L_i = L_i(\tau)$  denote the set of labels of the nodes at distance  $i$  from the root. For each  $i \geq 0$ ,  $L_i \in \text{Ind}([m])$ .*

*Proof.* We will show that for each  $i \geq 0$ , and each  $\alpha, \beta \in L_i$  we have that  $\alpha \approx \beta$ .

Let  $W = (w_1, w_2, \dots, w_t)$  be a witness sequence that can occur in an execution of our algorithm. Let  $\alpha, \beta$  be two distinct elements of  $L_i$ . By the definition of  $\tau$ , labels  $\alpha, \beta$  correspond to two indices  $w_{j_1}, w_{j_2}$  of  $W$ . Assume without loss of generality that  $j_1 < j_2$ . Then, according to the algorithm for constructing  $\tau$ , index  $w_{j_2}$  is ‘‘attached first’’ to the  $i$ -th level of  $\tau$ . The proof is concluded by noticing that if  $w_{j_1} = \alpha \sim \beta = w_{j_2}$  then the node corresponding to  $w_{j_1}$  is eligible to be a child of the node corresponding to  $w_{j_2}$  during the construction of  $\tau$  and, thus,  $\beta \notin L_i$ , which is a contradiction.  $\square$

**Proposition 4.3.** *For a witness sequence  $W$  of length  $t$  and any two distinct  $i, j \in [t]$  we have that  $\tau_W(i) \neq \tau_W(j)$ .*

*Proof.* Let  $W = (w_1, w_2, \dots, w_t)$ . Assume w.l.o.g. that  $i < j$ . If  $w_i \neq w_j$  then the claim is straightforward because the root of  $\tau_W(i)$  is  $w_i$  while the root of  $\tau_W(j)$  is  $w_j$ . If  $w_i = w_j = w$ , then there two cases. In the first case,  $w \in \Gamma(w)$ , and so tree  $\tau_W(j)$  has at least one more vertex than  $\tau_W(i)$ . In the second

case  $w \notin \Gamma(w)$ . This implies that at the  $i$ -th step of any trajectory  $\Sigma$  such that  $W(\Sigma) = W$ , flaw  $f_w$  was addressed and removed. However, the fact that  $w_j = w$  implies that there has to be  $k \in \{i + 1, \dots, j - 1\}$  such that addressing  $w_k$  introduced  $w$  and thus,  $w_k \sim w$ . Again, this means that  $\tau_W(j)$  has at least one more vertex than  $\tau_W(i)$ . □

#### 4.2.2 Stable Witness Sequences

We will now recall the definition of *stable sequences* [28], which have been used in [27, 30] to make the Shearer’s condition constructive.

**Definition 4.4.** *A sequence of subsets  $(I_1, \dots, I_k)$  of  $[m]$  with  $k \geq 1$  is called stable if*

1.  $I_r \in \text{Ind}([m]) \setminus \{\emptyset\}$  for each  $r \in [k]$ ;
2.  $I_{r+1} \subseteq \Gamma(I_r)$  for each  $r \in [k - 1]$ .

**Definition 4.5.** *A witness sequence  $W = (w_1, \dots, w_t)$  is called stable if it can be partitioned into non-empty sequences as  $W = (W_1, \dots, W_k)$  such that the elements of each sequence  $W_r$  are distinct, and the sequence  $\phi_W := (I_1, \dots, I_k)$  is stable, where  $I_r$  is the set of indices of flaws in  $W_r$  (for  $r \in [k]$ ).*

*For any arbitrary ordering  $\pi$  among indices of flaws, if in addition each sequence  $W_r = (w_i, \dots, w_j)$  satisfies  $w_i \prec_\pi \dots \prec_\pi w_j$  then  $W$  is called  $\pi$ -stable.*

**Proposition 4.6.** ([30]) *For a stable witness sequence the partitioning in Definition 4.5 is unique.*

*Proof Sketch.* Let  $W = (w_1, \dots, w_t)$  be a stable sequence and consider the following algorithm. We start with a single segment containing  $w_1$ . For  $i = 2$  to  $t$ , if there exists index  $w_k$  in the currently last segment such that  $w_k \sim w_i$  then we start a new segment containing  $w_i$ . Otherwise, we add  $w_i$  to the currently last segment. □

For a witness sequence  $W = (w_1, \dots, w_t)$  let  $\text{Rev}[W] = (w_t, \dots, w_1)$  denote the reverse sequence. Let also  $R_W$  denote the first set (the “root”) of the stable sequence  $\phi_W := (I_1, \dots, I_k)$ , i.e.,  $R_W = I_1$ . Finally, let  $\mathcal{R}_i^\pi$  be the set of witness sequences  $W$  such that  $\text{Rev}[W]$  is  $\pi$ -stable and  $R_{\text{Rev}[W]} = \{i\}$ .

There is a connection between stable sequences and witness trees that we will need for the proof of Theorem 3.2 and which we will describe below.

Let  $\mathcal{W}_i$  denote the set of witness trees with root labelled by  $i$ . For each  $\tau \in \mathcal{W}_i$ , let  $\chi_\pi(\tau)$  be the *ordered* witness tree that is induced by ordering the children of each node in  $\tau$  from left to right, increasingly according to  $\pi$ . Define  $\mathcal{W}_i^\pi := \chi_\pi(\mathcal{W}_i)$  and observe that  $\chi_\pi$  is a bijection. Finally, recall that for a witness tree  $\tau$  we denote by  $L_j(\tau)$  the set of labels of the nodes at distance  $j$  from the root.

**Lemma 4.7.** *There is a bijection  $\chi_i^\pi$  mapping  $\mathcal{R}_i^\pi$  to  $\mathcal{W}_i^\pi$  with the following property: Fix  $W \in \mathcal{R}_i^\pi$  and let  $(I_1 = \{i\}, I_2, \dots, I_k)$  be the unique partitioning of  $\text{Rev}[W]$  guaranteed by Proposition 4.6. Then,  $I_j = L_{j-1}(\chi_i^\pi(W))$  for each  $j \in [k]$ .*

*Proof.* Consider a witness sequence  $W \in \mathcal{R}_i^\pi$  of length  $t$ . We define  $\chi_i^\pi(W) := \chi_\pi(\tau_W(t))$ . That is, we map  $W$  to the  $\pi$ -ordered witness tree that is induced by applying the procedure that constructs witness trees to the final element of the sequence. Recall now the procedure in the proof of Proposition 4.6. The key observation is that the application of this procedure to  $\text{Rev}[W]$  is “equivalent” to the procedure that constructs  $\tau_W(t)$ , in the sense that the decisions taken for partitioning  $\text{Rev}[W]$  to segments by the procedure of Proposition 4.6 are identical to the decisions taken by the procedure that constructs  $\tau_W(t)$  in order to form

$L_j(\tau_W(t))$ ,  $j \geq 0$ . In particular, if  $\phi_{\text{Rev}[W]} = (I_1 = \{i\}, I_2, \dots, I_k)$  is the unique partitioning of  $\text{Rev}[W]$ , then  $I_j = L_{j-1}(\tau_W(t)) = L_{j-1}(\chi_i^\pi(W))$  for each  $j \in [k]$ .

It remains to show that  $\chi_i^\pi$  is bijection. To see this, at first observe that from  $\chi_i^\pi(W)$  one can uniquely reconstruct  $\phi_{\text{Rev}[W]}$ . Given  $\phi_{\text{Rev}[W]}$  one can reconstruct  $\text{Rev}[W]$  (and, thus,  $W$ ) by ordering each segment of  $\phi_{\text{Rev}[W]}$  according to  $\pi$ . □

### 4.2.3 Counting Witness Trees

In our proofs we will need to bound the sum over all trees  $\tau \in \mathcal{W}_i$  of the product of charges of the labels of the nodes of each tree  $\tau$ . Fortunately, the method for doing that is well trodden by now (see for example [36, 39]). Here we show the following lemma whose proof can be found in Appendix C. Recall that  $\mathcal{W}_i$  denotes the set of all possible witness trees with root that is labelled by  $i$ .

**Lemma 4.8.** *If the cluster expansion condition is satisfied then:*

$$\sum_{\tau \in \mathcal{W}_i} \prod_{v \in V(\tau)} \gamma(f(v)) \leq \psi_i .$$

We also show the following lemma that can be used whenever the Shearer's condition applies.

**Lemma 4.9.** *If the Shearer's condition is satisfied then:*

$$\sum_{\tau \in \mathcal{W}_i} \prod_{v \in V(\tau)} \gamma(f(v)) \leq \frac{q_{\{i\}}(\gamma)}{q_\emptyset(\gamma)} .$$

*Proof.* We first observe that due to Lemma 4.7 we have that

$$\sum_{\tau \in \mathcal{W}_i} \prod_{v \in V(\tau)} \gamma(f(v)) = \sum_{\tau \in \mathcal{W}_i^\pi} \prod_{v \in V(\tau)} \gamma(f(v)) = \sum_{W \in \mathcal{R}_i^\pi} \prod_{w \in W} \gamma(f_w) .$$

Now let  $\text{Stab}_i$  denote the set of stable set sequences whose first segment is  $\{i\}$  and also every segment is non-empty. For  $\phi = (I_1, I_2, \dots, I_k) \in \text{Stab}_i$  define  $\gamma_\phi = \prod_{i=1}^k \prod_{i \in I} \gamma(f_i)$ . Observe that there is a natural injection from  $\mathcal{R}_i^\pi$  to  $\text{Stab}_i$  which maps each sequence  $W \in \mathcal{R}_i^\pi$  to  $\phi_{\text{Rev}[W]}$ . This is because given  $\phi_{\text{Rev}[W]}$  one can reconstruct  $\text{Rev}[W]$  (and, thus,  $W$ ) by ordering each segment of  $\phi_{\text{Rev}[W]}$  according to  $\pi$ . The latter observation implies that:

$$\sum_{\tau \in \mathcal{W}_i} \prod_{v \in V(\tau)} \gamma(f(v)) \leq \sum_{\phi \in \text{Stab}_i} \gamma_\phi = \frac{q_{\{i\}}(\gamma)}{q_\emptyset(\gamma)} ,$$

where the proof of the last inequality can be found in Theorem 14 of [28] and in Lemmas 5.26, 5.27 of [27]. □

**Remark 4.1.** *We note that if we have assumed a stronger “cluster expansion condition” (namely, in (4) we have  $\Gamma(i) \cup \{i\}$ ) instead of  $\Gamma(i)$ , then Corollary 4.8 could have also been shown as an immediate application of Lemma 4.9, since it is known ([9, 27, 30]) that, in this case, for every  $i \in [m]$  we have that  $\frac{q_{\{i\}}(\gamma)}{q_\emptyset(\gamma)} \leq \psi_i$ .*

### 4.3 Proof of Theorem 3.2

We first prove Theorem 3.2. The first part follows by observing that if  $W$  is the witness sequence corresponding to the trajectory of the algorithm, then  $N_i$  is the number of occurrences of flaw  $f_i$  in  $W$ , and according to Proposition 4.3, also the number of distinct witness trees occurring that have their root labeled  $i$ . Therefore, one can bound the expectation of  $T_i$  by summing the bounds in Lemma 4.1. In particular, Lemma 4.8 concludes the proof of the first part.

To see the second part of Theorem 3.2, consider the new set of flaws  $F' = F \cup \{f_{m+1}\}$ , where  $f_{m+1} = E$ , as well as a “truncated” commutative extension  $\mathcal{A}'$  of  $\mathcal{A}$  with the following properties:

- (i) For each state  $\sigma \notin f_{m+1}$  algorithm  $\mathcal{A}'$  invokes  $\mathcal{A}$  to choose its next state;
- (ii)  $\gamma(E) := \gamma_{\mathcal{A}'}(f_{m+1})$ ;
- (iii)  $f_{m+1}$  is always of the highest priority: when at a state  $\sigma \in f_{m+1}$ ,  $\mathcal{A}'$  chooses to address  $f_{m+1}$ ;
- (iv)  $\mathcal{A}'$  stops after it addresses  $f_{m+1}$  for the first time.

By coupling  $\mathcal{A}$  and  $\mathcal{A}'$  we see that  $\Pr_{\mathcal{A}}[E] = \Pr_{\mathcal{A}'}[f_{m+1}]$ . Let  $\mathcal{W}_E$  be the set of witness trees that can occur in an execution of  $\mathcal{A}'$  and whose root is labelled by  $m + 1$ . Notice that, due to property (iv) of  $\mathcal{A}'$ , every tree  $\tau \in \mathcal{W}_E$  contains exactly one node (the root) labelled by  $m + 1$ , while every other node is labelled by elements in  $[m]$ . Furthermore, the set of labels of the children of the root of  $\tau$  is an element of  $\text{Ind}(\Gamma(E))$ . Finally, if  $v$  is a node that corresponds to a child of the root in  $\tau$ , then the subtree  $\tau_v$  that is rooted at  $v$  is an element of  $\mathcal{W}_{(v)}$ . Using Theorem 4.1 and the fact that  $\mathcal{A}'$  is commutative we get:

$$\Pr_{\mathcal{A}}[E] \leq \sum_{\tau \in \mathcal{W}_E} \Pr_{\mathcal{A}'}[\tau] \leq \lambda_{\text{init}} \gamma(E) \sum_{S \in \text{Ind}(\Gamma(E))} \left( \prod_{j \in S} \sum_{\tau \in \mathcal{W}_j} \prod_{v \in \tau} \gamma((v)) \right) \leq \lambda_{\text{init}} \gamma(E) \sum_{S \in \text{Ind}(\Gamma(E))} \prod_{j \in S} \psi_j ,$$

where the last equality follows from Lemma 4.8. The proof of Theorem 3.2 in the Shearer’s condition regime is identical, where instead of Lemma 4.8 we use Lemma 4.9.

### 4.4 Proof of Lemma 4.1

Throughout the proof, we will use ideas and definitions from [30]. We also note that we will assume w.l.o.g. that algorithm  $\mathcal{A}$  follows a deterministic flaw choice strategy. This is because randomized flaw choice strategies can equivalently be interpreted as convex combination of deterministic ones (and therefore, randomized strategies can be seen as taking expectation over deterministic ones).

For a trajectory  $\Sigma$  of length  $t$  we define

$$p(\Sigma) = \lambda_{\text{init}} \prod_{i=1}^t \rho_{w_i}(\sigma_i, \sigma_{i+1})$$

and notice that  $\Pr[\Sigma] \leq p(\Sigma)$ . Furthermore, we say that a trajectory  $\Sigma'$  is a proper prefix of  $\Sigma$  if  $\Sigma'$  is a prefix of  $\Sigma$  and  $\Sigma \neq \Sigma'$ .

**Definition 4.10** ([30]). *A set  $\mathcal{X}$  of trajectories of the algorithm will be called valid if (i) all trajectories in  $\mathcal{X}$  follow the same deterministic flaw choice strategy (not necessarily the same used by  $\mathcal{A}$ ); and (ii) for any  $\Sigma, \Sigma' \in \mathcal{X}$  trajectory  $\Sigma$  is not a proper prefix of  $\Sigma'$ .*

**Lemma 4.11** ([30]). *Consider a witness sequence  $W = (w_1, \dots, w_t)$  and a valid set of trajectories  $\mathcal{X}$  such that  $W$  is a prefix of  $W(\Sigma)$  for every  $\Sigma \in \mathcal{X}$ . We have that*

$$\sum_{\Sigma \in \mathcal{X}} p(\Sigma) \leq \lambda_{\text{init}} \prod_{i=1}^t \gamma(f_{w_i}) .$$

A *swap* is the operation of transforming a trajectory  $\Sigma = \dots \sigma_1 \xrightarrow{i} \sigma_2 \xrightarrow{j} \sigma_3 \dots$ , with  $i \approx j$ , to a trajectory  $\Sigma' = \dots \sigma_1 \xrightarrow{j} \sigma'_2 \xrightarrow{i} \sigma_3 \dots$ , where  $\sigma_1 \xrightarrow{j} \sigma'_2 \xrightarrow{i} \sigma_3 = \text{Swap}(\sigma_1 \xrightarrow{i} \sigma_2 \xrightarrow{j} \sigma_3)$ . A mapping  $\Phi$  on a set of trajectories will be called a *swapping mapping* if it operates by applying a sequence of swaps.

The main idea now will be to construct a swapping mapping whose goal will be to transform trajectories of the algorithm to a form that satisfies certain properties by applying swaps .

For a trajectory  $\Sigma$  in which a tree  $\tau \in \mathcal{W}_i$  occurs, we denote by  $W_\Sigma^\tau$  the prefix of  $W(\Sigma)$  up to the step that corresponds to the root of  $\tau$  (observe that Proposition 4.3 mandates that there exists a unique such step). Notice that, since  $\tau \in \mathcal{W}_i$ , the algorithm addresses flaw  $f_i$  at this step, and thus the final element of  $W_\Sigma^\tau$  is  $\{i\}$ . Finally, recall the definitions of  $\mathcal{R}_i^\pi$ ,  $\chi_\pi$  and  $\chi_i^\pi$ .

**Lemma 4.12.** *Fix a witness tree  $\tau \in \mathcal{W}_i$  and let  $\mathcal{X}^\tau$  be a valid set of trajectories in which  $\tau$  occurs. If  $\mathcal{A} = (F, \sim, \rho)$  is commutative then there exists a set of trajectories  $\mathcal{X}_\pi^\tau$  and a swapping mapping  $\Phi^\tau : \mathcal{X}^\tau \rightarrow \mathcal{X}_\pi^\tau$  which is a bijection such that*

- (a) *for any  $\Sigma \in \mathcal{X}_\pi^\tau$  we have that  $W_\Sigma^\tau$  is the unique witness sequence in  $\mathcal{R}_i^\pi$  such that  $\chi_i^\pi(W_\Sigma^\tau) = \chi_\pi(\tau)$ ;*
- (b) *for any witness sequence  $W$  the set  $\{\Sigma \in \mathcal{X}_\pi^\tau \mid \text{Rev}[W_\Sigma^\tau] = W\}$  is valid.*

We prove Lemma 4.12 in Section 4.5. To see Theorem 4.1, consider a witness tree  $\tau \in \mathcal{W}_i$ , and let  $\mathcal{Y}^\tau$  be the set of all trajectories that  $\mathcal{A}$  may follow in which  $\tau$  occurs. Now remove from  $\mathcal{Y}^\tau$  any trajectory  $\Sigma$  for which there exists a trajectory  $\Sigma'$  such that  $\Sigma$  is a proper prefix of  $\Sigma'$  to get  $\mathcal{X}^\tau$ . Clearly, this is a valid set and so recalling that  $\chi_\pi$  is a bijection and applying Lemma 4.12 we have that:

$$\Pr[\tau] = \sum_{\Sigma \in \mathcal{X}^\tau} \Pr[\Sigma] \leq \sum_{\Sigma \in \mathcal{X}^\tau} p(\Sigma) = \sum_{\Sigma \in \mathcal{X}_\pi^\tau} p(\Sigma) , \quad (6)$$

where to get the second equality we use the second requirement of Definition 2.1. Lemma 4.12 further implies that for every trajectory  $\Sigma \in \mathcal{X}_\pi^\tau$  we have that  $W_\Sigma^\tau$  is the (unique) witness sequence in  $\mathcal{R}_i^\pi$  such that  $\chi_i^\pi(W_\Sigma^\tau) = \chi_\pi(\tau)$ , i.e.,  $W_\Sigma^\tau = (\chi_i^\pi)^{-1}(\chi_\pi(\tau))$ . This means that the witnesses of the trajectories in  $\mathcal{X}_\pi^\tau$  have  $W := (\chi_i^\pi)^{-1}(\chi_\pi(\tau))$  as a common prefix. Since part (b) of Lemma 4.12 implies that  $\mathcal{X}_\pi^\tau$  is valid, applying Lemma 4.11 we get:

$$\sum_{\Sigma \in \mathcal{X}_\pi^\tau} p(\Sigma) \leq \lambda_{\text{init}} \prod_{w \in W} \gamma(f_w) = \lambda_{\text{init}} \prod_{v \in V(\tau)} \gamma(f_v) , \quad (7)$$

where the second inequality follows from the fact that  $\chi_i^\pi(W) = \chi_\pi(\tau)$  and  $V(\tau) = V(\chi_\pi(\tau))$ , concluding the proof.

## 4.5 Proof of Lemma 4.12

Our proof builds on the proof of Theorem 19 in [30]. We will be denoting witness sequences  $W = (w_1, w_2, \dots, w_t)$  as a sequence of *named indices of flaws*  $W = (\mathbf{w}_1, \dots, \mathbf{w}_t)$  where  $\mathbf{w}_j = (w_j, n_j)$  and  $n_j = |\{k \in [j] \mid w_k = w_j\}| \geq 1$  is the number of occurrences of  $w_j$  in the length- $j$  prefix of  $W$ . Note that a named index  $\mathbf{w}$  cannot appear twice in a sequence  $W$ . Finally, if  $\mathbf{w}$  is a named index of flaw we denote by  $w$  (that is, without bold font) the flaw index that is associated with it.

For a trajectory  $\Sigma$  such that  $W(\Sigma) = (\mathbf{w}_1, \dots, \mathbf{w}_t)$  we define a directed acyclic graph  $\mathbf{G}(\Sigma) = (\mathbf{V}(\Sigma), \mathbf{E}(\Sigma))$  where  $\mathbf{V}(\Sigma) = \{\mathbf{w}_1, \dots, \mathbf{w}_t\}$  and  $\mathbf{E}(\Sigma) = \{(\mathbf{w}_j, \mathbf{w}_k) \text{ s.t. } w_j \sim w_k \text{ and } j < k\}$ . This means that we have an edge from a named flow  $\mathbf{w}_i$  to another flow  $\mathbf{w}_j$  whenever their corresponding flow indices are related according to  $\sim$  and  $\mathbf{w}_j$  occurs in  $\Sigma$  before  $\mathbf{w}_k$ .

By Proposition 4.3, for any trajectory  $\Sigma$  in which  $\tau$  occurs there is a unique step  $t^* = t^*(\Sigma)$  such that  $\tau_{W(\Sigma)}(t^*) = \tau$ . For such a trajectory  $\Sigma$ , let  $\mathbf{Q}(\Sigma) \subseteq \mathbf{V}(\Sigma)$  be the set of flows from which the node  $\mathbf{w}_{t^*}$  can be reached in  $\mathbf{G}(\Sigma)$ , where  $\mathbf{w}_{t^*}$  is the named flow index that corresponds to the step  $t^*$ . Notice that  $w_{t^*} = i$  (since  $\tau \in \mathcal{W}_i$ ). For  $\mathbf{w} \in \mathbf{Q}(\Sigma)$  let  $d(\mathbf{w})$  be the length of the longest path from  $\mathbf{w}$  to  $\mathbf{w}_{t^*}$  in  $\mathbf{G}(\Sigma)$  plus one. For example,  $d(\mathbf{w}_{t^*}) = 1$ .

Let  $Q(\Sigma)$  denote the sequence consisting of the named flows in  $\mathbf{Q}(\Sigma)$  listed in the order they appear in  $\Sigma$ . The idea is to repeatedly apply the operation  $\text{Swap}$  to  $\Sigma$  so that we reach a trajectory  $\Sigma'$  that has a permutation  $Q_\pi(\Sigma)$  of  $Q(\Sigma)$  as a prefix. In particular, we will show that  $Q_\pi(\Sigma) \in \mathcal{R}_i^\pi$  and  $\chi_i^\pi(Q_\pi(\Sigma)) = \chi_\pi(\tau)$ .

To that end, for an integer  $r \geq 1$  define  $\mathbf{I}_r = \{\mathbf{w} \in \mathbf{Q}(\Sigma) \mid d(\mathbf{w}) = r\}$ , and let  $Q_r$  be the sequence consisting of the named flow indices in  $\mathbf{I}_r$  sorted in decreasing order with respect to  $\pi$ . Then, we define  $Q_\pi(\Sigma) = (Q_s, \dots, Q_1)$  where  $s = \max\{d(\mathbf{w}) \mid \mathbf{w} \in \mathbf{Q}(\Sigma)\}$ .

**Lemma 4.13.**  $Q_\pi(\Sigma) \in \mathcal{R}_i^\pi$  and  $\chi_i^\pi(Q_\pi(\Sigma)) = \chi_\pi(\tau)$ .

*Proof.* Let  $Y = Y(\Sigma) = \text{Rev}[Q_\pi(\Sigma)] = (Q_1, Q_2, \dots, Q_s)$  be the reverse sequence of  $\Sigma$ . By definition,  $R_Y = Q_1 = \{i\}$ . To show that  $Q \in \mathcal{R}_i^\pi$  it suffices to show that  $Q_{i+1} \subseteq \Gamma(Q_i)$  for each  $i \in [s-1]$ . To see this, recall the definitions of  $\mathbf{I}_{r+1}$  and  $Q_{r+1}$  and observe that, for each  $i_{r+1} \in Q_{r+1}$ , there must be a path of  $r$  indices of flows  $i_r, i_{r-1}, \dots, i_1$  such that for every  $j \in [r-1]$  we have that  $i_j \in Q_j$  and  $i_j \sim i_{j+1}$ .

Let  $k$  be the number of elements in witness sequence  $Q(\Sigma)$ . Recall that  $\chi_i^\pi(Q_\pi(\Sigma)) := \chi_\pi(\tau_{Q_\pi(\Sigma)}(k))$  (proof of Lemma 4.7). The proof is concluded by also recalling the algorithm for constructing witness trees and observing that  $\tau_{Q_\pi(\Sigma)}(k) = \tau_{W(\Sigma)}(t^*) = \tau$ . □

Note that applying  $\text{Swap}$  to  $\Sigma$  does not affect graph  $\mathbf{G}(\Sigma)$  and set  $\mathbf{Q}(\Sigma)$  and, thus, neither the sequence  $Q_\pi(\Sigma)$ . With that in mind, we show next how we could apply  $\text{Swap}$  repeatedly to  $\Sigma \in \mathcal{X}^\tau$  to reach a  $\Sigma'$  such  $Q_\pi(\Sigma)$  is a prefix of its witness sequence (that is,  $W(\Sigma') = (Q_\pi(\Sigma), U)$ ). We will do this by applying swaps to *swappable pairs* in  $\Sigma$ .

**Definition 4.14.** Consider a trajectory  $\Sigma \in \mathcal{X}^\tau$ . A pair  $(\mathbf{w}, \mathbf{y})$  of named indices of flows is called a *swappable pair* in  $\Sigma$  if it can be swapped in  $\Sigma$  (i.e.,  $W(\Sigma) = (\dots, \mathbf{w}, \mathbf{y}, \dots)$  and  $\mathbf{w} \approx \mathbf{y}$ ) and either

1.  $(\mathbf{w}, \mathbf{y}) \in (\mathbf{V}(\Sigma) \setminus \mathbf{Q}(\Sigma)) \times \mathbf{Q}(\Sigma)$ , or
2.  $(\mathbf{w}, \mathbf{y}) \in \mathbf{Q}(\Sigma) \times \mathbf{Q}(\Sigma)$  and their order in  $Q_\pi(\Sigma)$  is different:  $Q_\pi(\Sigma) = (\dots, \mathbf{y}, \mathbf{w}, \dots)$

The position of the rightmost swappable pair in  $\Sigma$  will be denoted as  $k(\Sigma)$ , where the position of  $(\mathbf{w}, \mathbf{y})$  in  $\Sigma$  is the number of named indices of flows that precede  $\mathbf{y}$  in  $W(\Sigma)$ . If  $\Sigma$  does not contain a swappable pair then  $k(\Sigma) = 0$ . Thus,  $k(\Sigma) \in [0, |\Sigma| - 1]$ .

We can only apply finite many swaps to swappable pairs in  $\Sigma$ . This is because swapping pairs of the first form moves a named index in  $\mathbf{Q}(\Sigma)$  to the left, while swapping pairs of the second one decrease the number of pairs whose relative order in  $Q(\Sigma)$  is not consistent with the one in  $Q_\pi(\Sigma)$ . Clearly, both of these actions can be performed only a finite number of times.

The following lemma shows how we can obtain a mapping  $\Phi^\tau$  such that  $\mathcal{X}_\pi^\tau := \Phi^\tau(\mathcal{X}^\tau)$  satisfies the first condition of Lemma 4.12. The proof is identical (up to minor changes) to the one of Lemma 27 of [30]. We include it in Section C for completeness.

**Lemma 4.15.** *Consider a trajectory  $\Sigma \in \mathcal{X}^\tau$  such that  $W(\Sigma) = (A, U)$  where  $A$  and  $U$  are some sequences of indices of flaws, and there are no swappable pairs inside  $U$ . Then  $U = (B, C)$  where sequence  $B$  is a subsequence of  $Q_\pi(\Sigma)$  and  $C$  does not contain named indices of flaws from  $\mathbf{Q}(\Sigma)$ .*

*In particular, if  $|A| = 0$  and  $W(\Sigma) = U$  does not contain a swappable pair then  $W(\Sigma) = (Q_\pi(\Sigma), C)$ .*

It remains to show that  $\Phi^\tau$  can be constructed so that is also a bijection and that it satisfies the second condition of Lemma 4.12. To do so, consider the following algorithm.

- Let  $\mathcal{X}_0 = \mathcal{X}^\tau$ .
- While  $k = \max_{\Sigma \in \mathcal{X}_p} k(\Sigma) \neq 0$ 
  - For each  $\Sigma \in \mathcal{X}_p$ : if  $k(\Sigma) = k$  then swap the pair  $(\mathbf{w}, \mathbf{y})$  at position  $k$  in  $\Sigma$ , otherwise leave  $\Sigma$  unchanged.
  - Let  $\mathcal{X}_{p+1}$  the new set of trajectories.

For a witness sequence  $W$  define  $\mathcal{X}_p[W] = \{\Sigma \in \mathcal{X}_p \mid Q_\pi(\Sigma) = W\}$  for an index  $p \geq 0$ . Now the following lemma concludes the proof since  $\mathcal{X}_0[W] \subseteq \mathcal{X}^\tau$  is valid. Its proof is identical (up to minor changes) to the proof of Lemma 28 in [30]. We also include it in Section C for completeness.

**Lemma 4.16.** *If set  $\mathcal{X}_p[W]$  is valid then so is  $\mathcal{X}_{p+1}[W]$ , and the mapping from  $\mathcal{X}_p[W]$  to  $\mathcal{X}_{p+1}[W]$  defined by the algorithm above is injective.*

## 5 Byproducts of Theorem 3.2

In this section we show two important byproducts of Theorem 3.2 which we will use in our applications.

### 5.1 Entropy of the Output Distribution

An important application of the known bounds for the Moser-Tardos distribution is estimating its randomness. In particular, Harris and Srinivasan [26] show that one can give lower bounds on the Rényi entropy of the output of the Moser-Tardos algorithm.

**Definition 5.1** ([12]). *Let  $\nu$  be a probability measure over a finite set  $S$ . The Rényi entropy with parameter  $\rho$  of  $\nu$  is defined to be*

$$H_\rho[\nu] = \frac{1}{1-\rho} \ln \sum_{s \in S} \nu(s)^\rho .$$

The min-entropy  $H_\infty$  is a special case defined as  $H_\infty[\nu] = \lim_{\rho \rightarrow \infty} H_\rho[\nu] = -\ln \max_{s \in S} \nu(s)$ .

Using the results of Section 3 we can show the analogous result in our setting.

**Theorem 5.2.** *Assume that  $\mathcal{A} = (F, \sim, \rho)$  is commutative, and the cluster expansion condition is satisfied. Let  $\nu$  be the output distribution of  $\mathcal{A}$ . Then, for  $\rho > 1$ ,*

$$H_\rho[\nu] \geq H_\rho[\mu] - \frac{\rho}{\rho-1} \ln \left( \sum_{S \in \text{Ind}([m])} \prod_{j \in S} \psi_j \right) - \frac{\rho}{\rho-1} \ln \lambda_{\text{init}} .$$

Given Theorem 3.2, the proof is akin to the analogous result in [26] and can be found in Appendix C .

**Remark 5.1.** Using Shearer’s condition we can replace  $\psi_i$  with  $\frac{q_{\{i\}}(\gamma)}{q_0(\gamma)}$ ,  $i \in [m]$ .

A straightforward application of having a lower bound on  $H_\rho[\nu]$  (for any  $\rho$ ), where  $\nu$  is the output distribution of the algorithm, is that there exist at least  $\exp(H_\rho[\nu])$  flawless objects. Before [26], the authors in [31] also used the (existential) LLL for enumeration of combinatorial structures by exploiting the fact that it guarantees a small probability  $p$  of avoiding all flaws when sampling from the uniform measure (and, thus, their number is at least  $p|\Omega|$ ).

## 5.2 Partially Avoiding Flaws

One of the main results of [21, 26] are LLL conditions for the existence of objects that avoid a large portion of the bad events. For instance, given a sufficiently sparse  $k$ -SAT formula which, nonetheless, violates the original LLL conditions, one can still find an assignment that satisfies many clauses. Using the results of Section 3, we are able to extend the (most general) result of Harris and Srinivasan [26] to the commutative setting.

Given a sequence of positive numbers  $\{\psi_i\}_{i=1}^m$ , for each  $i \in [m]$  define:

$$\zeta_i := \sum_{S \in \text{Ind}(\Gamma(i))} \prod_{j \in S} \psi_j ,$$

and notice that the cluster expansion condition can be expressed as requiring that for each  $i \in [m]$  we have that  $\gamma(f_i)\zeta_i \leq \psi_i$ .

**Theorem 5.3.** Assume that  $\mathcal{A} = (F, \sim, \rho)$  is commutative and  $\lambda_{\text{init}} = 1$ . Let  $\{\psi_i\}_{i=1}^m$  be a sequence of positive numbers. Then there is an algorithm  $\mathcal{A}'$  (which is a modification of  $\mathcal{A}$ ) and whose output distribution  $\nu$  has the property that for each  $i \in [m]$

$$\nu(f_i) \leq \max\{0, \gamma(f_i)\zeta_i - \psi_i\} .$$

Furthermore, the expected number of times a flaw  $f_i$  is addressed is at most  $\psi_i$ .

Given Theorem 3.2, the proof of Theorem 5.3 is akin to the one of [26] and can be found in Appendix C.

**Remark 5.2.** Using Shearer’s condition we can replace  $\psi_i$  with  $\frac{q_{\{i\}}(\gamma)}{q_0(\gamma)}$ ,  $i \in [m]$ .

## 6 Applications

In this section we show concrete applications of our main results in several problems.

### 6.1 Rainbow Matchings

In an edge-colored graph  $G = (V, E)$ , say that  $S \subseteq E$  is *rainbow* if its elements have distinct colors. In this section we consider the problem of finding rainbow matchings in complete graphs of size  $2n$ , where each color appears a limited amount of times.

Applying the cluster expansion condition, it can be shown [2, 27] that any edge-coloring of a complete graph of size  $2n$  in which each color appears on at most  $\frac{27}{128}n \approx 0.211n$  edges admits a rainbow perfect matching that can be found efficiently. Furthermore, in [30] it is shown that the resampling oracles defined by [27] for the space of matchings in a clique of even size, and which are used in this particular application, induce commutative algorithms. The latter fact will allow us to use our results to further study this problem.

### 6.1.1 Finding Rainbow Perfect Matchings

We first formulate the problem to fit our setting and use Theorem 3.2 to show that the algorithm of [2, 27] finds a perfect rainbow matching efficiently. Assuming a multiplicative slack in the cluster expansion conditions, a running time (number of steps) of  $O(n)$  can be given using the results of [2, 27, 30]. However, the best known upper bound without this assumption was given in [27] to be  $O(n^2)$ . Here we improve the latter to  $O(n)$ .

Let  $\phi$  be any edge-coloring of  $K_{2n}$  in which each color appears on at most  $\lambda n$  edges. Let  $P = P(\phi)$  be the set of all pairs of vertex-disjoint edges with the same color in  $\phi$ , i.e.,  $P = \{\{e_1, e_2\} : \phi(e_1) = \phi(e_2)\}$ . Let  $\Omega$  be the set of all perfect matchings of  $K_{2n}$ . For each  $\{e_i, e_j\} \in P$  let

$$f_{i,j} = \{M \in \Omega : \{e_i, e_j\} \subset M\} .$$

Thus, an element of  $\Omega$  is flawless iff it is a rainbow perfect matching. The algorithm that finds a rainbow perfect matching starts at a state of  $\Omega$  chosen uniformly at random and, in every subsequent step, it chooses (arbitrarily) a flaw to address. Algorithm 1 below describes the probability distributions  $\rho_{i,j}(M, \cdot)$ , where  $M \in f_{i,j}$ . This is a special case of the implementation of a general resampling oracle with respect to the uniform measure over  $\Omega$  for perfect matchings described in [27]. For the problem of rainbow matchings, the latter implies that  $\gamma(f_{i,j}) = \mu(f_{i,j}) = \frac{1}{(2n-1)(2n-3)}$ .

---

#### Algorithm 1 Probability Distribution $\rho_{i,j}(M, \cdot)$

---

- 1:  $M' := M, A := \{e_1, e_2\}, A' := A$ .
  - 2: **while**  $A' \neq \emptyset$  **do**
  - 3:   Pick  $(u, v) \in A'$  arbitrarily
  - 4:   Pick  $(x, y) \in M' \setminus A'$  uniformly at random, with  $(x, y)$  randomly ordered;
  - 5:   With probability  $1 - \frac{1}{2^{|M' \setminus A'|+1}}$ ,  
       Add  $(u, y), (v, x)$  to  $M'$  and remove  $(u, v), (x, y)$  from  $M'$ ;
  - 6:   Remove  $(u, v)$  from  $A'$ ;
  - 7: **Output**  $M'$ .
- 

For a vertex  $v$  let  $\Gamma(v)$  denote the set of indices of flaws that correspond to edges adjacent to  $v$ . By observing the algorithm it's not hard to verify (and is also proved in [2, 27, 30]) that the graph  $C$  over indices of flaws such that for each  $(e_i = (v_1, v_2), e_j = (v_3, v_4)) \in P$  we have that

$$\Gamma(i, j) = \bigcup_{k=1}^4 \Gamma(v_k)$$

is a causality graph. Furthermore, if  $S \in \text{Ind}(\Gamma(i, j))$ , then for each  $k \in \{1, 2, 3, 4\}$  we have that  $|S \cap \Gamma(v_k)| \leq 1$ . This means that  $|S| \leq 4$  and, moreover, for each  $j \in \{0, 1, 2, 3, 4\}$  there are at most  $\binom{4}{j} (2n-1)^j (\lambda n - 1)^j$  subsets  $S \in \text{Ind}(\Gamma(i, j))$  of size  $j$ . Choosing parameters  $\psi_{i,j} = \psi = \frac{3}{4n^2}$  we have that:

$$\gamma(f_{i,j}) \zeta_{i,j} := \gamma(f_{i,j}) \sum_{S \in \text{Ind}(\Gamma(i,j))} \psi^{|S|} \leq \frac{1}{(2n-3)(2n-1)} (1 + (2n-1)(\lambda n - 1))^4 ,$$

from which it can be seen that whenever  $\lambda \leq \frac{27}{128}$  we have that  $\gamma(f_{i,j}) \zeta_{i,j} \leq 1$  and so the cluster expansion condition is satisfied.

Since  $|P| \leq (2n)^2 \cdot (\lambda n - 1) < 4\lambda n^3$ , Theorem 3.2 implies that the algorithm terminates after an expected number of  $3\lambda n$  steps. Overall, we have showed the following theorem.

**Theorem 6.1.** For any  $\lambda \leq \frac{27}{128}$ , given any edge-coloring of the complete graph on  $2n$  vertices in which each color appears on at most  $\lambda n$  edges, there exists an algorithm that terminates in an expected number of at most  $3\lambda n$  steps and outputs a rainbow perfect matching.

### 6.1.2 Number of Rainbow Perfect Matchings

In this subsection we use Theorem 5.2 to give an exponential lower bound on the number of perfect matchings when each color appears at most  $\lambda n$  times, where  $\lambda \leq \frac{27}{128}$ , by bounding the entropy of the output distribution of the algorithm described in the previous subsection.

**Theorem 6.2.** For any  $\lambda \leq \frac{27}{128}$ , given any edge-coloring of the complete graph on  $2n$  vertices in which each color appears on at most  $\lambda n$  edges, there exist at least  $^1 e^{-3\lambda n} \cdot (2n-1)!!$  rainbow perfect matchings. Furthermore, there exists an algorithm that outputs each one of them with positive probability.

*Proof.* To apply Theorem 5.2, we will need to give an upper bound for  $\sum_{S \in \text{Ind}([m])} \psi^{|S|}$ . Similarly to applications in [26], we will find useful the following crude, but general upper bound:

$$\sum_{S \in \text{Ind}([m])} \prod_{j \in S} \psi_j \leq \sum_{S \subseteq [m]} \prod_{j \in S} \psi_j \leq \prod_{i \in [m]} (1 + \psi_i) \leq \exp\left(\sum_{i \in [m]} \psi_i\right). \quad (8)$$

Since the number of perfect matching in  $K_{2n}$  is  $(2n-1)!!$  and also  $|P| < 4\lambda n^3$ , Theorem 5.2 and (8) imply that the number of rainbow perfect matchings is at least

$$\exp\left(\ln |\Omega| - \sum_{i \in [m]} \psi\right) \geq \exp(\ln((2n-1)!!) - 3\lambda n) = \frac{(2n-1)!!}{e^{3\lambda n}},$$

concluding the proof. □

### 6.1.3 Low Weight Rainbow Perfect Matchings

Consider an arbitrary weighting function  $W : E \rightarrow \mathbb{R}$  over the edges of  $K_{2n}$ . Here we consider the problem of finding rainbow perfect matchings of low weight, where the weight of a matching is defined as the sum of weights of its edges. Clearly, there is a selection of  $n$  edges of  $K_{2n}$  whose total weight is at most  $\frac{1/2}{2n-1} \sum_{e \in K_{2n}} W(e)$ . We use Theorem 3.2 to show that, whenever  $\lambda \leq \frac{27}{128}$ , the algorithm of subsection 6.1.1 outputs a rainbow perfect matching of similar expected weight.

**Theorem 6.3.** For any  $\lambda \leq \frac{27}{128}$ , given any edge-coloring of the complete graph on  $2n$  vertices in which each color appears on at most  $\lambda n$  edges, there exists an algorithm that outputs a perfect rainbow matching  $M$  such that

$$\mathbb{E}[W(M)] \leq \frac{(1 + \frac{3}{2}\lambda)^2}{2n-1} \sum_{e \in K_{2n}} W(e).$$

*Proof.* Let  $A_e$  be the subset of  $\Omega$  that consists of the matchings that contain  $e$ . It is proven in [27], and it's also not hard to verify, that Algorithm 1 with  $A = \{e\}$  is a resampling oracle for this type of flaw. Moreover, using an identical counting argument to the one in subsection 6.1.1 we get that:

$$\sum_{S \in \text{Ind}(\Gamma(A_e))} \psi^{|S|} \leq (1 + (2n-1)(\lambda n - 1)\psi)^2.$$

---

<sup>1</sup>Recall that  $(2n-1)!! = 1 \cdot 3 \cdot \dots \cdot (2n-1) = \frac{(2n)!}{2^n n!}$ .

Applying Theorem 3.2 we get that:

$$\begin{aligned}
\mathbb{E}[W(M)] &\leq \sum_{e \in K_{2n}} W(e) \Pr[A_e] \\
&\leq \sum_{e \in K_{2n}} W(e) \mu(A_e) (1 + (2n-1)(\lambda n - 1)\psi)^2 \\
&< \frac{(1 + \frac{3}{2}\lambda)^2}{2n-1} \sum_{e \in K_{2n}} W(e) ,
\end{aligned}$$

concluding the proof.  $\square$

#### 6.1.4 Finding Rainbow Matchings with many edges

In this subsection we use Theorem 5.3 to show that whenever  $\lambda < 0.5$  we can find rainbow matchings with a linear number of edges.

**Theorem 6.4.** *Given any edge-coloring of the complete graph on  $2n$  vertices in which each color appears on at most  $\lambda n$  edges, where  $\lambda < 0.5$  and  $n$  is sufficiently large, there exists an algorithm that terminates within  $O(n)$  steps in expectation and finds a rainbow matching with an expected number of edges that is at least  $n \min\left(1, 0.94 \sqrt[3]{\frac{2}{\lambda}} - 1\right)$ .*

*Proof.* Let  $\phi$  be any edge-coloring of  $K_{2n}$  in which each color appears on at most  $\lambda n$  edges and recall the definitions of  $P = P(\phi)$ ,  $\Omega$ , and  $f_{i,j} = \{M \in \Omega : \{e_i, e_j\} \subset M\}$  from the proof of Theorem 6.2.

The idea is to apply Theorem 5.3 that guarantees that we can come up with a “truncated version”  $\mathcal{A}'$  of our algorithm for finding perfect rainbow matchings. In particular, if  $\nu$  is the output probability distribution of  $\mathcal{A}'$  and for each flaw  $f_{i,j}$  we set  $\psi_{i,j} = \alpha$  then:

$$\begin{aligned}
\nu(f_{i,j}) &\leq \max(0, \gamma(f_{i,j})\zeta_{i,j} - \alpha) \\
&\leq \max\left(0, \frac{(1 + (2n-1)(\lambda n - 1)\alpha)^4}{(2n-3)(2n-1)} - \alpha\right) .
\end{aligned} \tag{9}$$

Consider now the following strategy: We first execute algorithm  $\mathcal{A}'$  to get a perfect, possibly non-rainbow, matching  $M$  of  $K_{2n}$ . Then, for each flaw  $f_{i,j}$  that appears in  $M$ , we choose arbitrarily one of its corresponding edges and remove it from  $M$ , to get a non-perfect, but rainbow, matching  $M'$ . If  $S = S(M')$  is the random variable that equals the size (number of edges) of  $M'$  then by setting

$$\alpha = \frac{1}{(2n-1)(\lambda n - 1)} \left( \sqrt[3]{\frac{2n-3}{4(\lambda n - 1)}} - 1 \right)$$

we get:

$$\begin{aligned}
\mathbb{E}[S] &= n - \sum_{(e_i, e_j) \in P} \nu(f_{i,j}) \\
&\geq n - \max\left(0, |P| \left( \frac{(1 + (2n-1)(\lambda n - 1)\alpha)^4}{(2n-3)(2n-1)} - \alpha \right)\right) \\
&= n \min\left(1, 1 - \frac{4n}{2n-1} \left(1 - \frac{3}{4 \cdot 2^{2/3}} \sqrt[3]{\frac{2n-3}{\lambda n - 1}}\right)\right) .
\end{aligned}$$

For large enough  $n$ , the latter is  $\min\left(1, 0.94 \sqrt[3]{\frac{2}{\lambda}} - 1\right)$ . Finally, notice that (for large  $n$ )  $\alpha$  is positive whenever  $\lambda < 0.5$ .  $\square$

## 6.2 List-coloring Triangle-Free Graphs

In the problem of list-coloring one is given a graph  $G = G(V, E)$  over  $n$  vertices  $V = \{v_1, v_2, \dots, v_n\}$  and, for each  $v \in V$ , a list of colors  $\mathcal{L}_v$ . The goal is to find a list-coloring  $\sigma \in \mathcal{L}_{v_1} \times \dots \times \mathcal{L}_{v_n}$  of  $G$  such that  $\sigma(v) \neq \sigma(u)$  for any pair of adjacent vertices.

The *list chromatic number*  $\chi_\ell(G)$  of a graph  $G$  is the minimum number of colors for which such a coloring is attainable. A celebrated result of Johansson shows that there exist 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, Molloy [32] improved Johansson’s result showing that  $C$  can be replaced by  $(1 + \epsilon)$  for any  $\epsilon > 0$  assuming that  $\Delta \geq \Delta_\epsilon$ . (We note that, soon after, Bernshteyn [8] established the same bound for the list chromatic number using the LLL. However, his result is not constructive as it uses a sophisticated probability measure for which it is not clear how one could design “efficient” resampling oracles.)

Here we show how that the algorithm of Molloy is amenable to our analysis and, in particular, we prove that it can output exponentially many proper colorings with positive probability.

### 6.2.1 The Algorithm

The algorithm of [32] works in two stages. First, it finds a partial list-coloring which has the property that (i) each vertex  $v \in V$  has “many” available colors; (ii) there is not “too much competition” for the available colors of  $v$ , i.e., they do not appear in the list of available colors of its neighbors. Then, it completes the coloring via a fairly straightforward application of the Moser-Tardos algorithm.

To describe the algorithm formally, we will need some further notation. First, it will be convenient to treat Blank as a color that is in the list of every vertex. For each vertex  $v$  and *partial* proper coloring  $\sigma$  let

- $N_v$  denote the set of vertices adjacent to  $v$ ;
- $L_v(\sigma) \subseteq \mathcal{L}_v$  to be the set of *available* colors for  $v$  at state  $\sigma$ , i.e., the set of colors that we can assign to  $v$  in  $\sigma$  without making any edge monochromatic. Notice that Blank is always an available color;
- $T_{v,c}(\sigma)$  to be the set of vertices  $u \in N_v$  such that  $\sigma(u) = c$  and  $c \in L_u(\sigma)$ .

Let  $\Omega = \prod_{v \in V} \mathcal{L}_v$  and define  $L = \Delta^{\frac{\epsilon}{2}}$ . Given a partial list-coloring, we define the following flaws for any vertex  $v$ :

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

$$Z_v = \left\{ \sigma \in \Omega : \sum_{c \in L_v(\sigma) \setminus \text{Blank}} |T_{v,c}(\sigma)| > \frac{1}{10} L \cdot |L_v(\sigma)| \right\}.$$

**Lemma 6.5** (The Second Phase). *Given a flawless partial list-coloring, a complete list-coloring of  $G$  can be found in expected polynomial time.*

Lemma 6.5 was proved in [32] via a fairly straightforward application of the Lovász Local Lemma, and can be made constructive via the Moser-Tardos algorithm. We also present its proof in Appendix D.1, as it will be useful in our analysis. What is left is to describe the first phase of the algorithm.

- The initial distribution  $\theta$ , which is important in this case, is chosen to be the following: Fix an independent set  $S$  of  $G$  of size at least  $n/(\Delta + 1)$ . (This is trivial to find efficiently via a greedy algorithm). Choose one color from  $\mathcal{L}_u \setminus \text{Blank}$ ,  $u \in S$ , uniformly at random, and assign it to  $u$ ;
- to address a flaw  $f \in \{B_v, Z_v\}$  at state  $\sigma$ , for each  $u \in N_v$ , choose uniformly at random a color from  $L_u(\sigma)$  and assign it to  $u$ ;

- as a flaw choice strategy, the algorithm first fixes any ordering  $\pi$  over flaws. At every step, it chooses the lowest occurring flaw according to  $\pi$  and address it.

### 6.2.2 Proving Termination

Let  $\mathcal{A}_1$ ,  $\mathcal{A}_2$  denote the first and second phase of our algorithm, respectively. Here we prove that  $\mathcal{A}_1$  terminates in expected polynomial time. To do so, we will use the convergence result corresponding to equation (3). (Although, as we will see,  $\mathcal{A}_1$  is commutative for an appropriate choice of a causality graph, we won't use Theorem 3.2 to prove its convergence. This is because  $\lambda_{\text{init}}$  is exponentially large in this case).

The measure  $\mu$  we use for the analysis is the uniform measure over partial proper colorings. We will use the following lemma whose proof can be found in Section 6.3.1.

**Lemma 6.6.** *For each vertex  $v$  and flaw  $f \in \{B_v, Z_v\}$  we have that*

$$\gamma(f) \leq 2\Delta^{-4} .$$

Consider the causality graph such that  $f_v \sim f_u$ , if  $\text{dist}(u, v) \leq 3$ , where  $f_v, f_u$  are either  $B$ -flaws or  $Z$ -flaws. Notice that it has maximum degree at most  $2(\Delta^3 + 1)$ . Setting  $\psi_f = \psi = \frac{1}{2(\Delta^3 + 1)}$  for every flaw  $f \in \{B_v, Z_v\}$  and applying (3), we get that the algorithm converges in expected polynomial time since

$$\gamma(f) \sum_{S \subseteq \Gamma(f)} \prod_{g \in S} \psi_g \leq \frac{2}{\Delta^4} \cdot 2(\Delta^3 + 1) \cdot e < 4e \left( \frac{1}{\Delta} + \frac{1}{\Delta^4} \right) < 1 ,$$

for large enough  $\Delta$ , and  $\log_2 |\Omega| + 2n \log_2 (1 + \psi) = O(n \log n)$ .

### 6.3 A Lower Bound on the Number of Possible Outputs

In this section it will be convenient to assume that the list of each vertex has size exactly  $q$ .

Let  $\mathcal{A}_1$ ,  $\mathcal{A}_2$  denote the first and second phase of our algorithm, respectively. The bound regarding the number of list-colorings the algorithm can output with positive probability follows almost immediately from the two following lemmas.

**Lemma 6.7.** *Algorithm  $\mathcal{A}_1$  can output at least  $\exp\left(n \left(\frac{\ln q}{\Delta+1} - \frac{1}{\Delta^3}\right)\right)$  flawless partial colorings with positive probability.*

*Proof.* It is not hard to verify that algorithm  $\mathcal{A}_1$  is commutative with respect to the causality relation  $\sim$  induced by neighborhoods  $\Gamma(\cdot)$ . To see this, notice that for any two flaws  $f_v, f_u$  and any  $\sigma \in f_v \cap f_u$ , invoking procedure  $\text{RESAMPLE}(v, \sigma)$  does not change the list of available colors of the neighbors of  $u$ . Applying Theorem 5.2 (using the crude bound we saw in (8)) we get that the algorithm can output at least

$$\exp\left(\ln \frac{|\Omega|}{\lambda_{\text{init}}} - \sum_{f \in F} \psi_f\right) = \exp\left(\ln \frac{1}{\max_{\sigma \in \Omega} \theta(\sigma)} - \frac{2n}{2(\Delta^3 + 1)}\right) > \exp\left(n \left(\frac{\ln q}{\Delta + 1} - \frac{1}{\Delta^3}\right)\right) \quad (10)$$

flawless partial colorings. □

**Lemma 6.8.** *Suppose  $\mathcal{A}_1$  can output  $N$  flawless partial colorings with positive probability. Suppose further that among these partial colorings, the ones with the lowest number of colored vertices have exactly  $\alpha n$  vertices colored, where  $\alpha \in (0, 1)$ . Then,  $\mathcal{A}_2$  can output at least  $\max\left(Nq^{-(1-\alpha)n}, \left(\frac{8L}{11}\right)^{(1-\alpha)n}\right)$  list-colorings with positive probability.*

The proof of Lemma 6.8 can be found in Appendix D.

*Proof of Theorem 1.2.* According to Lemma 6.7 algorithm  $\mathcal{A}_1$  can output at least  $N := \exp\left(n\left(\frac{\ln q}{\Delta+1} - \frac{1}{\Delta^3}\right)\right)$  flawless partial colorings. Moreover, according to Lemma 6.8, algorithm  $\mathcal{A}_2$  can output at least

$$\min_{\alpha \in (0,1)} \left\{ \max \left( Nq^{-(1-\alpha)n}, \left(\frac{8L}{11}\right)^{(1-\alpha)n} \right) \right\}$$

distinct full-list colorings. Since  $Nq^{-(1-\alpha)n}$ ,  $\left(\frac{8L}{11}\right)^{(1-\alpha)n}$  are increasing and decreasing as functions of  $\alpha$ , respectively, the value of  $\alpha$  that minimizes our lower bound is the one that makes them equal, which can be seen to be

$$\alpha^* := 1 - \frac{\ln N}{n \ln\left(\frac{8Lq}{11}\right)} = 1 - \frac{\frac{\ln q}{\Delta+1} - \frac{1}{\Delta^3}}{\ln q + \ln\left(\frac{8L}{11}\right)}.$$

Therefore, algorithm  $\mathcal{A}_2$  can output at least

$$\exp\left(n\left(\frac{\ln q}{\Delta+1} - \frac{1}{\Delta^3}\right)\right) \cdot q^{-(1-\alpha^*)n} = \exp\left(n\left(\frac{q}{\Delta+1} - \frac{1}{\Delta^3}\right)(1-\delta)\right),$$

list-colorings, where  $\delta := \frac{1}{1 + \frac{\ln(8L/11)}{\ln q}} \in (0, 1)$ , concluding the proof. □

### 6.3.1 Proof of Lemma 6.6

It will be convenient to extend the notion of “addressing a flaw  $f$  in a state  $\sigma$ ” to arbitrary states  $\sigma \in \Omega$ , meaning that we recolor the vertices associated with  $f$  in the same way we would do it if the constraint corresponding to  $f$  was indeed violated. Consider the following random experiments.

- Address  $B_v$  at an arbitrary state  $\sigma \in \Omega$  to get a state  $\sigma'$ . Let  $\Pr_\sigma[B_v]$  denote the probability that  $\sigma' \in B_v$ .
- Address  $Z_v$  at an arbitrary state  $\sigma \in \Omega$  to get a state  $\sigma'$ . Let  $\Pr_\sigma[Z_v]$  denote the probability that  $\sigma' \in Z_v$ .

Our claim now is that

$$\gamma(B_v) \leq \max_{\sigma' \in \Omega} \Pr_{\sigma'}[B_v]; \tag{11}$$

$$\gamma(Z_v) \leq \max_{\sigma' \in \Omega} \Pr_{\sigma'}[Z_v]. \tag{12}$$

To see this, let  $f_v \in \{B_v, Z_v\}$  and observe that

$$\gamma(f_v) = \max_{\sigma' \in \Omega} \sum_{\sigma \in f_v} \frac{\mu(\sigma)}{\mu(\sigma')} \rho_v(\sigma, \sigma') = \max_{\sigma' \in \Omega} \sum_{\sigma \in \text{In}_{f_v}(\sigma')} \frac{1}{|\Lambda(\sigma)|},$$

where  $\Lambda(\sigma) := \prod_{u \in N_v} L_u(\sigma)$  is the cartesian product of the lists of available colors of each vertex  $u \in N_v$  at state  $\sigma$  and  $\text{In}_{f_v}(\sigma')$  is the set of states  $\sigma \in f_v$  such that  $\sigma' \in A(f_v, \sigma)$ .

The key observation now is that  $\Lambda(\sigma') = \Lambda(\sigma)$  for each state  $\sigma \in \text{In}_{f_v}(\sigma')$ . This is because any transition of the form  $\sigma \xrightarrow{f_v} \sigma'$  does not alter the lists of available colors of vertices  $u \in N_v$ , since the graph is triangle-free. Thus,

$$\gamma(f_v) = \max_{\sigma' \in \Omega} \frac{|\text{In}_{f_v}(\sigma')|}{\Lambda(\sigma')} = \max_{\sigma' \in \Omega} \Pr_{\sigma'}[f_v] ,$$

where the second equality follows from the fact that there is a bijection between  $\text{In}_{f_v}(\sigma')$  and the set of color assignments from  $\Lambda(\sigma')$  to the vertices of  $N_v$  that violate the constraint related to flaw  $f_v$ .

The following lemma concludes the proof.

**Lemma 6.9** ([32]). *For every state  $\sigma \in \Omega$  it holds that*

- (a)  $\Pr_{\sigma}[B_v] < \Delta^{-4}$ ;
- (b)  $\Pr_{\sigma}[Z_v] < \Delta^{-4}$ .

## 6.4 Acyclic Edge Coloring via the Clique Lovász Local Lemma

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 (AEC), was originally motivated by the work of Coleman et al. [15, 14] on the efficient computation of Hessians and, since then, there has been a series of works [5, 33, 37, 21, 29, 18] that upper bound  $\chi'_a(G)$  for graphs with bounded degree. The currently best result was given recently by Giotis et al. in [19] who showed that  $\chi'_a(G) \leq 3.74\Delta$  in graphs with maximum degree  $\Delta$ .

The analysis of [19], while inspired by the algorithmic LLL, uses a custom argument that does not correspond to any of its known versions. Furthermore, their algorithm does not correspond to an instantiation of the Moser Tardos algorithm and does not seem to be commutative (assuming the natural formulation in our setting) and, thus, it's not amenable to our analysis.

On the other hand, Kolipaka, Szegedy and Yixin Xu show in [29] that  $8.6(\Delta - 1)$  colors suffice for the Moser Tardos algorithm to converge in this setting. They do this by introducing the *Clique the Lovász Local Lemma*, a condition that is typically stronger than (although, technically, incomparable to) the cluster expansion condition, but weaker than the Shearer's condition. In fact, the Clique Lovász Local Lemma is a member of a "hierarchy" of LLL conditions that are increasingly complex and use an increasing amount of information about the structure of the dependency graph. On the limit, they give the Shearer's condition.

While the use of the Clique Lovász Local Lemma (or any other condition in the hierarchy of [29]) makes the results of [21] inapplicable, it does allow us to use Theorems 3.2 and 5.2 which capture the cases where the Shearer's condition is satisfied.

We show two results. Our first theorem says that the Moser Tardos algorithm applied on the acyclic edge coloring problem converges in polynomial time and has high output entropy whenever  $q \geq 8.6(\Delta - 1)$ .

**Theorem 6.10.** *Given a graph  $G = (V, E)$  with maximum degree  $\Delta$  and  $q \geq 8.6(\Delta - 1)$  colors, there exist at least  $(\frac{q}{4})^{|E|}$  acyclic edge colorings of  $G$ . Furthermore, there exists an algorithm with expected polynomial running time that outputs each one of them with positive probability.*

Our second theorem considers a problem of weighted acyclic edge colorings. In particular, given a graph  $G(V, E)$  let  $W = \sum_{v \in V} W_v$  be a weighting function over edge  $q$ -colorings of  $G$  such that each  $W_v, v \in V$ , is a function of the colors of the edges adjacent to  $v$ . By sampling uniformly at random, one can find an edge coloring  $\phi$  of weight  $\mathbb{E}_{\phi \sim \mu}[W(\phi)]$ , where  $\mu$  is the uniform distribution over the edge  $q$ -colorings of  $G$ . Using Theorem 3.2, we show that whenever  $q \geq 8.6(\Delta - 1)$  we can use the Moser-Tardos algorithm to find an acyclic edge coloring of similar weight (assuming that  $\Delta$  is constant).

**Theorem 6.11.** *Given a graph  $G(V, E)$  with maximum degree  $\Delta$  and  $q \geq 8.6(\Delta - 1)$  colors, there exist an algorithm with expected polynomial running time that outputs an acyclic edge coloring  $\phi_{\text{out}}$  of expected weight at most*

$$\mathbb{E}[\phi_{\text{out}}] < 1.3^\Delta \cdot \mathbb{E}_{\phi \sim \mu}[W(\phi)] .$$

The proofs of Theorems 6.10, 6.11, as well as more details on the Clique Lovász Local Lemma, can be found in Section D.

## 7 Acknowledgements

The author is grateful to Dimitris Achlioptas and Alistair Sinclair for detailed comments and feedback, as well as to anonymous reviewers for comments and remarks.

## References

- [1] Dimitris Achlioptas and Fotis Iliopoulos. Focused stochastic local search and the Lovász local lemma. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 2024–2038, 2016.
- [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, July 2016.
- [3] Dimitris Achlioptas, Fotis Iliopoulos, and Alistair Sinclair. A new perspective on stochastic local search and the lovasz local lemma. *CoRR*, abs/1805.02026, 2018.
- [4] Noga Alon. A parallel algorithmic version of the local lemma. *Random Struct. Algorithms*, 2(4):367–378, 1991.
- [5] Noga Alon. A parallel algorithmic version of the local lemma. *Random Structures & Algorithms*, 2(4):367–378, 1991.
- [6] Noga Alon, Colin McDiarmid, and Bruce Reed. Acyclic coloring of graphs. *Random Structures & Algorithms*, 2(3):277–288, 1991.
- [7] József Beck. An algorithmic approach to the Lovász local lemma. I. *Random Structures Algorithms*, 2(4):343–365, 1991.
- [8] Anton Bernshteyn. The johansson–molloy theorem for dp-coloring. *arXiv preprint arXiv:1708.03843*, 2017.
- [9] 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.
- [10] Karthekeyan Chandrasekaran, Navin Goyal, and Bernhard Haeupler. Deterministic algorithms for the Lovász local lemma. *SIAM J. Comput.*, 42(6):2132–2155, 2013.
- [11] Antares Chen, David G. Harris, and Aravind Srinivasan. Partial resampling to approximate covering integer programs. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1984–2003. SIAM, 2016.

- [12] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.
- [13] Kai-Min Chung, Seth Pettie, and Hsin-Hao Su. Distributed algorithms for the Lovász local lemma and graph coloring. In Magnús M. Halldórsson and Shlomi Dolev, editors, *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, pages 134–143. ACM, 2014.
- [14] Thomas F Coleman and Jin Yi Cai. The cyclic coloring problem and estimation of sparse hessian matrices. *SIAM J. Algebraic Discrete Methods*, 7(2):221–235, April 1986.
- [15] Thomas F. Coleman and Moré Jorge J. Estimation of sparse hessian matrices and graph coloring problems. *Mathematical Programming*, 28(3):243–270, 1984.
- [16] Artur Czumaj and Christian Scheideler. Coloring non-uniform hypergraphs: a new algorithmic approach to the general Lovász local lemma. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, CA, 2000)*, pages 30–39, 2000.
- [17] 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.
- [18] Louis Esperet and Aline Parreau. Acyclic edge-coloring using entropy compression. *European Journal of Combinatorics*, 34(6):1019–1027, 2013.
- [19] Ioannis Giotis, Lefteris M. Kirousis, Kostas I. Psaromiligkos, and Dimitrios M. Thilikos. Acyclic edge coloring through the Lovász local lemma. *Theor. Comput. Sci.*, 665:40–50, 2017.
- [20] Bernhard Haeupler and David G Harris. Parallel algorithms and concentration bounds for the Lovász local lemma via witness-DAGs. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1170–1187. SIAM, 2017.
- [21] Bernhard Haeupler, Barna Saha, and Aravind Srinivasan. New constructive aspects of the Lovász local lemma. *J. ACM*, 58(6):Art. 28, 28, 2011.
- [22] David G. Harris. New bounds for the Moser-Tardos distribution: Beyond the Lovasz local lemma. *CoRR*, abs/1610.09653, 2016.
- [23] David G. Harris. Oblivious resampling oracles and parallel algorithms for the lopsided Lovász local lemma. *CoRR*, abs/1702.02547, 2017.
- [24] David G. Harris and Aravind Srinivasan. The Moser-Tardos framework with partial resampling. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 469–478. IEEE Computer Society, 2013.
- [25] David G. Harris and Aravind Srinivasan. A constructive algorithm for the Lovász local lemma on permutations. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 907–925. SIAM, 2014.
- [26] David G. Harris and Aravind Srinivasan. Algorithmic and enumerative aspects of the Moser-Tardos distribution. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 2004–2023, 2016.

- [27] Nicholas J. A. Harvey and Jan Vondrák. An algorithmic proof of the Lovász local lemma via resampling oracles. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1327–1346. IEEE Computer Society, 2015.
- [28] Kashyap Babu Rao Kolipaka and Mario Szegedy. Moser and Tardos meet Lovász. In *STOC*, pages 235–244. ACM, 2011.
- [29] Kashyap Babu Rao Kolipaka, Mario Szegedy, and Yixin Xu. A sharper local lemma with improved applications. In Anupam Gupta, Klaus Jansen, José D. P. Rolim, and Rocco A. Servedio, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX 2012, and 16th International Workshop, RANDOM 2012, Cambridge, MA, USA, August 15-17, 2012. Proceedings*, volume 7408 of *Lecture Notes in Computer Science*, pages 603–614. Springer, 2012.
- [30] Vladimir Kolmogorov. Commutativity in the algorithmic Lovász local lemma. In Irit Dinur, editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 780–787. IEEE Computer Society, 2016.
- [31] Linyuan Lu and Laszlo A Szekely. A new asymptotic enumeration technique: the Lovász local lemma. *arXiv preprint arXiv:0905.3983*, 2009.
- [32] Michael Molloy. The list chromatic number of graphs with small clique number. *arXiv preprint arXiv:1701.09133*, 2017.
- [33] Michael Molloy and Bruce Reed. Further algorithmic aspects of the local lemma. In *STOC '98 (Dallas, TX)*, pages 524–529. ACM, New York, 1999.
- [34] Michael Molloy and Bruce Reed. *Graph colouring and the probabilistic method*, volume 23 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2002.
- [35] Robin A. Moser. A constructive proof of the Lovász local lemma. In *STOC'09—Proceedings of the 2009 ACM International Symposium on Theory of Computing*, pages 343–350. ACM, New York, 2009.
- [36] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2):Art. 11, 15, 2010.
- [37] Sokol Ndreca, Aldo Procacci, and Benedetto Scoppola. Improved bounds on coloring of graphs. *Eur. J. Comb.*, 33(4):592–609, May 2012.
- [38] Christos H. Papadimitriou. On selecting a satisfying truth assignment. In *FOCS*, pages 163–169. IEEE Computer Society, 1991.
- [39] Wesley Pegden. An extension of the Moser-Tardos algorithmic local lemma. *SIAM J. Discrete Math.*, 28(2):911–917, 2014.
- [40] J.B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3):241–245, 1985.
- [41] Aravind Srinivasan. Improved algorithmic versions of the Lovász local lemma. In Shang-Hua Teng, editor, *SODA*, pages 611–620. SIAM, 2008.

## A Dealing with Super-Polynomially Many Flaws

In this section we discuss how one can deal with problems where the number of flaws is super-polynomial in the natural size of the problem using commutative algorithms.

In such a setting, there are two issues to be resolved. The first issue is that one should be able to show that the expected number of steps until convergence is polynomial, and thus, much less than  $\Theta(|F|)$ . The second issue is that one should have an efficient procedure for finding a flaw that is present in the current state, or decide that no such flaw exists.

**Polynomial-Time Convergence.** As far as the issue of polynomial-time convergence is concerned, there are at least three approaches one can follow.

A first approach is to start the algorithm at a state  $\sigma_1$  in which the set of flaws present is of polynomial size, and then employ the main results from [2, 1, 30] which guarantee that the algorithm will terminate after  $O\left(|U(\sigma_1)| + \max_{\sigma \in \Omega} \log_2 \frac{1}{\mu(\sigma)}\right)$  steps with high probability. This approach does not require the algorithm to be commutative, but it does require that the LLL condition is satisfied with a slack in order to establish quick termination.

A second approach, which was first applied in the context of the Moser-Tardos algorithm by Haeupler, Saha and Srinivsan [21], is to find a *core* set of flaws of polynomial size and apply a modified version of the algorithm that effectively ignores any non-core flaw. The hope is that non-core flaws will never occur during the execution of this modified algorithm. Extended to our setting, one uses the following theorem which is a straightforward corollary of Theorem 3.2.

**Theorem A.1.** *Assume that  $\mathcal{A} = (F, \sim, \rho)$  is commutative. Let  $I \subseteq [m]$  be a set of indices that corresponds to a core subset of  $F$  and assume there exist positive real numbers  $\{\psi_i\}_{i=1}^m$  such that for every  $i \in [m]$*

$$\gamma(f_i) \sum_{S \in \text{Ind}(\Gamma(i) \cap I)} \prod_{j \in S} \psi_j \leq \psi_i .$$

*Then there exists a modification of  $\mathcal{A}$  that terminates in an expected number of  $O\left(\lambda_{\text{init}} \sum_{i \in I} \psi_i\right)$  steps and outputs a flawless element with probability at least  $1 - \sum_{i \in [m] \setminus I} \lambda_{\text{init}} \gamma(f_i) \sum_{S \in \text{Ind}(\Gamma(i) \cap I)} \prod_{j \in S} \psi_j$ .*

Finally, a third approach is to show that the causality graph can be decomposed into a set of cliques of polynomial size and then apply a result of [21] which states that, in this case, the running time of the algorithm is polynomial (roughly quadratic) in the size of the decomposition. To be more precise, we note that in [21] the latter result is shown for the Moser-Tardos algorithm in the variable setting and assuming the General LLL condition (1), where the clique decomposition considered is induced by the random variables that form the probability space (one clique per variable). However, the proof for the general case is identical. Using Theorem 3.2 and recalling Remark 2.1 we can extend this result to our setting to get the following theorem.

**Theorem A.2.** *Let  $\mathcal{A} = (D, F, \sim, \rho)$  be a commutative algorithm such that the causality graph induced by  $\sim$  can be partitioned into  $n$  cliques, with potentially further edges between them. Assume there exist real numbers  $\{x_i\}_{i=1}^m$  in  $(0, 1)$  such that for every  $i \in [m]$  we have that*

$$\gamma(f_i) \leq x_i \prod_{j \in \Gamma(i)} (1 - x_j) ,$$

*and let  $\delta := \min_{i \in [m]} x_i \prod_{j \in \Gamma(i)} (1 - x_j)$ . Then, the expected number of steps performed by  $\mathcal{A}$  is at most  $t = O\left(\lambda_{\text{init}} \cdot \frac{n}{\epsilon} \log \frac{n \log(1/\delta)}{\epsilon}\right)$ , and for any parameter  $\eta$ ,  $\mathcal{A}$  terminates within  $\eta t$  resamplings with probability  $1 - e^{-\eta}$ .*

**Remark A.1.** In [21] it is argued that in the vast majority of applications  $\delta = O(n \log n)$  and in many cases even linear in  $n$ .

Following Theorem 3.2, the proof of Theorem A.2 is identical to the analogous result of Haeupler, Saha and Srinivasan [21] for the Moser-Tardos algorithm and hence we omit it.

**Fast Search for Flaws.** Searching for occurring flaws efficiently can be a major obstacle in getting polynomial time algorithms, even in the case where convergence is guaranteed after a polynomial number of steps. Again, there is more than one approach one can follow to deal with this issue.

A first approach was introduced in [21] where it is shown that Theorems A.1 and A.2, in the context of the variable setting, can be combined into a single theorem that guarantees the existence of a Monte Carlo algorithm which runs in polynomial time, even in the presence of super-polynomially many flaws. The theorem assumes the existence of a polynomial size decomposition of the causality graph into cliques and, moreover, that the LLL conditions hold with an exponential slack. Using Theorem 3.2, we can extend this result in a straightforward way to our setting to get:

**Theorem A.3.** Let  $\mathcal{A} = (D, F, \sim, \rho)$  be a commutative algorithm such that the causality graph induced by  $\sim$  can be partitioned into  $n$ , with potentially further edges between them. Assume there exist real numbers  $\{x_i\}_{i=1}^m$  and  $\epsilon$  in  $(0, 1)$  be such that for every  $i \in [m]$  we have that

$$\gamma(f_i)^{1-\epsilon} \leq x_i \prod_{j \in \Gamma(i)} (1 - x_j) .$$

If we furthermore have that  $\log 1/\delta \leq \text{poly}(n)$ , where  $\delta = \min_{i \in [m]} x_i \prod_{j \in \Gamma(i)} (1 - x_j)$ , then for every  $\gamma \geq \frac{1}{\text{poly}(n)}$  the set  $\{i \in [m] \text{ s.t. } \gamma(f_i) \geq \gamma\}$  has size at most  $\text{poly}(n)$ . There also exists a Monte Carlo algorithm that terminates after  $O(\lambda_{\text{init}} \frac{n}{\epsilon} \log \frac{n}{\epsilon^2})$  steps and returns a perfect object with probability at least  $1 - n^{-c}$ , where  $c$  is any desired constant.

In a follow-up work [26], Harris and Srinivasan describe a general technique that yields efficient procedures for searching for flaws. The main building blocks of their technique is a “witness tree lemma for internal states” and problem-specific, possibly randomized, data-structures that contain the flaws that are present in each state. We refer the reader to [26] for more details, but we note that combining the proof of [26] with the proof of Theorems 4.1 and 3.2, one can show that the “witness tree lemma for internal states” holds for commutative algorithms.

## B Proofs Omitted from Section 4

### B.1 Proof of Lemma 4.8

We show the following more general lemma. The claim follows by applying this general lemma with  $\text{List}(j) = \text{Ind}(\Gamma(i))$  for every  $j \in [m]$ .

**Lemma B.1.** Assume that for every  $i \in [m]$  we are given a set  $\text{List}(i) \subseteq 2^{[m]}$  and there exist positive numbers  $\{\psi_i\}_{i=1}^m$  such that for each  $i$ :

$$\frac{\gamma(f_i)}{\psi_i} \sum_{S \in \text{List}(i)} \prod_{j \in S} \psi_j \leq 1 .$$

Let  $\mathcal{L}_i$  be the set of trees whose root is labelled by  $i$  and such that the set of labels of every node  $v$  with label  $(v) = j$  is in  $\text{List}(j)$ , for every  $j \in [m]$ . Then:

$$\sum_{\tau \in \mathcal{L}_i} \prod_{v \in V(\tau)} \gamma(f_{(v)}) \leq \psi_i .$$

*Proof of Lemma B.1.* To proceed, we use ideas from [36, 39]. Specifically, we introduce a branching process that produces only trees in  $\mathcal{L}_i$  and bound  $\sum_{\tau \in \mathcal{L}_i} \prod_{v \in V(\tau)} \gamma(f_{(v)})$  by analyzing it.

In particular, we start with a single node labelled by  $i$ . In each subsequent round each leaf  $u$  “gives birth” to a set of nodes whose set of (distinct) labels is a set  $S \in \text{List}((u))$  with probability proportional to  $\prod_{j \in S} \psi_j$ . It is not hard to see that this process creates every tree in  $\mathcal{L}_i$  with positive probability. To express the exact probability received by each  $S \subseteq [m]$  we define

$$Q(S) := \prod_{j \in S} \psi_j \tag{13}$$

and let  $Z_\ell = \sum_{S \in \text{List}((u))} Q(S)$ . Clearly, each  $S \in \text{List}((u))$  receives probability equal to  $\frac{Q(S)}{Z_\ell}$ . We now show the following lemma.

**Proposition B.2.** *The branching process described above produces every tree  $\tau \in \mathcal{L}_i$  with probability*

$$p_\tau = \frac{1}{\psi_i} \prod_{v \in V(\tau)} \frac{\psi_{(v)}}{\sum_{S \in \text{List}((v))} \prod_{j \in S} \psi_j} .$$

*Proof.* For each tree  $\tau \in \mathcal{L}_i$  and each node  $v$  of  $\tau$ , let  $N(v)$  denote the set of labels of its children. Then:

$$\begin{aligned} p_\tau &= \prod_{v \in V(\tau)} \frac{Q(N(v))}{\sum_{S \in \text{List}((v))} Q(S)} \\ &= \frac{1}{\psi_i} \prod_{v \in V(\tau)} \frac{\psi_{(v)}}{\sum_{S \in \text{List}((v))} Q(S)} . \end{aligned}$$

□

Notice now that

$$\sum_{\tau \in \mathcal{L}_i} \prod_{v \in V(\tau)} \gamma(f_{[v]}) \leq \sum_{\tau \in \mathcal{L}_i} \prod_{v \in V(\tau)} \frac{\psi_{(v)}}{\sum_{S \in \text{List}((v))} \prod_{j \in S} \psi_j} \tag{14}$$

$$\begin{aligned} &= \psi_i \sum_{\tau \in \mathcal{L}_i} p_\tau \\ &= \psi_i , \end{aligned} \tag{15}$$

where (14) follows by the hypothesis of Lemma B.1 while (15) by Proposition B.2.

□

## B.2 Proof of Lemma 4.15

Let  $\mathbf{u}_1, \dots, \mathbf{u}_m$  be the named indices of flaws of  $\mathbf{Q}(\Sigma)$  that occur in  $U$  listed in the order of their appearance in  $U$ . We claim that  $\mathbf{u}_1, \dots, \mathbf{u}_m$  is a prefix of  $U$ . To see this, assume for the sake of contradiction that  $U = \dots \mathbf{w} \mathbf{u}_i \dots$  where  $\mathbf{w} \notin \mathbf{Q}(\Sigma)$  and  $\mathbf{u}_i \in \mathbf{Q}(\Sigma)$ . Thus,  $(\mathbf{w}, \mathbf{u}_i) \notin \mathbf{E}(\Sigma)$  and so  $w \approx u_i$ . Therefore,  $(\mathbf{w}, \mathbf{u}_i)$  is a swappable pair in  $U$ , which is a contradiction.

Note that the latter observation implies that  $W(\Sigma) = (A, B, C)$  where  $B = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m)$ . It remains to show that  $B$  is a subsequence of  $Q_\pi(\Sigma)$ . In particular, it suffices to show that for any  $i \in [m-1]$  the relative order of  $\mathbf{u}_i$  and  $\mathbf{u}_{i+1}$  in  $B$  is the same as in  $Q_\pi(\Sigma)$ . Assume the opposite, i.e.,  $Q_\pi(\Sigma) = \dots \mathbf{u}_{i+1} \dots \mathbf{u}_i \dots$ . It has to be that  $\mathbf{u}_i \sim \mathbf{u}_{i+1}$ . For otherwise,  $(\mathbf{u}_i, \mathbf{u}_{i+1})$  would be a swappable pair, contradicting the assumption. This means that  $(\mathbf{u}_i, \mathbf{u}_{i+1}) \in \mathbf{E}(\Sigma)$ , implying that  $d(\mathbf{u}_i) > d(\mathbf{u}_{i+1})$ . Recalling the definition of  $Q_\pi(\Sigma)$  we see that  $\mathbf{u}_i$  should be to the left of  $\mathbf{u}_{i+1}$  in  $Q_\pi(\Sigma)$ , a contradiction.

### B.3 Proof of Lemma 4.16

First we prove that the mapping is injective. In particular, assume that two distinct trajectories  $\Sigma_1, \Sigma_2 \in \mathcal{X}_p[W]$  are transformed to the same trajectory  $\Sigma \in \mathcal{X}_{p+1}[W]$ . At least one of  $\Sigma_1, \Sigma_2$  must have changed. Without loss of generality, assume  $\Sigma \neq \Sigma_1$ . The latter implies that  $W(\Sigma_1) = (A, \mathbf{w}, \mathbf{y}, B)$  with  $w \approx y$  and  $\Sigma_1$  is transformed to a trajectory  $\Sigma$  with  $W(\Sigma) = (A, \mathbf{y}, \mathbf{w}, B)$ . Notice that it cannot be the case that  $\Sigma_2 = \Sigma$  since then  $\Sigma_1$  and  $\Sigma$  would both be in  $\mathcal{X}_p[W]$  without following the same deterministic flaw choice strategy, a contradiction. Thus,  $W(\Sigma_2) = (A, \mathbf{w}, \mathbf{y}, B)$ , and  $\Sigma$  was obtained from  $\Sigma_2$  by swapping  $\mathbf{w}$  and  $\mathbf{y}$ . Recalling that the Swap operation is an injection, we get that  $\Sigma_1 = \Sigma_2$ .

We will assume that  $\mathcal{X}_{p+1}[W]$  is not valid and reach a contradiction. The latter assumption implies that there should be trajectories  $\Sigma, \Sigma' \in \mathcal{X}_{p+1}[W]$  such that

$$\begin{aligned} W(\Sigma) &= (\mathbf{w}_1, \dots, \mathbf{w}_\ell, \mathbf{w}, \mathbf{y}, \dots) \\ W(\Sigma') &= (\mathbf{w}_1, \dots, \mathbf{w}_\ell, \bar{\mathbf{w}}, \bar{\mathbf{y}}, \dots) , \end{aligned}$$

with  $\mathbf{w} \neq \bar{\mathbf{w}}$  and the states in  $\Sigma$  to the left of  $\mathbf{w}$  match the corresponding states in  $\Sigma'$  to the left of  $\bar{\mathbf{w}}$ . Here it is assumed that some of  $\mathbf{w}, \mathbf{y}, \bar{\mathbf{w}}, \bar{\mathbf{y}}$  may equal  $\emptyset$ , which means they don't exist. It is also assumed that  $\mathbf{w} = \emptyset$  also implies that  $\mathbf{y} = \emptyset$  and similarly for  $\bar{\mathbf{w}}$  and  $\bar{\mathbf{y}}$ .

Let  $\Sigma_1$  and  $\Sigma_2$  be respectively the trajectories in  $\mathcal{X}_p[W]$  that were transformed to  $\Sigma$  and  $\Sigma'$ . Since  $\mathcal{X}_p[W]$  is valid, at least one of them must have changed. Assume, without loss of generality,  $\Sigma \neq \Sigma_1$ . We know that (i)  $\Sigma_1$  and  $\Sigma_2$  follow the same deterministic flaw choice strategy, and they are not proper prefixes of each other, as well as, that (ii) named flaws indices  $\mathbf{w}_1, \dots, \mathbf{w}_\ell, \mathbf{w}$  are distinct.

Some other useful facts to have in mind (and which we will implicitly use) are that the first  $\ell + 1$  states of  $\Sigma$  match those of  $\Sigma'$ , and also that swapping adjacent indices of flaws only affects the state between them in a deterministic way. Now there are four cases:

- (a) The swapped pair in  $\Sigma$  was  $(\mathbf{w}_i, \mathbf{w}_{i+1})$  for  $i \in [\ell - 1]$ . Thus,

$$W(\Sigma_1) = (\mathbf{w}_1, \dots, \mathbf{w}_{i+1}, \mathbf{w}_i, \dots, \mathbf{w}_\ell, \mathbf{w}, \mathbf{y}, \dots) .$$

Using (i) and (ii), we conclude that  $\Sigma' \neq \Sigma_2$  and, thus,

$$W(\Sigma_2) = (\mathbf{w}_1, \dots, \mathbf{w}_{i+1}, \mathbf{w}_i, \dots, \mathbf{w}_\ell, \bar{\mathbf{w}}, \dots, \mathbf{y}, \dots) .$$

To see this, recall that swaps are applied at the same positions for trajectories in  $\mathcal{X}_p$ . Condition (i) now implies that  $\mathbf{w} = \bar{\mathbf{w}}$ .

- (b) The swapped pair in  $\Sigma$  was  $(\mathbf{w}_\ell, \mathbf{w})$ , and so  $W(\Sigma) = (\mathbf{w}_1, \dots, \mathbf{y}, \mathbf{w}_\ell, \mathbf{w}, \dots)$ . Thus,  $\Sigma' \neq \Sigma_2$  by (i) and (ii). Now condition (i) implies that  $\mathbf{w} = \bar{\mathbf{w}}$ .

- (c) The swapped pair in  $\Sigma$  was  $(\mathbf{w}, \mathbf{y})$ , and so  $W(\Sigma_1) = (\mathbf{w}_1, \dots, \mathbf{w}_\ell, \mathbf{y}, \mathbf{w}, \dots)$  and  $\mathbf{w} \in \mathbf{Q}(\Sigma_1)$ . We now apply Lemma 4.15 to  $\Sigma_1$  and we notice that since there are no swappable pairs in  $\Sigma_1$  to the right of  $(\mathbf{y}, \mathbf{w})$  it should be that  $W(\Sigma_1) = (\mathbf{w}_1, \dots, \mathbf{w}_\ell, B, C)$  where  $B$  is a subsequence of  $Q_\pi(\Sigma_1)$  and  $C$

does not contain named indices of flaws from  $Q_\pi(\Sigma)$ . Observe that  $B$  should start with  $\mathbf{w}$ . Using (i) we get that  $W(\Sigma_2) = (\mathbf{w}_1, \dots, \mathbf{w}_\ell, \mathbf{y}, \dots)$  and again by Lemma 4.15:  $W(\Sigma_2) = (\mathbf{w}_1, \dots, \mathbf{w}_\ell, \mathbf{y}, \overline{B}, \overline{C})$  where  $\overline{B}$  is a subsequence of  $Q_\pi(\Sigma_2)$  and  $\overline{C}$  does not contain named indices of flaws from  $\mathbf{Q}(\Sigma_2)$ .

By the definition of  $\mathcal{X}_p[W]$  we know that  $Q_\pi(\Sigma_1) = Q_\pi(\Sigma_2) = W$ . The latter fact, along with the forms of  $\Sigma_1, \Sigma_2$  imply that  $B$  should be a permutation of  $\overline{B}$ . Notice though that  $B$  and  $\overline{B}$  are subsequences of  $W$  and, furthermore, all elements of  $W$  are distinct. Therefore, it has to be that  $B = \overline{B}$ .

The latter observation implies that  $W(\Sigma_2) = (\mathbf{w}_1, \dots, \mathbf{w}_\ell, \mathbf{y}, \mathbf{w}, \dots)$ . Now since  $(\mathbf{y}, \mathbf{w})$  is a swappable pair in  $\Sigma_1$  it should also be a swappable pair in  $\Sigma_2$ . Thus,  $\Sigma' = (\mathbf{w}_1, \dots, \mathbf{w}_\ell, \mathbf{w}, \mathbf{y}, \dots)$ . This means that  $\mathbf{w} = \overline{\mathbf{w}}$ .

(d) The swapped pair in  $\Sigma$  was to the right of  $(\mathbf{w}, \mathbf{y})$ . In this case condition (i) implies that  $\mathbf{w} = \overline{\mathbf{w}}$ .

## C Proofs Omitted from Section 5

### C.1 Proof of Theorem 5.2

To lighten the notation, let  $u := \lambda_{\text{init}} \sum_{S \in \text{Ind}([m])} \prod_{j \in S} \psi_j$ . For each  $\sigma \in \Omega$ , define a flaw  $f_\sigma = \{\sigma\}$  and consider the extended algorithm that addresses it by sampling from  $\mu$ , as well as the extended causality graph that connects  $f_\sigma$  with every flaw in  $F$ . Clearly, we have that  $\gamma(f_\sigma) = \mu(\sigma)$ . Moreover, if the original algorithm is commutative, so is the extended one since the commutativity condition is trivially true for flaws  $\{f_\sigma\}_{\sigma \in \Omega}$ . Observe now that for every  $\sigma \in \Omega$ , Theorem 3.2 yields  $\nu(\sigma) \leq \Pr[\sigma] \leq u \cdot \mu(\sigma)$ . Thus:

$$H_\rho[\nu] = \frac{1}{1-\rho} \ln \sum_{\sigma \in \Omega} \nu(\sigma)^\rho \geq \frac{1}{1-\rho} \ln \sum_{\sigma \in \Omega} (u\mu(\sigma))^\rho = \frac{1}{1-\rho} \ln \sum_{\sigma \in \Omega} \mu(\sigma)^\rho - \frac{\rho}{\rho-1} \ln u ,$$

concluding the proof.

### C.2 Proof of Theorem 5.3

For each flaw  $f_i$  we define a Bernoulli variable  $Y_i$  with probability of success  $p_i = \min \left\{ 1, \frac{\psi_i}{\zeta_i \gamma(f_i)} \right\}$ . The sequence  $\{Y_i\}_{i=1}^m$  and  $\Omega$  induce a new space  $\Omega' = \Omega \times \{0, 1\}^m$  which can be thought as a “labelled” version of  $\Omega$ , where each state  $\sigma$  is labelled with a binary vector of length  $m$  whose  $i$ -th bit describes the state of  $Y_i$ . Similarly, measure  $\mu$  and  $\{Y_i\}_{i=1}^m$  induce a measure  $\mu'$  over  $\Omega'$ .

In this new state space we introduce a new family of flaws  $F' = \{f'_1, f'_2, \dots, f'_m\}$ , where  $f'_i$  is defined as the subset of  $\Omega'$  where  $f_i$  is present and  $Y_i = 1$ . Consider now the algorithm  $\mathcal{A}'$  that is induced by  $\mathcal{A}$  as follows: Each time we want to address flaw  $f'_i$  we move in  $\Omega$  by invoking  $\mathcal{A}$  to address  $f_i$  and also take a sample from  $Y_i$  to update the value of the  $i$ -th entry of the label-vector.

It is not hard to verify that (i) the charge of each flaw  $f'_i$  is  $\gamma(f'_i) = \gamma(f_i)p_i$ ; (ii) any causality graph for  $(\Omega, F, \mathcal{A})$  is also a causality graph for  $(\Omega', F', \mathcal{A}')$  (and, in particular, so is the one induced by  $\sim$ ); (iii) if  $\mathcal{A}$  is commutative then so is  $\mathcal{A}'$ ; and that (iv) the cluster expansion condition with respect to the causality graph induced by  $\sim$  is satisfied.

To conclude the proof, consider a flaw  $f'_i$  and notice that in order for  $f_i$  to be present in the output of  $\mathcal{A}'$  it has to be the case that  $Y_i = 0$ . Notice now that Theorem 3.2 implies:

$$\nu(f_i \cap Y_i = 0) \leq (1 - p_i)\gamma(f_i)\zeta_i = \max \{0, \gamma(f_i)\zeta_i - \psi_i\} .$$

## D Proofs Omitted from Section 6

### D.1 Proof of Lemma 6.5

We will use the following LLL condition which is obtained from (1) by setting  $\psi_i = \frac{2\mu(A_i)}{1-2\mu(A_i)}$ . It's proof can be found in [34].

**Proposition D.1.** *Let  $(\Omega, \mu)$  be an arbitrary probability space and let  $\mathcal{A} = \{A_1, \dots, A_m\}$  be a set of (bad) events. For each  $i \in [m]$  let  $D(i) \subseteq ([m] \setminus \{i\})$  be such that  $\mu(A_i \mid \cap_{j \in S} \overline{A_j}) = \mu(A_i)$  for every  $S \subseteq (D(i) \cup \{i\})$ . If*

$$\sum_{j \in D(i) \cup \{i\}} \mu(f_j) < \frac{1}{4} \text{ for each } i \in [m] \text{ ,}$$

*then the probability that none of the events in  $\mathcal{A}$  occurs is strictly positive.*

Let  $\mu$  be the probability distribution induced by giving each Blank vertex  $v$  a color from  $L_v(\sigma) \setminus \text{Blank}$  uniformly at random. For any edge  $e$  and color  $c \in \cap_{u \in e} L_u(\sigma) \setminus \text{Blank}$  we define  $A_{e,c}$  to be the event that all vertices of  $e$  receive  $c$ . We also define  $\text{Blank}(e)$  to be the set of vertices of  $e$  that are Blank in  $\sigma$ . Observe now that

$$\mu(A_{e,c}) \leq \frac{1}{\prod_{v \in \text{Blank}(e)} (|L_v(\sigma)| - 1)} \text{ .}$$

Furthermore,  $A_{e,c}$  is mutually independent of all events with which it does not share a vertex. The lemma follows from Proposition D.1 (and can be made constructive using the Moser-Tardos algorithm) as flaws  $B_v, Z_v$  are not present for every vertex  $v \in V$  and so

$$\begin{aligned} \sum_{v \in \text{Blank}(e)} \sum_{c \in L_v(\sigma) \setminus \text{Blank}} \sum_{u \in T_{v,c}(\sigma)} \mu(A_{\{u,v\},c}) &= \sum_{v \in \text{Blank}(e)} \sum_{c \in L_v(\sigma) \setminus \text{Blank}} \sum_{u \in T_{v,c}(\sigma)} \frac{1}{(|L_v(\sigma)| - 1)(|L_u(\sigma)| - 1)} \\ &\leq 2 \max_{v \in \text{Blank}(e)} \frac{1}{(|L_v(\sigma)| - 1)(L - 1)} \sum_{c \in L_v(\sigma) \setminus \text{Blank}} |T_{v,c}(\sigma)| \\ &\leq \frac{2}{10} \max_{v \in \text{Blank}(e)} \frac{L \cdot |L_v(\sigma)|}{(L - 1) \cdot (|L_v(\sigma)| - 1)} \\ &\leq \frac{1}{5} \left( \frac{L}{L - 1} \right)^2 < \frac{1}{4} \text{ ,} \end{aligned}$$

for large enough  $\Delta$ , concluding the proof.

### D.2 Proof of Lemma 6.8

Recall the description of  $\mathcal{A}_2$  from the proof of Lemma 6.5.

First, we show that  $\mathcal{A}_2$  is able to output at least  $Nq^{-(1-\alpha)^n}$  list-colorings with positive probability. Let  $\Omega_{\mathcal{A}_1}^*$  denote the set of flawless partial list-colorings algorithm  $\mathcal{A}_1$  can output with positive probability, and note that, according to our assumption,  $|\Omega_{\mathcal{A}_1}^*| = N$ . To see the idea behind the bound, observe that given two colorings  $\sigma_1, \sigma_2 \in \Omega_{\mathcal{A}_1}^*$ , applying  $\mathcal{A}_1$  to each one of them is guaranteed to result in different full list-colorings unless there is a way to start from  $\sigma_1$  (respectively, from  $\sigma_2$ ) and assign colors to Blank vertices so that we reach  $\sigma_2$  (respectively, to  $\sigma_1$ ). In this bad case we write  $\sigma_1 \bowtie \sigma_2$ . Consider now the graph  $H$  over  $\Omega_{\mathcal{A}_1}^*$  in which two colorings  $\sigma_1, \sigma_2$  are adjacent iff  $\sigma_1 \bowtie \sigma_2$ , and observe that the size of any independent

set of  $H$  is a lower bound on the number of list-colorings  $\mathcal{A}_2$  can output. Since we have assumed that every coloring in  $\Omega_{\mathcal{A}_1}^*$  has at least  $\alpha n$  vertices colored, we see that the maximum degree of  $H$  is at most  $D := q^{(1-\alpha)n} - 1$  and, therefore, there exists an independent set of size at least  $\frac{|\Omega_{\mathcal{A}_1}^*|}{D+1} = Nq^{-(1-\alpha)n}$ , concluding the proof of the first part of Lemma 6.5.

Second, we show that  $\mathcal{A}_2$  is able to output at least  $\left(\frac{8L}{11}\right)^{(1-\alpha)n}$  list colorings with positive probability. To do that, we will need the following theorem regarding the output distribution of the Moser-Tardos algorithm that was proved in [26], and which we rephrase here to fit our needs.

**Theorem D.2** ([26]). *Consider a constraint satisfaction problem on a set of variables  $\mathcal{V}$  and set of constraints  $\mathcal{C}$ . Assume we have a flaw  $f_c$  for each constraint  $c$ , comprising the set of states that violate  $c$ . We are also given an undirected causality graph such that two constraints are connected with an edge iff they share variables. For each constraint  $c$  define*

$$y_c = (1 + \psi_c)^{\frac{1}{|\text{var}(c)|}} - 1 ,$$

where  $\text{var}(c)$  denotes the set of variables that correspond to constraint  $c$ . Then:

$$\sum_{S \in \text{Ind}(\mathcal{C})} \prod_{c \in S} \psi_c \leq \prod_{v \in \mathcal{V}} \left( 1 + \sum_{\substack{c \in \mathcal{C} \\ v \in \text{var}(c)}} y_c \right) ,$$

where  $\text{Ind}(\mathcal{C})$  denotes the set of independent sets of  $\mathcal{C}$ .

Observe that the hypothesis implies that  $\mathcal{A}_1$  can output a flawless partial coloring  $\sigma$  where exactly  $\alpha n$  vertices are colored with positive probability. We apply  $\mathcal{A}_2$  to  $\sigma$ , which recall that is an instantiation of the Moser-Tardos algorithm using the uniform measure  $\mu$  over the the cartesian product,  $\Omega'$ , of the lists of non-Blank available colors of the Blank vertices of  $\sigma$ , and where we have a bad event  $A_{e,c}$  for any edge  $e$  and color  $c \in \bigcap_{u \in e} L_u(\sigma) \setminus \{\text{Blank}\}$ . Recall further that the general (and, thus, also the cluster expansion) LLL condition (??) is satisfied with  $\psi_{e,c} = \frac{2\mu(A_{e,c})}{1-2\mu(A_{e,c})}$ . Thus, we can combine Theorem 5.2 and Theorem D.2 to get that  $\mathcal{A}_2$  can output at least

$$|\Omega'| \left( \prod_{\substack{v \in V \\ \sigma(v) = \text{Blank}}} \left( 1 + \sum_{c \in L_v(\sigma) \setminus \text{Blank}} \sum_{u \in T_{v,c}(\sigma)} y_{\{u,v\},c} \right) \right)^{-1} \geq \frac{L^{(1-\alpha)n}}{\left(1 + \frac{3}{8}\right)^{(1-\alpha)n}} \geq \left(\frac{8L}{11}\right)^{(1-\alpha)n} ,$$

list-colorings with positive probability. To see this, notice that for any  $v \in V$  that is Blank in  $\sigma$ , and sufficiently large  $\Delta$ ,

$$\begin{aligned} \sum_{c \in L_v(\sigma) \setminus \text{Blank}} \sum_{u \in T_{v,c}(\sigma)} y_{\{u,v\},c} &= \sum_{c \in L_v(\sigma) \setminus \text{Blank}} \sum_{u \in T_{v,c}(\sigma)} \left( \sqrt{1 + \frac{2\mu(A_{\{u,v\},c})}{1-2\mu(A_{\{u,v\},c})}} - 1 \right) \\ &\leq \sum_{c \in L_v(\sigma) \setminus \text{Blank}} \sum_{u \in T_{v,c}(\sigma)} 3\mu(A_{\{u,v\},c}) \\ &\leq 3 \cdot \left( \frac{1}{2} \cdot \frac{1}{4} \right) = \frac{3}{8} , \end{aligned} \tag{16}$$

where to obtain (16) we perform identical calculations to the ones in Lemma 6.5.

### D.3 Proofs of Theorems 6.10 and 6.11

#### D.3.1 The Clique Lovász Local Lemma

We first state the Clique Lovász Local Lemma (reformulated to fit our setting) assuming as input a commutative algorithm  $(F, C, \rho)$ , where  $C$  is a causality graph.

**Theorem D.3** (The Clique Lovász Local Lemma). *Let  $\{K_1, K_2, \dots, K_n\}$  be a set of cliques in  $C$  covering all the edges (not necessarily disjointly). If there exists a set of vectors  $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$  from  $(0, 1)^m$  such that the following condition are satisfied:*

- for each  $v \in [n]$ :  $\sum_{i \in K_v} x_{i,v} < 1$ ;
- for each  $i \in [m]$ ,  $\forall v$  such that  $i \in K_v$ :

$$\gamma(f_i) \leq x_{i,v} \prod_{u \neq v: K_u \ni i} \left(1 - \sum_{j \in K_u \setminus \{i\}} x_{j,u}\right)$$

then:

1.  $\mu \left( \bigcap_{i \in [m]} \overline{f_i} \right) \geq \prod_{v \in [m]} \left(1 - \sum_{i \in K_v} x_{i,v}\right) > 0$
2. The algorithm terminates after an expected number of at most

$$\sum_{i \in [m]} \min_{v: K_v \ni i} \frac{x_{i,v}}{1 - \sum_{j \in K_v \setminus \{i\}} x_{j,v}},$$

steps.

We note that in [29] the authors first prove the first part of their theorem, which implies the *existence* of perfect objects, and then they invoke the results of [28] which imply that the Moser-Tardos algorithm converges under the Shearer's condition. In particular, they use the following fact, which we will also find useful in our applications:

$$\frac{q_{\{i\}}}{q_{\emptyset}} \leq \min_{v: K_v \ni i} \frac{x_{i,v}}{1 - \sum_{j \in K_v \setminus \{i\}} x_{j,v}}, \text{ for every } i \in [m]. \quad (17)$$

To prove Theorem D.3 in our setting we can follow the same strategy, invoking Theorem 3.2 (in the Shearer's regime) instead of the main result of [28]. In fact, the proof of the first part of Theorem D.3 is identical to the one of [29] assuming that the input algorithm is a resampling oracle with respect to  $\mu$  for each flaw  $f_i$ . In the general case, some extra work is required. Since in our application we will be using the Moser-Tardos algorithm (which is a resampling oracle for every flaw), and in the interest of brevity, we omit it.

Finally, we note that the authors provide a canonical way of decomposing the causality graph and applying the Clique Lovász Local Lemma in the variable setting of Moser and Tardos. Specifically, recall that in the variable setting the family of bad events is determined by a set of independent discrete random variables  $\{v_1, \dots, v_m\}$ , and that two events are adjacent in the dependency graph whenever they share a variable. Thus, each variable  $v$  forms a clique  $K_v$  in the dependency graph consisting of the events that are dependent on this variable.

### D.3.2 Finding Acyclic Edge Colorings

We now recall the proof of [29] for Acyclic Edge Coloring. The proof is the same as the one in [6, 33] and the improvement comes from the use of Clique LLL instead of condition (1).

Given  $q$  colors and a graph  $G$  with maximum degree  $\Delta$  let  $\Omega$  be the set of all edge  $q$ -colorings of  $G$ . We identify the two following types of flaws:

1. For a path  $P$  of length 2 let  $f_P$  be the set of states in  $\Omega$  in which  $P$  is monochromatic.
2. For a cycle  $C$  of even length let  $f_C$  comprise the set of states in  $\Omega$  in which  $C$  is bicolored.

Clearly a flawless element of  $\Omega$  is an acyclic edge coloring of  $G$ . Our algorithm is the MT algorithm (the variables that correspond to each event are the edges of the path/cycle),  $\mu$  is the uniform measure and  $\theta = \mu$ . Therefore:

$$\begin{aligned}\gamma(f_P) &= \mu(f_P) = \frac{1}{q} , \\ \gamma(f_C) &= \mu(f_C) \leq \frac{1}{q^{|C|-2}} .\end{aligned}$$

Furthermore, two flaws are connected in the causality graph iff they share an edge. We now follow the canonical way of decomposing the causality graph into cliques by having one clique  $K_e$  for each edge  $e$  of  $G$ . Moreover:

- For each path  $P \ni e$  of length 2 we set  $x_{f_P,e} = x_{P,e} = \frac{c}{1+\epsilon} \frac{1}{2\Delta-2}$
- For a cycle  $C$  of even length we set  $x_{f_C,e} = x_{C,e} = \frac{c}{(1+\epsilon)^{|C|/2}} \frac{1}{(\Delta-1)^{|C|-2}}$

for some positive  $c, \epsilon$  to be determined later.

Observe that the number of cycles of length  $2\ell$ , where  $\ell \geq 2$ , that contain any given edge  $e$  is at most  $(\Delta - 1)^{2\ell-2}$ , while the number of paths of length 2 that contain  $e$  is at most  $2\Delta - 2$ . Thus, it suffices to show that for each edge  $e$  and each path of length  $P$  and cycle  $C$  of length  $2\ell$  that contain  $e$  we have that:

$$\begin{aligned}\gamma(f_P) &\leq x_{P,e} \prod_{e' \in P \setminus \{e\}} \left( 1 - \sum_{j \in K_e} x_{j,e'} \right) , \\ \gamma(f_C) &\leq x_{C,e} \prod_{e' \in C \setminus \{e\}} \left( 1 - \sum_{j \in K_e} x_{j,e'} \right) .\end{aligned}$$

which is

$$\begin{aligned}\frac{1}{q} &\leq \frac{c}{1+\epsilon} \frac{1}{2\Delta-2} \left( 1 - c \sum_{j=1}^{\infty} (1+\epsilon)^{-j} \right) , \\ \frac{1}{q^{2\ell-2}} &\leq \frac{c}{(1+\epsilon)^\ell} \frac{1}{(\Delta-1)^{2\ell-2}} \left( 1 - c \sum_{j=1}^{\infty} (1+\epsilon)^{-j} \right)^{2\ell-1} .\end{aligned}$$

The latter imply that for the Moser-Tardos algorithm to converge it suffices to have  $c < \epsilon$  and also

$$\begin{aligned} \frac{q}{\Delta - 1} &\geq \max \left\{ \frac{2}{c} (1 + \epsilon) \frac{\epsilon}{\epsilon - c}, \max_{\ell \geq 2} \left\{ (1 + \epsilon)^{\frac{\ell}{2\ell-2}} \frac{1}{c^{\frac{1}{2\ell-2}}} \left( \frac{\epsilon}{\epsilon - c} \right)^{\frac{2\ell-1}{2\ell-2}} \right\} \right\} \\ &\geq \max \left\{ \frac{2}{c} (1 + \epsilon) \frac{\epsilon}{\epsilon - c}, \frac{(1 + \epsilon)}{\sqrt{c}} \left( \frac{\epsilon}{\epsilon - c} \right)^{\frac{3}{2}} \right\} \geq 8.59 . \end{aligned}$$

for  $\epsilon = 2.05869$  and  $c = 0.8282$ .

We now show that the expected running time of our algorithm is polynomial. To do so, we will need to be careful in our flaw choice strategy and also deal with the fact that the number of flaws is super-polynomial.

As far as the flaw choice strategy is concerned, we choose to give priority to flaws of the form  $f_P$ , whose presence in the current state can be verified in polynomial time in the number of edges of  $G$ . This means we only address bichromatic cycles when the underlying edge-coloring is proper. Note that in order to find bichromatic cycles in a properly edge-colored graph we can simply consider each of the  $\binom{q}{2}$  pair of distinct colors and seek cycles in the subgraph of the correspondingly colored edges.

One way to address the fact that the number of flaws is super-polynomial is to invoke Theorem 7 of [28] that states that whenever the LLL condition is satisfied with a multiplicative slack (notice that we have showed that the Clique LLL is satisfied with  $q \geq 8.59$ ) then the expected number of resampling of the Moser-Tardos algorithm is polynomial in the number of the independent random variables (in our case, the edges of  $G$ ) with constant probability. The probability of polynomial convergence can be boosted by repetition. Overall, we have shown the following theorem.

**Theorem D.4.** *Given a graph  $G$  with maximum degree  $\Delta$  and  $q \geq 8.6$  colors there exists an algorithm that outputs an acyclic edge coloring in expected polynomial time.*

### D.3.3 Proof of Theorem 6.10

To prove Theorem 6.10 we need to estimate  $\sum_{S \in \text{Ind}(F)} \prod_{f \in S} \frac{q_{\{f\}}}{q_0}$  per Theorem 5.2 and Remark 5.1 (we slightly abuse the notation and indicate  $\psi$ 's using flaws instead of indices of flaws). To do so, we will use Theorem D.2 where we replace  $\psi_c$  with  $q_{\{c\}}/q_0$ .

Using the result of the previous subsection along with (17) we have that for each path  $P$  of length two and each cycle  $C_\ell$  of length  $\ell$ :

$$\begin{aligned} y_P &\leq \left( 1 + \frac{c}{1 + \epsilon} \frac{1}{2(\Delta - 1)} \frac{\epsilon}{\epsilon - c} \right)^{\frac{1}{2}} - 1 < \left( 1 + \frac{1}{\Delta - 1} \right)^{1/2} - 1 < \frac{1}{2(\Delta - 1)} \\ y_{C_\ell} &\leq \left( 1 + \frac{c}{(1 + \epsilon)^\ell} \frac{1}{(\Delta - 1)^{2\ell-2}} \right)^{\frac{1}{2\ell}} - 1 < \left( 1 + \frac{3^{-\ell}}{(\Delta - 1)^{2\ell-2}} \right)^{\frac{1}{2\ell}} - 1 < \frac{1}{(2(\Delta - 1))^{2\ell-2}} \end{aligned}$$

and, thus, we get:

$$\sum_{S \in \text{Ind}(F)} \prod_{f \in S} \frac{q_{\{f\}}}{q_0} \leq \prod_{e \in E} \left( 1 + 2(\Delta - 1) \frac{1}{2(\Delta - 1)} + \sum_{i=2}^{\infty} (\Delta - 1)^{2i-2} \frac{1}{(2(\Delta - 1))^{2i-2}} \right) < 4^{|E|} .$$

Now Theorem 5.2 and the fact that  $|\Omega| = q^{|E|}$  conclude the proof.

### D.3.4 Proof of Theorem 6.11

For a vertex  $v$  and a function  $W_v$  let  $N(v)$  denote the set of edges adjacent to  $v$  and  $\mathcal{A}_v$  denote the set of possible edge  $q$ -colorings of the edges in  $N(v)$ . For  $\alpha \in \mathcal{A}_v$  let  $E_v(\alpha)$  be the subset of  $\Omega$  whose elements assign  $\alpha$  to the edges in  $N(v)$ . Moreover, consider the resampling probability distributions induced by the Moser-Tardos algorithm for  $E_v(\alpha)$ . Observe that using (17) we get (again, slightly abusing the notation):

$$\begin{aligned} \sum_{S \in \text{Ind}(\Gamma(E_v(\alpha)))} \prod_{f \in S} \frac{q_{\{f\}}}{q_\emptyset} &\leq \prod_{e \in N(v)} \left( 1 + \sum_{P \ni e} x_{P,e} + \sum_{C \ni e} x_{C,e} \right) \\ &\leq \prod_{e \in N(v)} \left( 1 + \left( \frac{c}{1+\epsilon} + \sum_{\ell=2}^{\infty} \frac{c}{(1+\epsilon)^\ell} \right) \frac{\epsilon}{\epsilon-c} \right) \\ &< 1.3^\Delta . \end{aligned}$$

Now applying Theorem 3.2 we get:

$$\begin{aligned} \mathbb{E}[W_v] &= \sum_{\alpha \in \mathcal{A}_v} \Pr[E_v(\alpha)] W_v(\alpha) \\ &< 1.3^\Delta \sum_{\alpha \in \mathcal{A}_v} \mu(E_v(\alpha)) W_v(\alpha) \\ &= 1.3^\Delta \cdot \mathbb{E}_{\phi \sim \mu}[W_v(\phi)] . \end{aligned}$$

Linearity of expectation concludes the proof.