

TIGHT BOUNDS FOR SAMPLING q -COLORINGS VIA COUPLING FROM THE PAST

TIANXING DING HONGYANG LIU YITONG YIN CAN ZHOU

ABSTRACT. The Coupling from the Past (CFTP) paradigm is a canonical method for perfect Sampling. For uniform sampling of proper q -colorings in graphs with maximum degree Δ , the bounding chains of [Huber, STOC'98] provide a systematic framework for efficiently implementing CFTP algorithms within the classical regime $q \geq (1 + o(1))\Delta^2$. This was subsequently improved to $q > 3\Delta$ by [Bhandari and Chakraborty, STOC'20] and to $q \geq (\frac{8}{3} + o(1))\Delta$ by [Jain, Sah, and Sawhney, STOC'21].

In this work, we establish the asymptotically tight threshold for bounding-chain-based CFTP algorithms for graph colorings. We prove a lower bound showing that all such algorithms satisfying the standard contraction property require $q \geq 2.5\Delta$, and we present an efficient CFTP algorithm that achieves this asymptotically optimal threshold $q \geq (2.5 + o(1))\Delta$ via an optimal design of bounding chains.

1. INTRODUCTION

Uniform sampling of graph colorings is a fundamental problem that has attracted extensive attention in combinatorics, probability theory, and theoretical computer science. Given a graph $G = (V, E)$ and an integer $q \geq 1$, a *proper q -coloring* (or simply a *q -coloring*) is a mapping $\omega : V \rightarrow [q]$ such that $\omega(u) \neq \omega(v)$ for every edge $(u, v) \in E$. The set of all proper q -colorings of G is

$$\Omega = \{\omega : V \rightarrow [q] \mid \forall (u, v) \in E, \omega(u) \neq \omega(v)\}.$$

Let $\Delta = \Delta(G)$ denote the maximum degree of G . It is well known that the condition $q \geq \Delta + 1$ guarantees the existence of a proper q -coloring, i.e., $\Omega \neq \emptyset$.

A canonical approach to sampling from Ω is the *Markov chain Monte Carlo* (MCMC) method, in which a Markov chain (X_t) is designed to mix rapidly to the uniform distribution π over Ω . Starting from an arbitrary initial state $X_0 \in \Omega$, simulating the chain for a sufficiently long time T yields a sample X_T approximately distributed as π . A well studied chain is the *Glauber dynamics*, which at each step selects a vertex uniformly at random and recolors it with a color not currently used by its neighbors. When $q \geq \Delta + 2$, the chain is ergodic and converges to the uniform distribution π over Ω . Determining the critical threshold on q for rapid mixing is a long-standing open problem. Over the past decades, a sequence of works has progressively improved this threshold, from the classical bound $q > 2\Delta$ [Jer95, SS97, BD97] to the current best bound $q > 1.809\Delta$ [Vig99, CDM⁺19, CV25].

1.1. Coupling from the past and grand coupling. The *coupling from the past* (CFTP) technique, introduced by Propp and Wilson [PW96], provides a general framework for the *perfect simulation* of the stationary distribution of a Markov chain. The key idea is to consider an idealized Markov chain that has been running from time $-\infty$ up to time 0, so that its final state $X_0 \sim \pi$ is an exact (perfect) sample from the stationary distribution π . To generate such a sample, the CFTP algorithm searches for a (possibly random) time $-T$ in the past such that, when the chain is simulated forward from every possible initial state at time $-T$ using the same random bits, all trajectories

$$X_{-T} \rightarrow X_{-T+1} \rightarrow \cdots \rightarrow X_{-1} \rightarrow X_0$$

coalesce into the same state at time 0. Once such a coalescence time is identified, the final state $X_0 \sim \pi$ is completely determined by the random bits used in the simulation, independent of the initial state at time $-T$. Consequently, X_0 is a perfect sample from the stationary distribution π .

At the heart of CFTP lies the notion of a *grand coupling*, which provides a unified probabilistic construction that simultaneously drives the evolution of the chain from all possible initial states.

State Key Laboratory for Novel Software Technology, New Cornerstone Science Laboratory, Nanjing University, 163 Xianlin Avenue, Nanjing, Jiangsu Province, China. Emails: 652024330006@smail.nju.edu.cn, liuhongyang@smail.nju.edu.cn, yinyt@nju.edu.cn, 5020243300075@smail.nju.edu.cn.

Formally, for a Markov chain with transition kernel $P(x, y) = \Pr[X_{t+1} = y \mid X_t = x]$ on state space Ω , a grand coupling is a mapping $g : \Omega \times [0, 1] \rightarrow \Omega$, such that for $U \sim \text{Uniform}[0, 1]$ and every $x, y \in \Omega$,

$$(1) \quad \Pr[g(x, U) = y] = P(x, y).$$

In other words, $g(\cdot, U)$ simulates a single Markov chain transition under shared randomness U , thereby coupling all trajectories in a common probability space. This notion extends naturally to the T -step transitions, represented by a grand coupling $F : \Omega \times [0, 1] \rightarrow \Omega$ for the evolution $X_{-T} \rightarrow X_0$. The CFTP algorithm identifies a time $-T$ in the past and terminates once $F(\cdot, U)$ becomes constant, that is, when all trajectories coalesce into a single state at time 0.

Designing grand couplings that achieve fast coalescence is therefore the central problem in CFTP. Beyond CFTP, grand couplings have found applications in a variety of contexts, including derandomization of MCMC [FGW⁺23], analytic stability (absence of complex zeros) [LWYY25], and local algorithms for sampling [LWY26]. Despite being a natural and fundamental concept, their systematic design and analysis remain poorly understood. While fast-coalescing grand couplings immediately imply rapid mixing of the underlying Markov chain, essentially nothing is known about the converse in general. This raises a fundamental question:

Does fast coalescence of grand couplings require stronger conditions than rapid mixing?

1.2. Bounding chains for q -colorings. A key algorithmic challenge in implementing CFTP lies in detecting coalescence, that is, determining whether the induced random mapping on the state space Ω has become a constant function. Formally, let $\Omega_t \subseteq \Omega$ denote the set of all possible states of X_t that can be reached from different initial configurations $X_{-T} \in \Omega$ under the same randomness. Specifically,

$$\Omega_t := \{g_t \circ g_{t-1} \circ \dots \circ g_{-T}(\omega) : \omega \in \Omega\},$$

where each $g_t = g_t(\cdot, U)$ denotes the grand coupling for the transition $X_t \rightarrow X_{t+1}$ with $U \sim \text{Uniform}[0, 1]$. Initially, $\Omega_{-T} = \Omega$, and clearly $|\Omega_{t+1}| \leq |\Omega_t|$ since g_t is a function. Coalescence occurs once $|\Omega_t| = 1$ for some $t \leq 0$. In certain special cases, most notably in monotone systems such as the ferromagnetic Ising model, where a natural partial order on Ω is preserved by the grand coupling, checking coalescence is straightforward. In contrast, for general systems, verifying coalescence is substantially more intricate and can be computationally intractable.

The *bounding chains* method, introduced by Huber [Hub98, Hub04], provides an efficient way to detect coalescence by tracking a “bounding box” of the subset Ω_t in the space $\Omega \subseteq [q]^V$. The bounding chain maintains, for each vertex $v \in V$, a *bounding list* $\mathcal{L}_t(v) \subseteq [q]$ of colors, which contains all colors that vertex v may take under any configuration in Ω_t . The corresponding bounding box is given by

$$\mathcal{B}_t := \bigotimes_{v \in V} \mathcal{L}_t(v).$$

Initially, we set $\mathcal{L}_{-T}(v) = [q]$ for every vertex v . The bounding list configuration $\mathcal{L}_t = (\mathcal{L}_t(v))_{v \in V}$ evolves as a Markov chain: at each time step $t \geq -T$, the next bounding list $\mathcal{L}_{t+1}$ is constructed from the current list $\mathcal{L}_t$, under the conservative assumption that every configuration $\omega \in \mathcal{B}_t$ may appear at time t . By induction, we have $\Omega_t \subseteq \mathcal{B}_t$ for any $t \geq -T$. Thus, coalescence at time t occurs if the bounding box has collapsed to the size $|\mathcal{B}_t| = 1$, or equivalently, when $|\mathcal{L}_t(v)| = 1$ for all $v \in V$.

The fast coalescence of bounding chains relies on the following standard contraction property.

Condition 1.1 (Contraction of bounding chains). *Let $(\mathcal{L}_t)$ be a bounding chain. At any time t , it holds*

$$(2) \quad \mathbb{E} \left[\sum_{v \in V} |\mathcal{L}_{t+1}(v)| \mid \mathcal{L}_t \right] \leq \sum_{v \in V} |\mathcal{L}_t(v)|.$$

Intuitively, this condition requires the total size of the bounding lists to be monotonically non-increasing in expectation at each step, even in the worst case. All existing CFTP algorithms based on bounding chains satisfy this condition, including [Hub98, BC20, JSS21]. The seminal work of Huber [Hub98] proposed a bounding chain algorithm for uniform sampling of q -coloring that is fast coalescing under the condition $q \geq (1 + o(1))\Delta^2$. For many years, this remained the best achievable bound for any CFTP algorithm for q -coloring, until recent breakthroughs [BC20, JSS21] improved the condition for fast-coalescing bounding chains first down to $q > 3\Delta$ and later to $q \geq (\frac{8}{3} + o(1))\Delta$.

1.3. Our results. We establish an asymptotically optimal threshold of $q = 2.5\Delta$ for CFTP algorithms based on bounding chains for proper q -colorings. In particular, we design an efficient CFTP algorithm that samples uniform proper q -colorings under the improved condition $q \geq (2.5 + o(1))\Delta$.

Theorem 1.2 (Upper bound). *There exists a CFTP algorithm that, given an undirected graph $G = (V, E)$ with maximum degree $\Delta \geq 3$ and $q > 2.5\Delta + 2\sqrt{(\log \Delta + 1)\Delta}$, outputs a uniformly random proper q -coloring of G in expected time $\tilde{O}(n\Delta^2)$, where $\tilde{O}(\cdot)$ hides poly-logarithmic factors.*

The algorithm in Theorem 1.2 is based on a new construction of bounding chains for the Glauber dynamics on uniform proper q -colorings, which coalesce rapidly using substantially fewer colors.

This result advances the state of the art for CFTP algorithms and, more broadly, for all perfect sampling algorithms for q -colorings that achieve exponential convergence.

Remark 1.3 (Perfect sampling via deterministic counting or approximate sampling). *Perfect sampling can also be achieved via generic reductions to deterministic approximate counting or approximate sampling. In the classical work of Jerrum, Valiant, and Vazirani [JVV86], a general reduction from perfect sampling to deterministic approximate counting (FPTAS) was established using a rejection-sampling framework. More recently, Göbel, Liu, Manurangsi, and Pappik [GLMP24] presented a reduction from perfect sampling to rapidly mixing Markov chains: once the chain has mixed sufficiently well (so that the total variation error is exponentially small), a JVV-style filtering step (implemented via exponentially exhaustive enumeration) is applied with exponentially small probability, keeping the overall expected runtime polynomial.*

Through these reductions, the known thresholds for polynomial-time perfect sampling of q -colorings have been improved to $q > 2\Delta$ via deterministic approximate counting [LSS19] and to $q > 1.809\Delta$ via rapidly mixing Markov chains [CV25].

Despite these advances, the coupling-from-the-past (CFTP) paradigm remains of great theoretical and practical significance. It provides a systematic framework for designing Las Vegas samplers with exponentially convergent runtime. Specifically, letting T denote the random runtime, we have

$$\Pr[T \geq k \cdot \mathbb{E}[T]] \leq \exp(-\Omega(k)), \quad \text{for all } k \geq 1.$$

where in our case the expected runtime satisfies $\mathbb{E}[T] = \tilde{O}(n\Delta^2)$. This sharp concentration arises from the efficient certificate of perfect sampling inherently provided by CFTP through coalescence, a key advantage over the certificates in [JVV86, GLMP24], which rely on deterministic counting.

Beyond its role in perfect sampling, the grand coupling is a natural and fundamental construct for Markov chains, and has found implications in several key aspects, including studies of derandomization [FGW⁺23], phase transitions [LWYY25], and local algorithms for sampling [LWY26]. It is therefore of central theoretical interest to understand the algorithmic power of grand couplings.

We show that the upper bound in Theorem 1.2 is asymptotically tight by establishing a matching lower bound for bounding-chain-based grand couplings satisfying the contraction property in Condition 1.1.

Theorem 1.4 (Lower bound). *Let $\Delta \geq 3$, $q < 2.5\Delta - 1$, and let n be sufficiently large. There exists a Δ -regular graph $G = (V, E)$ with $|V| = n$, together with a bounding list configuration $\mathcal{L} = (\mathcal{L}(v))_{v \in V}$, such that for any bounding chain for the Glauber dynamics on proper q -colorings, the contraction property (2) in Condition 1.1 is violated when the bounding chain is in state $\mathcal{L}_t = \mathcal{L}$.*

Together, our results establish a tight bound for CFTP for q -colorings based on bounding chains. They show that contractive bounding chains require a significantly stronger condition on the number of colors q than is needed for rapid mixing of the underlying Markov chain. In particular, the critical threshold $q > 2.5\Delta$ identified here is substantially stronger than the classical Dobrushin condition for q -colorings, which holds when $q > 2\Delta$ and corresponds to the worst-case contraction condition in path coupling [Jer95, SS97, BD97], thereby guaranteeing rapid mixing.

Remark 1.5. *In [JSS21], Jain, Sah, and Sawhney also identified $q = 2.5\Delta$ as a natural barrier for their bounding chain construction, reflecting a limitation of the algorithmic framework adopted in [BC20, JSS21]. In contrast, our lower bound shows that this threshold is inherent to all contractive bounding chains: no grand coupling based on such constructions can surpass it. Both our general lower bound and the optimal bounding chain design arise from a new characterization of bounding-chain-based grand couplings as an optimization problem, detailed in Section 3.*

1.4. Overview of conceptual and technical contributions. In the CFTP framework, the random evolution $X_{-T} \rightarrow X_0$ of the Markov chain (X_t) is captured by a random mapping $F : \Omega \rightarrow \Omega$, generated by a grand coupling $F = F(\cdot, U)$ using a random seed $U \in [0, 1]$. Equivalently, F can be expressed as a composition of single-step couplings, $F = g_{-1} \circ g_{-2} \circ \dots \circ g_{-T}$, where each $g_t = g_t(\cdot, U_t)$ corresponds to the grand coupling of the one-step transition at time t , determined by an independent random seed $U_t \sim \text{Uniform}[0, 1]$. Conceptually, the global seed U can thus be viewed as an infinite sequence of independent random variables $(U_{-1}, U_{-2}, \dots)$, each determining one step of the coupled evolution.

When the Markov chain (X_t) is a single-site dynamics, each transition g_t is specified by a pair (v_t, f_{v_t}) , where $v_t \in V$ denotes the vertex selected for update at time t , and f_{v_t} is a local random update function that recolors v_t based on the current coloring of its neighbors. Once the sequence of update vertices $v_{-1}, v_{-2}, \dots$ is fixed, the CFTP process is fully determined by these local updates.

Conceptually, the bounding-chain construction can then be viewed as a *localization* of this global coupling: rather than tracking the evolution of the entire configuration space $\Omega \subseteq [q]^V$, it maintains for each vertex v a local bounding list $\mathcal{L}(v) \subseteq [q]$ of possible colors, which compactly encodes all configurations consistent with the current stage of the coupled evolution.

Remark 1.6. *In the original formulation of bounding chains for single-site dynamics [Hub98], the update vertex is chosen uniformly at random from all vertices. Subsequent developments [BC20, JSS21] introduced more flexible update schedules to accelerate coalescence. For correctness, however, the update schedule must be specified a priori, independent of the evolution of the bounding lists during execution.*

In the bounding-chain framework, the sequence of update vertices $(v_{-1}, v_{-2}, \dots)$ forms part of the shared randomness and is identically coupled across all global configurations. That is, every coupled trajectory uses the same pre-specified sequence of update locations. This identically coupled choice of the update sequence (v_t) ensures that each update (v_t, f_{v_t}) acts purely locally and that the bounding-chain coupling remains well defined regardless of how the bounding lists evolve.

Formulation of grand coupling via bounding chains. Building on the local perspective above, we introduce a unified formulation of the bounding-chain method. In this framework, a global grand coupling of the Markov chain is specified via a collection of *local grand couplings* defined over the bounding boxes of neighborhood configurations, as formalized in Definition 3.2.

This subsumes all prior constructions of bounding chains, including those based on the standard uniform random update schedule as well as more structured, designed schedules, and enables a formal analysis of the fundamental limitations of the bounding-chain approach.

Lower bound: limitations of bounding chains. Theorem 1.4 establishes a lower bound for bounding chains for q -coloring Glauber dynamics satisfying the standard contraction condition (Condition 1.1). This applies to any schedule, as long as updates follow the Glauber dynamics. It stems from a fundamental obstruction to the coalescence of bounding chains: the reduction of bounding lists from size two to one.

Theorem 1.7 (Lower bound: 2-to-1 contraction). *Let $\Delta \geq 3$, $q < 2.5\Delta - 1$, and let n be sufficiently large. There exist a Δ -regular graph $G = (V, E)$ with $|V| = n$ and a bounding list configuration $\mathcal{L} = (\mathcal{L}(v))_{v \in V}$ with $|\mathcal{L}(v)| = 2$ for all $v \in V$, such that for any grand coupling of q -coloring Glauber dynamics given $\mathcal{L}$, the updated $\mathcal{L}'$ satisfies: there exists a $v \in V$ with $\mathbb{E}[|\mathcal{L}'(v)|] > 2$, while for all $u \in V \setminus \{v\}$, $|\mathcal{L}'(u)| = 2$.*

The construction in Theorem 1.7 results in an expected increase in the total size $\sum_v |\mathcal{L}'(v)|$ relative to $\sum_v |\mathcal{L}(v)|$, violating Condition 1.1. Thus, Theorem 1.4 follows from Theorem 1.7. Moreover, since adding colors to $\mathcal{L}$ can only increase size of $\mathcal{L}'$, it blocks coalescence from any bounding list sizes > 1 .

Upper bound: optimal design of grand coupling. We formalize the design of a grand coupling as an explicit optimization problem (Problem 3.4). To make the problem tractable, we introduce a simplified variant (Problem 4.1) that captures the key design issues underlying prior bounding-chain algorithms.

Solutions to this simplified problem can be characterized via a linear program, which yields a grand coupling that optimally solves Problem 4.1. This optimal construction, called SEEDING, provides a key component: Combined with several existing bounding-chain building blocks, including the COMPRESS coupling [BC20] and the DISJOINT coupling [JSS21], we finally obtain a fast coalescing grand coupling for q -coloring when $q > (2.5 + o(1))\Delta$, matching, up to lower-order terms, the lower bound established for any bounding-chain-based CFTP.

1.5. Related work. The critical condition for rapid mixing of local Markov chains on proper q -colorings, such as Glauber dynamics, has been a central question in the study of sampling algorithms. A classic result of Jerrum [Jer95] (and independently, Salas and Sokal [SS97]) established that when $q > 2\Delta$, the Glauber dynamics mixes in $O(n \log n)$ time. This foundational work laid the groundwork for coupling-based analyses, notably inspiring the path coupling technique of Bubley and Dyer [BD97] and the coupling criterion of Hayes [Hay06]. Vigoda [Vig99] introduced a coupling for the *flip dynamics* and proved an $O(n^2)$ mixing-time bound for the Glauber dynamics when $q > \frac{11}{6}\Delta$. This threshold remained the state of the art for nearly two decades until Chen et al. [CDM⁺19] further refined it to $q > (\frac{11}{6} - \varepsilon_0)\Delta$ for an explicit $\varepsilon_0 = 1/84000$. Most recently, Carlson and Vigoda [CV25] improved the bound to $q > 1.809\Delta$, which remains the best-known rapid-mixing threshold for general graphs.

Under additional structural assumptions, such as a lower bound on the girth, sharper thresholds are known. For instance, on triangle-free graphs, the Glauber dynamics mixes optimally whenever $q > \alpha^*\Delta$ for some constant $\alpha^* \approx 1.763$; see [HV06, CFYZ21, FGYZ22, CLV21, JPV22]. Even stronger results hold for graphs of larger girth [HV03, DFHV13, CLMM23]. In particular, the recent breakthrough of Chen et al. [CLMM23] establishes an almost tight condition $q \geq \Delta + 3$ for an $O(n \log n)$ mixing time of the Glauber dynamics on graphs with constant maximum degree $\Delta = O(1)$ and sufficiently large girth.

The conjectured optimal threshold for rapid mixing is $q > \Delta$, corresponding to the uniqueness threshold of the Gibbs measure on the infinite Δ -regular tree [Jon02]. In contrast, approximate sampling from q -colorings becomes computationally intractable when $q < \Delta$, unless $\text{NP} = \text{RP}$ [GŠV15].

In a separate line of work focused on deterministic approximate counting, Gamarnik and Katz [GK12] gave a polynomial-time approximate counting algorithm for proper q -colorings when $q > 2.84\Delta$ for bounded-degree graphs. Subsequent improvements were obtained by Lu and Yin [LY13] ($q > 2.58\Delta$) and by Liu, Sinclair, and Srivastava [LSS19] ($q > 2\Delta$).

Related works on CFTP. Coupling from the past (CFTP) is a central technique for perfect sampling and has been widely applied across statistical physics and combinatorial models. Classical applications include the Ising and Potts models [HN98, Hub04], while extensive developments on graph colorings have leveraged the bounding-chain framework [Hub04, Hub16, BC20, JSS21]. More recently, CFTP methods have been extended to the sampling Lovász Local Lemma, leading to perfect samplers for solutions to atomic constraint satisfaction problems (CSPs) [HSW21, QWZ22].

Other perfect sampling paradigms. Beyond CFTP, several alternative paradigms for perfect sampling have been developed, including the *random recycler* [FH00], *partial rejection sampling* [Wil96, CPP02, GJL19, FVY21], and approaches based on the *Bayes filter* [FGY22]. More recently, *local perfect samplers* have been introduced, which generate samples from marginal distributions at a local cost. Examples include the *lazy depth-first sampler* (also known as the Anand–Jerrum algorithm) [AJ22, HWY22] and the *coupling towards the past* framework [FGW⁺23, LWYY25, LWY26], the latter extending the philosophy of coupling from the past to more general local settings.

2. PRELIMINARIES

2.1. Notations for coloring. Let $G = (V, E)$ be an undirected graph and $[q] = \{1, 2, \dots, q\}$ a set of $q \geq 2$ colors. The configuration space of proper q -colorings of G is

$$\Omega = \{\omega : V \rightarrow [q] \mid \omega(u) \neq \omega(v) \text{ for all } \{u, v\} \in E\}.$$

For each vertex v , let $\Gamma(v) := \{u \in V \mid \{u, v\} \in E\}$ denote v 's neighborhood.

2.2. Glauber dynamics. A *single-site dynamics* on Ω is defined by a predetermined *update schedule*, which is a (possibly random) sequence of vertices (v_t) , together with a collection of *local update distributions*:

$$\left\{ p_v^\omega : v \in V, \omega \in [q]^{\Gamma(v)} \right\},$$

where p_v^ω specifies a distribution over $[q]$ that depends only on the vertex v and the configuration ω of its neighbors. The resulting dynamics (X_t) evolves as follows: at each time step t , the vertex v_t is selected according to the update schedule, and its color is updated by sampling from $p_{v_t}^{X_t(\Gamma(v_t))}$, where $X_t(\Gamma(v_t))$ denotes the restriction of the current configuration X_t to the neighborhood $\Gamma(v_t)$ of v_t . r

The *Glauber dynamics* for proper q -colorings is a single-site dynamics in which, for each vertex $v \in V$ and neighborhood coloring $\omega \in [q]^{\Gamma(v)}$, the local update distribution at v is uniform over the set of colors available to v given ω :

$$(3) \quad p_v^\omega = \text{Uniform}([q] \setminus \omega),$$

where $\text{Uniform}(S)$ denotes the uniform distribution over a set S . Here, by slight abuse of notation, we treat the tuple ω equivalently as a set, so that $[q] \setminus \omega = [q] \setminus \{\omega(u) : u \in \Gamma(v)\}$.

2.3. Bounding chains. For a single-site dynamics (X_t) on a state space $\Omega \subseteq [q]^V$ with update schedule (v_t) and local update distributions $p_v^\omega : v \in V, \omega \in [q]^{\Gamma(v)}$, let $g_t : \Omega \times [0, 1] \rightarrow \Omega$ denote a grand coupling for the transition $X_t \rightarrow X_{t+1}$, as defined in (1).

Suppose that the randomness in the update schedule (v_t) is *identically coupled* across all trajectories, i.e., at each time t , every coupled trajectory updates the same randomly chosen vertex v_t .

Under this natural assumption, the coupling $g_t(\cdot, U)$ becomes a local function, such that:

- the image $g_t(\omega, U)$ depends only on $\omega \in \Omega$ restricted to $\Gamma(v_t)$;
- the image $g_t(\omega, U)$ differs from $\omega \in \Omega$ only at the update vertex v_t .

Hence, without loss of generality, we may write

$$g_t(\omega, U)(v_t) = g_t(\omega_{\Gamma(v_t)}, U)(v_t),$$

since the outcome of $g_t(\omega, U)$ at v_t depends solely on the neighborhood configuration $\omega_{\Gamma(v_t)}$.

This locality property of grand couplings for single-site dynamics naturally motivates the definition of *bounding chains*. A bounding chain $(\mathcal{L}_t)$ is naturally constructed from a grand coupling (g_t) of the single-site dynamics. At each time step t , the bounding chain maintains, for every vertex $v \in V$, a *bounding list* (or *bounding set*) $\mathcal{L}_t(v) \subseteq [q]$ representing the possible colors that v may take at time t under the coupled evolution, inferred based on the bounding lists.

Formally, given the shared random seed U_t and the update vertex v_t , the bounding chain evolves as:

$$(4) \quad \mathcal{L}_{t+1}(v) = \begin{cases} \left\{ g_t(\omega, U_t)(v) : \omega \in \bigotimes_{u \in \Gamma(v)} \mathcal{L}_t(u) \right\}, & \text{if } v = v_t, \\ \mathcal{L}_t(v) & \text{otherwise.} \end{cases}$$

In words, when vertex v_t is updated, the new bounding list $\mathcal{L}_{t+1}(v_t)$ consists of all possible colors that v_t can take, given the current bounding lists of its neighbors and the shared randomness U_t at time t . For all other vertices, the bounding lists remain unchanged. In practice, one may instead maintain a superset of the set defined in (4), as it can be easier to compute. Such relaxations preserve the correctness of sampling, though they may affect the efficiency of coalescence.

In coupling from the past (CFTP), the bounding chain is initialized with full uncertainty: $\mathcal{L}_{-T}(v) = [q]$ for all $v \in V$. As the chain $(\mathcal{L}_t)$ evolves forward from time $-T$ to time 0, the bounding lists $\mathcal{L}_t(v)_{v \in V}$ gradually shrink as uncertainty about each vertex's color decreases. When all bounding lists collapse to singletons, i.e., $|\mathcal{L}_t(v)| = 1$ for every $v \in V$, the coupled single-site dynamics have coalesced to a single configuration. At this point, the bounding-chain-based CFTP procedure outputs a perfect sample from the stationary distribution of the Glauber dynamics.

3. GRAND COUPLING VIA BOUNDING CHAINS

In this section, we formalize a unified framework for grand couplings arising from bounding chains. We begin by formally defining the notion of a grand coupling.

Definition 3.1 (Grand coupling). *Let Ω be a finite sample space, and let $\mathcal{P}$ be a family of probability distributions on Ω . A measurable function $f : \mathcal{P} \times [0, 1] \rightarrow \Omega$ is called a grand coupling on $\mathcal{P}$ if, for a uniform random variable $U \sim \text{Uniform}[0, 1]$, it holds that for every $p \in \mathcal{P}$ and every $y \in \Omega$,*

$$\Pr[f(p, U) = y] = p(y).$$

Equivalently, for each fixed $p \in \mathcal{P}$, the random variable $f(p, U)$ has the distribution p .

A grand coupling thus provides a unified deterministic rule that simultaneously generates a family of probability distributions on the same sample space Ω using shared randomness $U \sim \text{Uniform}[0, 1]$. In particular, the grand coupling for a Markov chain transition P defined in (1) is a special case of Definition 3.1, where the family $\mathcal{P}$ consists of the update distributions $P(x, \cdot)$ for all $x \in \Omega$.

3.1. Local grand coupling via bounding chains. We introduce a general formulation for grand couplings constructed using bounding chains, which applies to single-site dynamics. Let $G = (V, E)$ be a graph and $[q] = \{1, 2, \dots, q\}$ a set of $q \geq 2$ colors. Recall that a single-site dynamics on a state space $\Omega \subseteq [q]^V$ is specified via a collection of local update distributions: $\{p_v^\omega : v \in V, \omega \in [q]^{\Gamma(v)}\}$.

In the bounding-chain framework, we maintain for each vertex $v \in V$ a *bounding list* $\mathcal{L}(v) \subseteq [q]$, representing the set of possible colors that v may take across coupled trajectories. The grand coupling is then constructed based on these bounding lists.

Definition 3.2 (Grand coupling via bounding chain). *Let $G = (V, E)$ be a graph and $q \geq 2$. Consider a single-site dynamics on the state space $\Omega \subseteq [q]^V$ with local update distributions $\{p_v^\omega : v \in V, \omega \in [q]^{\Gamma(v)}\}$.*

Fix a vertex $v \in V$ and a bounding list configuration $(\mathcal{L}(u))_{u \in \Gamma(v)}$, with each $\mathcal{L}(u) \subseteq [q]$. Define

$$(5) \quad \mathcal{P}_v := \left\{ p_v^\omega : \omega \in \bigotimes_{u \in \Gamma(v)} \mathcal{L}(u) \right\},$$

the family of all possible local update distributions at v consistent with the bounding lists.

A local grand coupling at vertex v given the bounding list $(\mathcal{L}(u))_{u \in \Gamma(v)}$ is a measurable mapping

$$f_v : \mathcal{P}_v \times [0, 1] \rightarrow [q]$$

such that f_v is a valid grand coupling on $\mathcal{P}_v$ in the sense of Definition 3.1.

Previous bounding-chain based algorithms (e.g., [BC20, JSS21]) implicitly employ two stages. These stages can be unified in our framework for grand couplings via bounding chains.

- *Generating stage:* Once a random seed $U \sim \text{Uniform}[0, 1]$ is generated, the function f_v determines, for every $p \in \mathcal{P}_v$, a possible update outcome of v . The collection of all outcomes $\{f_v(p, U) : p \in \mathcal{P}_v\}$ forms the new bounding list $\mathcal{L}'(v)$ of vertex v , representing all colors that could be assigned under the current randomness.
- *Decoding stage:* Once the actual neighbor configuration $\omega \in [q]^{\Gamma(v)}$ (and thus the corresponding update distribution $p_v^\omega \in \mathcal{P}_v$) is revealed, the updated color of v is determined by $f_v(p_v^\omega, U)$.

This formulation unifies the generating–decoding paradigm into a single mathematical framework, enabling an analysis of the optimality of grand coupling designs.

3.1.1. Grand coupling for Glauber dynamics. Now consider the Glauber dynamics on proper q -colorings. For each vertex v and neighborhood configuration $\omega \in [q]^{\Gamma(v)}$, the local update distribution p_v^ω at v given ω is uniform over the set of colors available to v given ω , that is, $p_v^\omega = \text{Uniform}([q] \setminus \omega)$, where we adopt the abuse of notation as in (3).

Under this specialization, the grand coupling f_v in Definition 3.2 admits an *interval-based* model: The preimages of $f_v(p, \cdot)$ induce a measurable partition of $[0, 1]$ corresponding to the outcomes of p . Let $k := |[q] \setminus \omega|$ denote the number of available colors at v . The mapping $f_v(p_v^\omega, \cdot)$ defines a partition of the unit interval $[0, 1]$ into k subsets of equal measure $1/k$, each labeled by a distinct available color. A shared random seed $U \sim \text{Uniform}[0, 1]$ specifies a point in $[0, 1]$, and the update color $f_v(p_v^\omega, U)$ is the color assigned to that point in the partition corresponding to p_v^ω .

Example 3.3. Suppose $q = 5$, and vertex v has two neighbors with bounding lists $\{1, 2\}$ and $\{2, 3\}$, respectively. Then the possible sets of colors used by these two neighbors are $\{2, 3\}$, $\{1, 3\}$, $\{1, 2\}$, $\{2\}$; and the corresponding sets of available colors at v are their complements $\{1, 4, 5\}$, $\{2, 4, 5\}$, $\{3, 4, 5\}$, $\{1, 3, 4, 5\}$. In the interval-based model, the grand coupling f_v associates each case with a colored partition of $[0, 1]$, where the available colors are distributed uniformly over the interval. Given the shared randomness $U \in [0, 1]$, a new color for v is determined in each case by the color at position U . The updated bounding list $\mathcal{L}'(v)$ at vertex v is then the set of all such colors across all cases. This is illustrated in Figure 1.




FIGURE 1. Each line illustrates how the unit interval $[0, 1]$ is partitioned and assigned to the available colors, thereby defining a uniform distribution over them using the shared source of randomness $U \sim \text{Uniform}[0, 1]$. A single random position $U \in [0, 1]$ determines the color selected for each case, and the union of these selected colors forms the updated bounding list $\mathcal{L}'(v)$. In this example, for any choice of U , the updated bounding list satisfies $|\mathcal{L}'(v)| = 2$.

3.1.2. *Classical construction of grand coupling for q -colorings.* We next illustrate how the formulation in Definition 3.2 specializes to the classical grand coupling of [Hub98]. In this coupling, the shared randomness is a uniformly random permutation π of $[q]$. For each vertex v with neighborhood coloring $\omega \in [q]^{\Gamma(v)}$, the local rule f_v outputs the first color in π that does not appear in ω . It is straightforward to verify that, for any ω , this rule samples uniformly from the set of available colors $[q] \setminus \omega$.

In the interval-based interpretation, the unit interval $[0, 1]$ is partitioned uniformly into $q!$ subintervals, each corresponding to a distinct permutation of $[q]$. For each neighborhood coloring ω , and for each subinterval, let π denote the permutation associated with that subinterval; the subinterval is then assigned the first color in π that does not appear in ω . Given a shared random seed $U \sim \text{Uniform}[0, 1]$, the subinterval containing U determines the color selected for each local update distribution $p \in \mathcal{P}_v$. Ideally, the new bounding list at vertex v is

$$\mathcal{L}'(v) = \{f_v(p, U) : p \in \mathcal{P}_v\},$$

where $\mathcal{P}_v$ contains all local update distributions given the current bounding lists, as defined in (5).

Although $|\mathcal{P}_v|$ may be exponentially large, it is not necessary to evaluate $f_v(p, U)$ for every $p \in \mathcal{P}_v$. A key observation here is that including extra colors into $\mathcal{L}'(v)$ does not affect correctness of CFTP (though it may slow down coalescence). Let π denote the permutation corresponding to the interval containing U . Since the number of distinct colors in a neighborhood satisfies $|\{\omega(u) : u \in \Gamma(v)\}| \leq \Delta$, at least one of the first $\Delta + 1$ entries of π is available for v . Hence, it suffices to take

$$\{\pi_1, \pi_2, \dots, \pi_{\Delta+1}\},$$

as the updated bounding list $\mathcal{L}'(v)$, since $\{f_v(p, U) : p \in \mathcal{P}_v\} \subseteq \{\pi_1, \pi_2, \dots, \pi_{\Delta+1}\}$.

For further refinement, observe that if a color c does not appear in the bounding list $\mathcal{L}(u)$ of any neighbor $u \in \Gamma(v)$, then c cannot appear in any neighborhood coloring $\omega \in \bigotimes_{u \in \Gamma(v)} \mathcal{L}(u)$. Let C denote the set of all such colors. If there exists an index $i \leq \Delta + 1$ such that $\pi_i \in C$, it suffices to take

$$\{\pi_1, \pi_2, \dots, \pi_i\}$$

as the updated bounding list. Indeed, $\pi_i \in C$ guarantees that for every neighborhood coloring $\omega \in \bigotimes_{u \in \Gamma(v)} \mathcal{L}(u)$, the grand coupling outcome $f_v(p_v^\omega, U)$ never appears after position i in π . Hence, $\{f_v(p, U) : p \in \mathcal{P}_v\} \subseteq \{\pi_1, \pi_2, \dots, \pi_i\}$.

The grand coupling described above, together with this refinement, precisely recovers the bounding chain of [Hub98] for q -colorings. The effectiveness of a bounding chain depends on how well it tracks the bounding list $\{f_v(p, U) : p \in \mathcal{P}_v\}$. Smaller bounding lists may result in a tighter bounding chain and thus faster coalescence in the CFTP process. Subsequent works [BC20, JSS21] refined this idea by designing more sophisticated grand couplings. Their algorithms partition the process into multiple *phases*, each employing a distinct grand coupling, allowing the bounding lists to shrink gradually across

phases. Despite differences in implementation, all such couplings admit a natural interpretation within the framework in Definition 3.2.

Within this formulation, designing an effective grand coupling can be quantitatively phrased as minimizing the size of $\{f_v(p, U) : p \in \mathcal{P}_v\}$ for $U \sim \text{Uniform}[0, 1]$. In the following sections, we analyze this objective under specific structural assumptions on the neighboring bounding lists $(\mathcal{L}(u))_{u \in \Gamma(v)}$ and establish both upper and lower bounds for grand couplings within this model.

3.1.3. An optimization problem on grand coupling. As discussed above, the performance of a grand coupling is naturally quantified by the size of its updated bounding list:

$$\mathcal{L}'(v) = \{f_v(p, U) : p \in \mathcal{P}_v\}, \quad U \sim \text{Uniform}[0, 1].$$

A smaller expected size of $\mathcal{L}'(v)$ leads to faster coalescence in the coupling-from-the-past procedure. This motivates designing grand couplings that minimize the updated bounding list given the current neighborhood bounding lists. We formalize this task as the following optimization problem.

Problem 3.4 (Local grand coupling optimization). *Fix a vertex $v \in V$ and a bounding list configuration $(\mathcal{L}(u))_{u \in \Gamma(v)}$, with each $\mathcal{L}(u) \subseteq [q]$. Let*

$$(6) \quad \mathcal{P}_v = \left\{ \text{Uniform}([q] \setminus \omega) : \omega \in \bigotimes_{u \in \Gamma(v)} \mathcal{L}(u) \right\}$$

denote the set of all possible local update distributions at v in the Glauber dynamics on q -colorings.

The objective is to design a local grand coupling $f_v : \mathcal{P}_v \times [0, 1] \rightarrow [q]$ at vertex v given the bounding lists $(\mathcal{L}(u))_{u \in \Gamma(v)}$ (cf. Definition 3.2) that minimizes

$$\mathbb{E}_{U \sim \text{Uniform}[0, 1]}[|\{f_v(p, U) : p \in \mathcal{P}_v\}|].$$

That is, f_v seeks to minimize the expected size of the updated bounding list at v .

3.2. Lower bounds for grand coupling. In this subsection we prove Theorem 1.4 by establishing a fundamental lower bound for Problem 3.4. Our analysis reveals an intrinsic obstacle that limits the effectiveness of all known bounding-chain constructions [Hub98, BC20, JSS21]:

The difficulty of reducing every bounding list from size 2 to 1.

This step represents the critical bottleneck in coalescence of bounding chains. To understand this limitation, we focus on the configuration where each neighbor $u \in \Gamma(v)$ has a bounding list $\mathcal{L}(u)$ of size 2. This setting captures the worst-case local uncertainty at vertex v , and hence determines the strongest contraction that any grand coupling can possibly attain. Our lower-bound argument is therefore developed under this extremal configuration.

Theorem 3.5. *Let $q < 2.5\Delta - 1$, and let v be a vertex with degree $|\Gamma(v)| = \Delta$. There exists a bounding list configuration $(\mathcal{L}(u))_{u \in \Gamma(v)}$ with $|\mathcal{L}(u)| = 2$ for all $u \in \Gamma(v)$ such that, for every local grand coupling $f_v : \mathcal{P}_v \times [0, 1] \rightarrow [q]$ of the Glauber dynamics update at v given $(\mathcal{L}(u))_{u \in \Gamma(v)}$, where $\mathcal{P}_v$ is as in (6),*

$$\mathbb{E}_{U \sim \text{Uniform}[0, 1]}[|\{f_v(p, U) : p \in \mathcal{P}_v\}|] > 2.$$

Theorem 3.5 is established by explicitly constructing a bounding-list configuration exhibiting a structural bottleneck, formalized as follows.

Definition 3.6 (Worst-case triangle configuration). *We say that two vertices $u, v \in V$ form a triangle on a bounding list configuration $\mathcal{L}$ if $|\mathcal{L}(u)| = |\mathcal{L}(v)| = 2$ and $|\mathcal{L}(u) \cup \mathcal{L}(v)| = 3$.*

For a vertex $v \in V$, we say that $\mathcal{L}$ is in a worst-case configuration at v if the neighbors of v can be partitioned into disjoint pairs, each forming a triangle on $\mathcal{L}$.

Proof of Theorem 3.5. We consider the worst-case configuration at v from Definition 3.6, which realizes the scenario where no grand coupling can reduce the expected bounding list size at v to 2 or below.

Without loss of generality, assume Δ is even, write $\Delta = 2m$ and express $q = 3m + r$ for some integer $r \geq 0$. Partition v 's neighbors into m disjoint pairs. Index the pairs by $i = 1, \dots, m$. For the i -th pair in each side we assign bounding lists as follows: the two vertices in the i -th pair receive bounding lists

$\{3i - 2, 3i - 1\}$ and $\{3i - 1, 3i\}$ respectively. The colors $\{1, 2, \dots, 3m\}$ are called the *constrained* colors, while the remaining r colors are called *free* colors.

We first note that the construction above indeed satisfies Definition 3.6: each neighbor pair indexed by $i = 1, \dots, m$ has bounding lists $\{3i - 2, 3i - 1\}$ and $\{3i - 1, 3i\}$, so the two vertices in the pair form a triangle in the sense of Definition 3.6. Let f_v be an arbitrary grand coupling for the Glauber update at v , and write

$$\mathcal{L}_U(v) = \{f_v(p, U) : p \in \mathcal{P}_v\}$$

for the bounding list produced by f_v when the shared seed equals $U \sim \text{Uniform}[0, 1]$.

For $i = 1, \dots, m$, define the following events, where all probabilities are over $U \sim \text{Uniform}[0, 1]$:

$$\mathcal{A}_i := \{\mathcal{L}_U(v) \cap \{3i - 2, 3i\} \neq \emptyset\},$$

$$\mathcal{B}_i := \{3i - 1 \in \mathcal{L}_U(v)\},$$

$$\mathcal{C}_i := \{\{3i - 2, 3i\} \subseteq \mathcal{L}_U(v)\},$$

and let $\mathcal{D} := \{\mathcal{L}_U(v) \cap \{3m + 1, \dots, 3m + r\} \neq \emptyset\}$ denote the event that at least one *free* color (a color in $\{3m + 1, \dots, 3m + r\}$) appears in $\mathcal{L}_U(v)$. We claim the following two lower bounds:

$$(7) \quad \Pr[\mathcal{A}_i] \geq \frac{2}{m + r + 1}, \quad \forall i \in \{1, 2, \dots, m\};$$

$$(8) \quad \Pr[\mathcal{B}_i] \geq \frac{1}{m + r}, \quad \forall i \in \{1, 2, \dots, m\};$$

and a covering property:

$$(9) \quad \Pr[\mathcal{C}_1 \vee \dots \vee \mathcal{C}_m \vee \mathcal{D}] = 1, \quad \text{and consequently, } \sum_{i=1}^m \Pr[\mathcal{C}_i] + \Pr[\mathcal{D}] \geq 1.$$

Assuming these inequalities, we obtain a lower bound on the expected bounding list size. Observe

$$\Pr[3i - 2 \in \mathcal{L}_U(v)] + \Pr[3i \in \mathcal{L}_U(v)] = \Pr[\mathcal{A}_i] + \Pr[\mathcal{C}_i], \quad \forall i \in \{1, 2, \dots, m\};$$

since the left-hand side counts the indicators of the two events and the right-hand side decomposes that sum into the indicator of “at least one” plus the indicator of “both”. Hence

$$\begin{aligned} \mathbb{E}[|\mathcal{L}_U(v)|] &= \sum_{c \in [q]} \Pr[c \in \mathcal{L}_U(v)] \\ &\geq \left(\sum_{c \in \{1, \dots, 3m\}} \Pr[c \in \mathcal{L}_U(v)] \right) + \Pr[\mathcal{L}_U(v) \cap \{3m + 1, \dots, 3m + r\} \neq \emptyset] \\ &= \sum_{i=1}^m (\Pr[3i - 2 \in \mathcal{L}_U(v)] + \Pr[3i - 1 \in \mathcal{L}_U(v)] + \Pr[3i \in \mathcal{L}_U(v)]) + \Pr[\mathcal{D}] \\ &= \left(\sum_{i=1}^m \Pr[\mathcal{A}_i] \right) + \left(\sum_{i=1}^m \Pr[\mathcal{B}_i] \right) + \left(\sum_{i=1}^m \Pr[\mathcal{C}_i] \right) + \Pr[\mathcal{D}]. \end{aligned}$$

Applying the bounds from (7)–(9) yields

$$\mathbb{E}[|\mathcal{L}_U(v)|] \geq \frac{2m}{m + r + 1} + \frac{m}{m + r} + 1.$$

Since $q < 2.5\Delta$ implies $r < 2m$, it follows that

$$\frac{2m}{m + r + 1} + \frac{m}{m + r} + 1 > \frac{2m}{3m} + \frac{m}{3m} + 1 = 2.$$

Hence for every $q < 2.5\Delta$ we have $\mathbb{E}[|\mathcal{L}_U(v)|] > 2$ as claimed.

It remains to prove (7)–(9). These are simple consequences of the marginal correctness of the grand coupling and the explicit worst-case assignment of neighbor lists:

- **Proof of (7).** Fix $i \in \{1, \dots, m\}$. For every pair $j \neq i$, assigns the two neighbors the colors $3j - 2$ and $3j$, and for the i -th pair assigns both neighbors the color $3i - 1$. Let ω be the neighborhood coloring constructed above. Then $|\omega| = \Delta - 1 = 2m - 1$, so the update distribution is $\text{Uniform}([q] \setminus \omega)$, which places mass $1/(q - (2m - 1)) = 1/(m + r + 1)$ on each of the two colors $3i - 2$ and $3i$.

By marginal correctness for this fixed distribution, each of the events “ $3i - 2 \in \mathcal{L}_U(v)$ ” and “ $3i \in \mathcal{L}_U(v)$ ” has measure at least $1/(m + r + 1)$, hence

$$\Pr[\mathcal{A}_i] \geq \Pr[3i - 2 \in \mathcal{L}_U(v)] + \Pr[3i \in \mathcal{L}_U(v)] \geq \frac{2}{m + r + 1}.$$

- **Proof of (8).** Fix $i \in \{1, \dots, m\}$. Let ω be the neighbor configuration that, for every pair j , assigns the two neighbors the colors $3j - 2$ and $3j$. Then $|\omega| = \Delta$, so the update distribution is $\text{Uniform}([q] \setminus \omega)$, which places mass $1/(q - \Delta) = 1/(m + r)$ on each middle color $3i - 1$.

By marginal correctness for this fixed distribution we obtain

$$\Pr[\mathcal{B}_i] \geq \Pr[3i - 1 \in \mathcal{L}_U(v)] \geq \frac{1}{m + r}.$$

- **Proof of (9).** Suppose for contradiction that for some U none of $C_1, \dots, C_m, \mathcal{D}$ occurs. Then for every i at most one of the colors $3i - 2, 3i$ belongs to $\mathcal{L}_U(v)$, and no free color appears in $\mathcal{L}_U(v)$. For each pair i choose neighbor colors from the lists $\{3i - 2, 3i - 1\}$ and $\{3i - 1, 3i\}$ so as to ensure every element of $\mathcal{L}_U(v) \cap \{3i - 2, 3i - 1, 3i\}$ is included in the resulting neighborhood set; this is always possible when at most one of $\{3i - 2, 3i\}$ is present. Doing this for all i yields a neighborhood coloring ω with $\mathcal{L}_U(v) \subseteq \omega$. But the coupling output $f_v([q] \setminus \omega, U)$ must lie in $[q] \setminus \omega$, contradicting the definition of $\mathcal{L}_U(v)$ as the set of possible outputs at seed U . Hence the assumption is false, so for every U at least one of $C_1, \dots, C_m, \mathcal{D}$ occurs, thus

$$\Pr[C_1 \vee \dots \vee C_m \vee \mathcal{D}] = 1.$$

Combining the three items completes the proof of Theorem 3.5. $\square$

We now present a global version of the above lower bound, which extends the local obstruction in Theorem 3.5 to the entire graph, thereby proving Theorem 1.4.

Theorem 3.7 (Worst-case lower bound). *Let $\Delta \geq 3$, $q < 2.5\Delta - 1$, and let n be sufficiently large. There exists a Δ -regular graph $G = (V, E)$ with $|V| = n$, together with a bounding list configuration $\mathcal{L} = (\mathcal{L}(u))_{u \in V}$ satisfying $|\mathcal{L}(u)| = 2$ for all $u \in V$, such that for every vertex $v \in V$ and every local grand coupling $f_v : \mathcal{P}_v \times [0, 1] \rightarrow [q]$ of the Glauber dynamics update at v given $\mathcal{L}$, where $\mathcal{P}_v$ is as in (6),*

$$\mathbb{E}_{U \sim \text{Uniform}[0, 1]}[|\{f_v(p, U) : p \in \mathcal{P}_v\}|] > 2.$$

Moreover, the worst-case configuration at each vertex (defined in Definition 3.6) can be simultaneously realized across all vertices of G .

Proof. Without loss of generality assume Δ is even and write $\Delta = 2m$. We first construct a Δ -regular graph on 2Δ vertices that realizes the worst-case bounding-list configuration at every vertex, and then obtain larger graphs by taking disjoint copies.

Let $G = (V_1 \cup V_2, E)$ be the complete bipartite graph with bipartition sets V_1 and V_2 , where $|V_1| = |V_2| = \Delta$. Then G is Δ -regular. Partition the vertices of V_1 into m disjoint pairs and likewise partition V_2 into m disjoint pairs. Index the pairs by $i = 1, \dots, m$. For the i -th pair in each side we assign bounding lists as follows: the two vertices in the i -th pair receive bounding lists $\{3i - 2, 3i - 1\}$ and $\{3i - 1, 3i\}$ respectively. Do this assignment independently for V_1 and V_2 (using the same indexing convention).

By construction each bounding list has size 2. Moreover, for any fixed vertex $v \in V_1$ its neighborhood $\Gamma(v) = V_2$ is partitioned into the m disjoint pairs on V_2 , and each such pair forms a triangle in the sense of Definition 3.6: the two lists in the pair each have size 2 and their union has size 3. Hence the collection of bounding lists on G is in the worst-case configuration at every $v \in V_1$. The same argument applies to every $v \in V_2$ because the construction on V_1 is symmetric. Therefore every vertex of G is in the worst-case configuration of Definition 3.6.

By Theorem 3.5, for any vertex v in such a worst-case configuration and for any grand coupling f_v for the single-site Glauber update at v we have

$$\mathbb{E}_{U \sim \text{Uniform}[0,1]}[|\{f_v(p, U) : p \in \mathcal{P}_v\}|] > 2.$$

Thus the 2Δ -vertex graph G constructed above satisfies the claimed property for every vertex.

Finally, for any sufficiently large n that is a multiple of 2Δ we obtain an n -vertex Δ -regular graph with the same property by taking the disjoint union of $n/(2\Delta)$ independent copies of G , and by carrying the same bounding-list assignment on each copy. Every vertex in this disjoint union is in the same local worst-case configuration, hence the conclusion of the theorem holds simultaneously for every vertex of the n -vertex graph. This completes the proof. $\square$

As noted in Section 2.3, any bounding-chain-based grand coupling selects a vertex v at random in each step and then applies a local grand coupling at v . Thus Theorem 3.7 is in fact a formal restatement of Theorem 1.7, and the latter can be derived as a straightforward corollary of the former.

4. DESIGN OF GRAND COUPLINGS

In this section, we explore the design of grand couplings with desirable contraction properties, a key ingredient for CFTP algorithms that succeed near the critical threshold.

4.1. A simplified formulation for grand coupling design. Based on the bounding-chain framework, we formulated in Problem 3.4 an optimization problem for designing grand couplings.

To obtain a more tractable and intuitive formulation, we approximate $\mathcal{P}_v$ by considering only the set of colors $\mathcal{S} = \bigcup_{u \in \Gamma(v)} \mathcal{L}(u)$ appearing in the neighborhood bounding list of v , which coarsely tracks the local uncertainty around v . This motivates the following simplified grand coupling design problem, which serves as an approximation of Problem 3.4.

Problem 4.1 (Grand coupling design problem). *Fix a subset of colors $\mathcal{S} \subseteq [q]$ and an integer Δ . Let $\mathcal{S}_\Delta := \{C \subseteq \mathcal{S} : |C| \leq \Delta\}$ denotes the collection of all subsets of $\mathcal{S}$ of size at most Δ , and define*

$$\mathcal{P} = \mathcal{P}(\mathcal{S}_\Delta) := \{\text{Uniform}([q] \setminus C) : C \in \mathcal{S}_\Delta\}.$$

The objective is to design a grand coupling $f : \mathcal{P} \times [0, 1] \rightarrow [q]$ (cf. Definition 3.1) that minimizes

$$\mathbb{E}_{U \sim \text{Uniform}[0,1]}[|\{f(p, U) : p \in \mathcal{P}\}|].$$

It is straightforward to verify that any grand coupling $f : \mathcal{P} \times [0, 1] \rightarrow [q]$ that is a feasible solution to Problem 4.1 on an instance $(\mathcal{S}, \Delta)$ is also a feasible solution to Problem 3.4 for any vertex $v \in V$ with neighborhood bounding list configuration $(\mathcal{L}(u))_{u \in \Gamma(v)}$ satisfy $|\Gamma(v)| = \Delta$ and $\mathcal{S} = \bigcup_{u \in \Gamma(v)} \mathcal{L}(u)$. Indeed, for any $\omega \in \bigotimes_{u \in \Gamma(v)} \mathcal{L}(u)$, we have

$$\{\omega_u : u \in \Gamma(v)\} \subseteq \mathcal{S} \quad \text{and} \quad |\{\omega_u : u \in \Gamma(v)\}| \leq |\Gamma(v)| = \Delta.$$

Hence $\mathcal{P}_v \subseteq \mathcal{P}(\mathcal{S}_\Delta)$, where $\mathcal{P}_v$ denotes the family of local update distributions at v defined in (6). Consequently, any grand coupling f constructed on $\mathcal{P}(\mathcal{S}_\Delta)$ in Problem 4.1 remains a valid grand coupling on $\mathcal{P}_v$ in Problem 3.4.

The simplified formulation in Problem 4.1 is not only conceptually cleaner, but also captures the key structure underlying previous designs. In particular, the core grand couplings employed in the algorithms of Bhandari et al. [BC20] (Algorithm 3, CONTRACT) and Jain et al. [JSS21] (Algorithm 3, SEEDING) can both be viewed within the framework of Problem 4.1. Building on this connection, we present a new construction that optimally solves Problem 4.1.

4.2. Optimal design of grand coupling. In this subsection, we present an optimal construction of the grand coupling for Problem 4.1, focusing on the nontrivial case where $\Delta < |\mathcal{S}| < q$.

For a grand coupling $f : \mathcal{P} \times [0, 1] \rightarrow [q]$ that is feasible for Problem 4.1, define, for each $U \in [0, 1]$, the bounding set of possible outcomes

$$\mathcal{L}_U := \{f(p, U) : p \in \mathcal{P}\}.$$

The grand coupling f naturally induces a probability distribution $\{r_k\}_{k \geq 1}$ over the sizes of $\mathcal{L}_U$, where

$$r_k = \Pr_{U \sim \text{Uniform}[0,1]} [|\mathcal{L}_U| = k].$$

We next present a system of linear constraints on the probabilities $\{r_k\}$ which correspond to the grand couplings that are feasible for Problem 4.1.

Lemma 4.2 (Linear system for grand couplings). *Fix an instance $(\mathcal{S}, \Delta)$ for Problem 4.1, and let $\mathcal{P} = \mathcal{P}(\mathcal{S}_\Delta)$. Let $\{r_k\}_{k=1}^\Delta$ be a probability distribution on $\{1, 2, \dots, \Delta\}$ (so that $\sum_{k=1}^\Delta r_k = 1$).*

Consider the system of linear constraints

$$(10) \quad \sum_{k=1}^\Delta r_k \cdot \frac{\binom{j}{k-1}}{\binom{|\mathcal{S}|}{k-1}} \leq \frac{q - |\mathcal{S}|}{q - j}, \quad \forall 1 \leq j \leq \Delta.$$

Then the following statements hold:

(1) **(Necessity)** *Let $f : \mathcal{P} \times [0, 1] \rightarrow [q]$ be a grand coupling that is feasible for Problem 4.1 on $(\mathcal{S}, \Delta)$. The sequence of probabilities $r_k = \Pr_{U \sim \text{Uniform}[0,1]} [|\mathcal{L}_U| = k]$, for $k = 1, \dots, \Delta$, satisfies (10).*

(2) **(Sufficiency)** *Conversely, if a probability distribution $\{r_k\}_{k=1}^\Delta$ satisfies (10), then there exists a grand coupling $f : \mathcal{P} \times [0, 1] \rightarrow [q]$ that is feasible for Problem 4.1 on $(\mathcal{S}, \Delta)$ and realizes the size law*

$$\Pr_{U \sim \text{Uniform}[0,1]} [|\mathcal{L}_U| = k] = r_k, \quad k = 1, 2, \dots, \Delta.$$

Remark 4.3. *In Lemma 4.2 we restrict attention to grand couplings for which the random bounding list $\mathcal{L}_U$ always has size at most Δ . This restriction is without loss of generality: the permutation-based “unconstrained” coupling from Section 3.1.2 already produces a bounding list of size $\Delta + 1$ for every seed. Thus, any coupling that sometimes outputs size $\Delta + 1$ can be treated separately as the straightforward (unrefined) case. Consequently, when searching for non-straightforward, size-reducing couplings, it suffices to consider laws supported on $\{1, 2, \dots, \Delta\}$.*

Since the objective of Problem 4.1 is to minimize $\mathbb{E}[|\mathcal{L}_U|] = \sum_k k \cdot r_k$ where $r_k = \Pr [|\mathcal{L}_U| = k]$, we can equivalently reformulate Problem 4.1 as a linear program over $\{r_k\}$ subject to (10).

Theorem 4.4. *Fix integers $\Delta \geq 1$ and $q > \Delta$, and let $\mathcal{S} \subseteq [q]$ satisfy $\Delta < |\mathcal{S}| < q$. For each $j, k \in \{1, 2, \dots, \Delta\}$, define*

$$z_j(k) := \frac{\binom{j}{k-1}}{\binom{|\mathcal{S}|}{k-1}}.$$

Consider the linear program

$$(11) \quad \begin{aligned} & \text{minimize} && \sum_{k=1}^\Delta k \cdot r_k \\ & \text{subject to} && \sum_{k=1}^\Delta r_k = 1, \\ & && \sum_{k=1}^\Delta r_k \cdot z_j(k) \leq \frac{q - |\mathcal{S}|}{q - j} && j = 1, 2, \dots, \Delta, \\ & && r_k \geq 0 && k = 1, 2, \dots, \Delta. \end{aligned}$$

Let OPT_{LP} denote the optimal objective value of (11).

Then OPT_{LP} equals the minimum expected number of coupling outcomes over all grand couplings $f : \mathcal{P}(\mathcal{S}_\Delta) \times [0, 1] \rightarrow [q]$ that are feasible for Problem 4.1 on the instance $(\mathcal{S}, \Delta)$, i.e.,

$$\text{OPT}_{\text{LP}} = \min_{f: \text{a grand coupling on } \mathcal{P}(\mathcal{S}_\Delta)} \mathbb{E}_{U \sim \text{Uniform}[0,1]} [|\{f(p, U) : p \in \mathcal{P}\}|].$$

Theorem 4.4 follows directly from Lemma 4.2, since Lemma 4.2 provides both necessary and sufficient conditions for any sequence $\{r_k\}_{k=1}^\Delta$ that correspond to a feasible solution to Problem 4.1.

We next provide a proof of Lemma 4.2, addressing both the necessity and sufficiency of the linear constraints.

Proof of Lemma 4.2. We split the proof into two parts corresponding to the two statements of the lemma.

Proof of item (1). Let $f : \mathcal{P} \times [0, 1] \rightarrow [q]$ be a grand coupling that is feasible for Problem 4.1 on $(\mathcal{S}, \Delta)$. For each $U \in [0, 1]$, recall that $\mathcal{L}_U := \{f(p, U) : p \in \mathcal{P}\}$ denote the bounding set of possible outcomes of the grand coupling f for the random seed U .

For convenience, define $\mathcal{T} := [q] \setminus \mathcal{S}$. Intuitively, $\mathcal{T}$ consists of colors that are guaranteed to be unused by neighbors, as they do not appear in the neighborhood bounding list.

Given any bounding list $\mathcal{L}_U$, we can always decompose it into the part contained in $\mathcal{S}$ and the part contained in $\mathcal{T}$, that is,

$$\mathcal{L}_U = \mathcal{L}_U^{\mathcal{S}} \cup \mathcal{L}_U^{\mathcal{T}}, \quad \mathcal{L}_U^{\mathcal{S}} \subseteq \mathcal{S}, \quad \mathcal{L}_U^{\mathcal{T}} \subseteq \mathcal{T}.$$

We first show that the component $\mathcal{L}_U^{\mathcal{T}}$ must be non-empty for any $U \in [0, 1]$. By contradiction, suppose $\mathcal{L}_U^{\mathcal{T}} = \emptyset$ for some $U \in [0, 1]$. Then we have $\mathcal{L}_U = \mathcal{L}_U^{\mathcal{S}} \subseteq \mathcal{S}$ and $|\mathcal{L}_U| \leq \Delta$. Taking $C = \mathcal{L}_U$, the distribution $\text{Uniform}([q] \setminus C) \in \mathcal{P}$ would be a valid input of Problem 4.1, but all colors in its support would lie outside the bounding list $\mathcal{L}_U$, violating the correctness of the grand coupling.

Recall that the probability sequence $\{r_k\}_{k=1}^{\Delta}$ is defined by $r_k = \Pr_{U \sim \text{Uniform}[0,1]} [|\mathcal{L}_U| = k]$, for $k = 1, \dots, \Delta$. Fix an integer $j \in \{1, 2, \dots, \Delta\}$. Let $\mathcal{S}_j^* := \{C \subseteq \mathcal{S} : |C| \leq \Delta\}$ denotes the collection of all subsets of $\mathcal{S}$ of size exactly j .

Now consider sampling a neighbor configuration $C \in \mathcal{S}_j^*$ uniformly at random. Then for any fixed $\mathcal{L}_U$ of size k , note that the subset $\mathcal{L}_U^{\mathcal{S}} \subseteq \mathcal{S}$ has size at most $k - 1$, since we have already established that $\mathcal{L}_U^{\mathcal{T}}$ is non-empty. The probability that a uniformly sampled configuration C contains all colors in $\mathcal{L}_U^{\mathcal{S}}$ is therefore lower bounded by

$$(12) \quad \Pr_{C \sim \text{Uniform } \mathcal{S}_j^*} [\mathcal{L}_U^{\mathcal{S}} \subseteq C] = \frac{\binom{j}{|\mathcal{L}_U^{\mathcal{S}}|}}{\binom{j}{|\mathcal{S}|}} \geq \frac{\binom{j}{k-1}}{\binom{j}{|\mathcal{S}|}}.$$

On the other hand, by marginal correctness of grand coupling we therefore have, for every fixed C , let $p_C := \text{Uniform}([q] \setminus C)$,

$$\Pr_{U \sim \text{Uniform}[0,1]} [f(p_C, U) \in \mathcal{T}] = \frac{q - |\mathcal{S}|}{q - j}.$$

Averaging over $C \sim \text{Uniform}(\mathcal{S}_j)$ yields the identity

$$(13) \quad \Pr_{C \sim \text{Uniform}(\mathcal{S}_j), U \sim \text{Uniform}[0,1]} [f(p_C, U) \in \mathcal{T}] = \frac{q - |\mathcal{S}|}{q - j}.$$

Finally observe that the event $\mathcal{L}_U^{\mathcal{S}} \subseteq C$ implies $f(p_C, U) \in \mathcal{T}$: if every color from $\mathcal{L}_U^{\mathcal{S}}$ is banned by C , then no color in $\mathcal{S}$ can be produced by $f(p_C, U)$. Combining (12), (13) and the above observation, using the law of total probability yields

$$\begin{aligned} \frac{q - |\mathcal{S}|}{q - j} &= \Pr_{C, U} [f(p_C, U) \in \mathcal{T}] \\ &= \sum_{k=1}^{\Delta} \Pr_U [|\mathcal{L}_U| = k] \cdot \Pr_{C, U} [f(p_C, U) \in \mathcal{T} \mid |\mathcal{L}_U| = k] \\ &\geq \sum_{k=1}^{\Delta} \Pr_U [|\mathcal{L}_U| = k] \cdot \Pr_{C, U} [\mathcal{L}_U^{\mathcal{S}} \subseteq C \mid |\mathcal{L}_U| = k] \\ &\geq \sum_{k=1}^{\Delta} r_k \cdot \frac{\binom{j}{k-1}}{\binom{j}{|\mathcal{S}|}}. \end{aligned}$$

Since the above inequality holds for each fixed $j \in \{1, 2, \dots, \Delta\}$, we conclude that the collection $\{r_k\}_{k=1}^{\Delta}$ satisfies exactly the family of inequalities specified in (10).

Proof of item (2). Recall that $\mathcal{S}_\Delta := \{C \subseteq \mathcal{S} : |C| \leq \Delta\}$ and $\mathcal{T} := [q] \setminus \mathcal{S}$. Let $\{r_k\}_{k=1}^\Delta$ be a probability distribution satisfies (10), and for each $C \in \mathcal{S}_\Delta$, define

$$P_C = \sum_{k=1}^\Delta r_k \cdot \frac{\binom{|C|}{k-1}}{\binom{|\mathcal{S}|}{k-1}}, \quad Q_C = \frac{q - |\mathcal{S}|}{q - |C|}.$$

By (10) we have $P_C \leq Q_C$ for every such C . (Note that $|\mathcal{S}| > \Delta$ in our setting, so $Q_C < 1$.)

We now construct a grand coupling $f : \mathcal{P} \times [0, 1] \rightarrow [q]$ that realizes the prescribed distribution $\{r_k\}$. Concretely, for each seed $U \in [0, 1]$ we deterministically extract a where

- $K \in \{1, \dots, \Delta\}$ is drawn with $\Pr[K = k] = r_k$;
- π is a uniformly random permutation of $\mathcal{S}$;
- c_0 is drawn uniformly from $\mathcal{T}$;
- $U' \sim \text{Uniform}[0, 1]$ is independent of the previous choices.

As usual, such independent randomness can be encoded measurably into a single uniform seed U .

Given an input distribution $p = \text{Uniform}([q] \setminus C) \in \mathcal{P}$ (where $C \in \mathcal{S}_\Delta$), and given the internal randomness (K, π, c_0, U') obtained from U , let $A = \{\pi_1, \dots, \pi_{K-1}\}$ denote the prefix of π of length $K - 1$. Define

$$\alpha_C = \frac{1 - Q_C}{1 - P_C} \in (0, 1].$$

We then set

$$(14) \quad f(\text{Uniform}([q] \setminus C), U) = \begin{cases} \text{the first element of } A \text{ that lies in } \mathcal{S} \setminus C, & \text{if } A \not\subseteq C \text{ and } U' < \alpha_C, \\ c_0, & \text{otherwise.} \end{cases}$$

We verify that f is a valid grand coupling and that it realizes the law $\Pr[|\mathcal{L}_U| = k] = r_k$.

First we check marginal correctness. Let X denote the random variable corresponding to the output of f above. Condition on C . By construction,

$$\Pr[A \subseteq C] = \sum_{k=1}^\Delta r_k \cdot \frac{\binom{|C|}{k-1}}{\binom{|\mathcal{S}|}{k-1}} = P_C,$$

hence $\Pr[A \not\subseteq C] = 1 - P_C$. The algorithm outputs a color in $\mathcal{T}$ exactly when either $A \subseteq C$ (in which case we always output c_0) or when $A \not\subseteq C$ but $U' \geq \alpha_C$. Therefore

$$\begin{aligned} \Pr[X \in \mathcal{T}] &= \Pr[A \subseteq C] + \Pr[A \not\subseteq C] \cdot (1 - \alpha_C) \\ &= P_C + (1 - P_C) \left(1 - \frac{1 - Q_C}{1 - P_C}\right) = Q_C. \end{aligned}$$

When the output falls in $\mathcal{T}$ it equals the independent uniform color c_0 , so conditional on this event the distribution over $\mathcal{T}$ is uniform; consequently each color $c \in \mathcal{T}$ is output with marginal probability

$$\Pr[X = c] = \frac{Q_C}{|\mathcal{T}|} = \frac{1}{q - |C|}.$$

It remains to check the colors in $\mathcal{S} \setminus C$. The event that the output lies in $\mathcal{S} \setminus C$ occurs exactly when $A \not\subseteq C$ and $U' < \alpha_C$. Conditioned on $A \not\subseteq C$ and $K = k$, the decoder returns the first element of A that lies in $\mathcal{S} \setminus C$. By symmetry of π , for fixed k this "first-not-in- C " element is uniformly distributed over $\mathcal{S} \setminus C$ (whenever $A \not\subseteq C$). Therefore each color $c \in \mathcal{S} \setminus C$ has conditional probability $1/(|\mathcal{S}| - |C|)$ given that the output is in $\mathcal{S} \setminus C$. Since $\Pr[X \in \mathcal{S} \setminus C] = 1 - Q_C$, we deduce that for every $c \in \mathcal{S} \setminus C$

$$\Pr[X = c] = (1 - Q_C) \cdot \frac{1}{|\mathcal{S}| - |C|} = \frac{1}{q - |C|}.$$

Combining with the uniformity on $\mathcal{T}$, we conclude that the distribution of X equals $\text{Uniform}([q] \setminus C)$, as required. Hence f is a valid grand coupling.

Finally we verify that the coupling indeed realizes the prescribed size distribution. Recall that the quadruple (K, π, c_0, U') is deterministically decoded from the seed $U \sim \text{Uniform}[0, 1]$, and $A = \{\pi_1, \dots, \pi_{K-1}\}$ denote the prefix of π of length $K - 1$. By the construction of f (see (14)), for every choice of set $C \subseteq \mathcal{S}$ with $|C| \leq \Delta$ the output of the decoder is either c_0 or the first element of A that

does not lie in C . Consequently, the image of $\mathcal{P}$ under $f(\cdot, U)$ is contained in $A \cup \{c_0\}$. Moreover, it can be easily verified that for each color in $A \cup \{c_0\}$ there exists some set C for which the decoder outputs that color, so every element of $A \cup \{c_0\}$ is actually attainable. Therefore for this fixed seed U we have

$$\mathcal{L}_U = \{f(p, U) : p \in \mathcal{P}\} = A \cup \{c_0\},$$

and hence $|\mathcal{L}_U| = |A| + 1 = K$.

Since K was drawn according to $\Pr[K = k] = r_k$, it follows that $\Pr[|\mathcal{L}_U| = k] = \Pr[K = k] = r_k$ for every $k = 1, 2, \dots, \Delta$, as required.

This completes the construction and the proof of Item (2). $\square$

Solving the linear program in (11). We now discuss how to approach the solution of the linear program (11), which characterizes the optimal grand coupling via the distribution $\{r_k\}_{k=1}^\Delta$.

We begin by examining a relaxed version of the program that retains only the constraint corresponding to $j = \Delta$ in (11). Set

$$z(k) = z_\Delta(k) = \frac{\binom{\Delta}{k-1}}{\binom{|\mathcal{S}|}{k-1}}, \quad w = \frac{q - |\mathcal{S}|}{q - \Delta},$$

and consider the reduced linear program

$$(15) \quad \begin{aligned} & \text{minimize} && \sum_{k=1}^{\Delta} k \cdot r_k \\ & \text{subject to} && \sum_{k=1}^{\Delta} r_k = 1, \\ & && \sum_{k=1}^{\Delta} r_k \cdot z(k) \leq w, \\ & && r_k \geq 0 \quad (k = 1, \dots, \Delta). \end{aligned}$$

The function $z(k)$ defined above is convex in k for fixed Δ and $|\mathcal{S}| > \Delta$. By standard extremal properties of linear programs with a single convex moment constraint (see, e.g., [HLP52] and related moment-LP arguments), any optimal solution of (15) is supported on at most two consecutive values of k . In other words, there exists an index $1 < i \leq \Delta$ such that the optimal distribution $\{r_k\}$ has $r_k = 0$ for all $k \notin \{i-1, i\}$. Consequently the optimal $\{r_k\}$ (when it exists) can be constructed as follows. Let $i \in \{2, \dots, \Delta\}$ be the smallest index such that $z(i) \leq w$. If such an i exists, set

$$(16) \quad r_{i-1} = \frac{w - z(i)}{z(i-1) - z(i)}, \quad r_i = \frac{z(i-1) - w}{z(i-1) - z(i)},$$

with all other $r_k = 0$. By construction $r_{i-1}, r_i \geq 0$, $r_{i-1} + r_i = 1$, and this choice satisfies $\sum_k r_k \cdot z(k) = w$, hence is optimal for (15).

Remark 4.5. *The program (15) is a relaxation of the original LP (11). In our subsequent analysis we work with the relaxed program (15) because it admits the explicit solution (16). In typical parameter regimes of interest (e.g. when $q = 2\Delta$), one can verify that satisfying the $j = \Delta$ constraint (i.e. the single constraint used in the relaxed program) generally implies the remaining constraints for $j = 1, 2, \dots, \Delta - 1$. Thus the relaxation is non-pathological for the regimes we care about, and the explicit solution (16) yields the optimal (or near-optimal) distribution of $|\mathcal{L}_U|$ in those settings.*

Combining the explicit probabilities in (16) with the constructive proof in Lemma 4.2 thus yields an explicit optimal grand coupling for Problem 4.1 under the stated assumptions.

4.3. Constructions of grand couplings. In this subsection, we present several concrete constructions of grand couplings that serve as building blocks for the CFTP algorithm stated in Theorem 1.2.

4.3.1. *Grand coupling: SEEDING.* The grand coupling *SEEDING* refers to the optimal construction introduced in Section 4.2. This coupling is the key to our improvement of the coloring threshold, reducing from the previous bound $q > (8/3 + o(1)) \Delta$ [JSS21] to the nearly optimal bound $q > (2.5 + o(1)) \Delta$.

Lemma 4.6 (Seeding coupling). *Fix any vertex $v \in V$ and neighborhood bounding lists $(\mathcal{L}(u))_{u \in \Gamma(v)}$. Let $\mathcal{P}_v$ denote the set of all possible local update distributions at v in the Glauber dynamics on q -colorings, given the bounding lists $(\mathcal{L}(u))_{u \in \Gamma(v)}$ on v 's neighbors. Formally,*

$$(17) \quad \mathcal{P}_v = \left\{ \text{Uniform}([q] \setminus \omega) : \omega \in \bigotimes_{u \in \Gamma(v)} \mathcal{L}(u) \right\}$$

Assume

$$q \geq \frac{7}{3} \Delta \quad \text{and} \quad |\mathcal{S}| \leq 2\Delta, \quad \text{where } \mathcal{S} := \bigcup_{u \in \Gamma(v)} \mathcal{L}(u).$$

Then there exists a grand coupling $\text{SEEDING}_v : \mathcal{P}_v \times [0, 1] \rightarrow [q]$ satisfying:

- (a) (**Bounding list reduction**) For random seed $U \in [0, 1]$, let $\mathcal{L}_U(v) = \{\text{SEEDING}_v(p, U) : p \in \mathcal{P}_v\}$ denote the set of possible coupling outcomes. Then $|\mathcal{L}_U(v)| \in \{2, 3\}$, and

$$\mathbb{E}_{U \sim \text{Uniform}[0, 1]}[|\mathcal{L}_U(v)|] \leq 2 + \frac{(|\mathcal{S}| + \Delta - q)(|\mathcal{S}| - 1)}{(q - \Delta) \Delta}.$$

- (b) (**Efficiency**) The grand coupling SEEDING_v can be implemented so that each update of the bounding list can be performed in expected time $O(\Delta \log \Delta)$.

Proof. Recall that solving Problem 4.1 for the input $(\mathcal{S}, \Delta)$ yields a grand coupling that is valid for any neighborhood bounding lists with union equal to $\mathcal{S}$. Hence it suffices to produce an admissible probability law $\{r_k\}_{k=1}^\Delta$ satisfying the feasibility constraints of Lemma 4.2. The lemma then guarantees existence of a corresponding grand coupling realizing $\Pr[|\mathcal{L}_U(v)| = k] = r_k$ for every k .

We take $\{r_k\}$ to be the two-point law obtained from the explicit formula (16) with index $i = 3$. Concretely define

$$r_3 = \frac{(|\mathcal{S}| + \Delta - q)(|\mathcal{S}| - 1)}{(q - \Delta) \Delta}, \quad r_2 = 1 - r_3, \quad r_k = 0 \quad (k \notin \{2, 3\}).$$

Under the parameter restrictions of the lemma ($q \geq 7/3 \Delta$ and $|\mathcal{S}| \leq 2\Delta$) one checks routinely that $r_2, r_3 \in [0, 1]$. It can also be easily verified that the above choice of $\{r_k\}$ satisfies the linear constraint (10) throughout this regime, as earlier mentioned in Remark 4.5.

By Lemma 4.2, there exists a valid grand coupling f whose explicit construction was given in the proof of that lemma. For $U \sim \text{Uniform}[0, 1]$, this coupling satisfies $\Pr[|\mathcal{L}_U(v)| = k] = r_k$ for every k . We take precisely this constructed coupling as our SEEDING_v .

Proof of item a. Because $\Pr[|\mathcal{L}_U(v)| = k] = r_k$ and only r_2, r_3 are nonzero, the size $|\mathcal{L}_U(v)|$ takes values in $\{2, 3\}$ and its expectation equals

$$\mathbb{E}[|\mathcal{L}_U(v)|] = 2 \cdot r_2 + 3 \cdot r_3 = 2 + r_3 = 2 + \frac{(|\mathcal{S}| + \Delta - q)(|\mathcal{S}| - 1)}{(q - \Delta) \Delta},$$

which establishes the claimed bound.

Proof of item b. Since $|\mathcal{S}| \leq 2\Delta$, constructing $\mathcal{S} = \bigcup_{u \in \Gamma(v)} \mathcal{L}(u)$ requires $O(\Delta \log \Delta)$ time. Within SEEDING_v , apart from generating a random permutation of $\mathcal{S}$ in $O(\Delta)$ time, all encoding and decoding operations take constant time. Hence, the overall computational complexity is $O(\Delta \log \Delta)$. $\square$

4.3.2. *Grand Coupling: COMPRESS and DISJOINT.* Next, we present two canonical constructions of grand couplings, referred to as *COMPRESS* and *DISJOINT*, which serve as fundamental components in the CFTP algorithm discussed later.

The *COMPRESS* and *DISJOINT* algorithms were originally introduced in [BC20] and [JSS21], respectively. Both algorithms can be naturally interpreted as specific instances of grand couplings via the bounding chains within our framework.

In particular, the COMPRESS coupling can be viewed as a refinement of Huber’s permutation-based grand coupling [Hub98] discussed in Section 3.1.2, optimized for the bounding-chain setting.

Lemma 4.7 (Compress coupling). *Fix any vertex $v \in V$ and neighborhood bounding lists $(\mathcal{L}(u))_{u \in \Gamma(v)}$. Let $\mathcal{P}_v$ be defined as (17). Assume $q \geq \Delta + 1$. For any fixed subset $A \subseteq [q]$ with $|A| = \Delta$, called a reference color set, there exists a grand coupling $\text{COMPRESS}_v^A : \mathcal{P}_v \times [0, 1] \rightarrow [q]$ satisfying:*

(a) (**Bounding property**) *For random seed $U \sim \text{Uniform}[0, 1]$, the set of possible outcomes is*

$$\{\text{COMPRESS}_v^A(p, U) : p \in \mathcal{P}_v\} = A \cup \{c\},$$

where c is distributed uniformly in $[q] \setminus A$ (with the randomness induced by U).

(b) (**Efficiency**) *The grand coupling COMPRESS_v^A can be implemented so that each update of the bounding list can be performed in expected time $O(\Delta \log \Delta)$.*

Proof. We encode the internal randomness of the COMPRESS procedure into a single random seed $U \in [0, 1]$. Concretely, for each seed U we deterministically extract a triple (π, x', U') where

- π is a uniformly random permutation of A ,
- x' is a uniformly random element of $[q] \setminus A$,
- U' is an independent uniform draw from $[0, 1]$.

For each feasible neighbor configuration $\omega \in \text{Conf}(v)$ —with induced forbidden set $\Gamma_\omega(v)$ and corresponding distribution $p_\omega = \text{Uniform}([q] \setminus \Gamma_\omega(v)) \in \mathcal{P}_v$ —define

$$\text{COMPRESS}_v^A(p_\omega, U) = \begin{cases} x', & \text{if } x' \notin \Gamma_\omega(v) \text{ and } U' \leq \frac{q - \Delta}{q - |\Gamma_\omega(v)|}, \\ \text{the first } y \text{ in } \pi \text{ with } y \notin \Gamma_\omega(v), & \text{otherwise.} \end{cases}$$

Grand coupling property (correctness). Fix a configuration ω and let $\Gamma = \Gamma_\omega(v)$. Conditional on (π, x') (or marginalizing over them), the above rule returns x' with probability

$$\Pr[x' \notin \Gamma] \cdot \Pr\left[U' \leq \frac{q - \Delta}{q - |\Gamma|}\right] = \frac{1}{q - \Delta} \cdot \frac{q - \Delta}{q - |\Gamma|} = \frac{1}{q - |\Gamma|}.$$

If the first branch is not taken, the decoder returns the first element of π outside Γ ; since π is a uniform permutation of A , each color in $[q] \setminus \Gamma$ (that lies within A) is equally likely to be the first available, and its probability is again $1/(q - |\Gamma|)$. Hence marginally $\text{COMPRESS}_v^A(p_\omega, U) \sim \text{Uniform}([q] \setminus \Gamma)$, verifying that COMPRESS_v^A is a valid grand coupling in the sense of Definition 3.2.

Proof of item a. For any fixed seed (equivalently fixed (π, x', U')), the image of $\mathcal{P}_v$ under $\text{COMPRESS}_v^A(\cdot, U)$ is contained in $A \cup \{x'\}$. Conversely, x' itself is attainable (for those configurations with $x' \notin \Gamma$ and satisfying the U' -condition), and every element of A is attainable (for configurations that force the decoder to use π). Thus for each seed the image equals $A \cup \{c\}$ for some $c \in [q] \setminus A$, and the distribution of c over seeds is uniform on $[q] \setminus A$.

Proof of item b. Decoding the seed (recovering (π, x', U')) and forming the predicted set $A \cup \{x'\}$ both take $O(\Delta)$ time. Given a concrete configuration ω , evaluating $\text{COMPRESS}_v^A(p_\omega, U)$ requires scanning π until the first element not in $\Gamma_\omega(v)$, which can be implemented in expected $O(\Delta \log \Delta)$ time using standard hashing or balanced lookup structures. This yields the claimed efficiency bound.

This completes the proof of Lemma 4.7. $\square$

The DISJOINT coupling corresponds to a key step in [JSS21], enabling the reduction of all bounding list sizes from 2 to 1 in the regime where $q > 2.5\Delta$.

Lemma 4.8. (Disjoint coupling [JSS21, Lemma 4.3]) Fix any vertex $v \in V$ and neighborhood bounding lists $(\mathcal{L}(u))_{u \in \Gamma(v)}$. Let $\mathcal{P}_v$ be defined as (17). Define

$$\begin{aligned}\mathcal{S} &:= \bigcup_{u \in \Gamma(v)} \mathcal{L}(u), \\ \mathcal{Q} &:= \bigcup_{u \in \Gamma(v): |\mathcal{L}(u)|=1} \mathcal{L}(u), \\ \mathcal{D} &:= \bigcup_{\substack{u \in \Gamma(v): \\ \forall w \in \Gamma(v), \mathcal{L}(u) \cup \mathcal{L}(w) = \emptyset}} \mathcal{L}(u).\end{aligned}$$

Assume $q \geq 2.5\Delta$ and $|\mathcal{L}(u)| \leq 2$ for all $u \in \Gamma(v)$. Then, there exists a grand coupling $\text{DisJOINT}_v : \mathcal{P}_v \times [0, 1] \rightarrow [q]$ satisfying:

- (a) (**Bounding list reduction**) For random seed $x \in [0, 1]$, let $\mathcal{L}_U(v) = \{\text{DisJOINT}_v(p, x) : p \in \mathcal{P}_v\}$ denote the set of possible coupling outcomes. Then $|\mathcal{L}_U(v)| \in \{1, 2\}$, and

$$\Pr_{U \sim \text{Uniform}[0, 1]} [|\mathcal{L}_U(v)| = 1] \geq 1 - \frac{|\mathcal{S}| - |\mathcal{Q}|}{q - \Delta} + \frac{|\mathcal{D}|/2}{q - |\mathcal{Q}| - |\mathcal{D}|/2}.$$

- (b) (**Efficiency**) The grand coupling DisJOINT_v can be implemented so that each update of the bounding list can be performed in expected time $O(\Delta \log \Delta)$.

Remark 4.9. The disjoint coupling builds on Jerrum’s pairing coupling analysis of Glauber dynamics [Jer95], refined under the bounding-list framework. Its key idea is to perfectly couple the pair of colors within the same independent bounding set, making them to be regarded as single color, effectively reducing the neighborhood size to $E = |\mathcal{S}| - \frac{1}{2}|\mathcal{D}|$. While [BC20] treated $E = |\mathcal{S}|$ and thus required $q > 3\Delta$, the refinement of [JSS21] showed that when $|\mathcal{L}(u)| \leq 2$, one has $E \leq 1.5\Delta$, so $q > 2.5\Delta$ ensures contraction. Our theorem 3.5 show that this 2.5Δ threshold is tight, making the disjoint coupling optimal under these assumptions.

5. CFTP FOR q -COLORINGS WITH $q > 2.5\Delta$

In this section, we present the construction of a perfect sampler for proper q -colorings within the classical *Coupling From The Past* (CFTP) framework, implemented using the grand couplings developed in Section 4. We show that this construction gives an efficient CFTP algorithm for exact sampling from the uniform distribution over proper q -colorings, under the condition $\Delta > (2.5 + \eta)\Delta$ where $\eta = 2\sqrt{(\log \Delta + 1)/\Delta} = o(1)$, thus completing the proof of Theorem 1.2.

5.1. Algorithm overview. Let Ω be a finite configuration space endowed with the uniform measure μ_Ω , and let $F : \Omega \times [0, 1] \rightarrow \Omega$ be a randomized *update operator* driven by a random seed $U \in [0, 1]$. Associated with F is a *coalescence predicate* $\Phi_F : [0, 1] \rightarrow \Omega \cup \{\perp\}$, which detects whether the update $F(\cdot, U)$ maps all initial configurations in Ω to a common state. Specifically, $\Phi_F(U) = \omega^* \neq \perp$ indicates that all trajectories under $F(\cdot, U)$ coalesce at the same configuration $\omega^* \in \Omega$.

Algorithm 1: CFTP Sampler

Input: Random seed $U \sim \text{Uniform}[0, 1]$, interpreted as i.i.d. $U_{-1}, U_{-2}, \dots \sim \text{Uniform}[0, 1]$; update operator $F : \Omega \times [0, 1] \rightarrow \Omega$; coalescence predicate $\Phi_F : [0, 1] \rightarrow \Omega \cup \{\perp\}$.

Output: A configuration ω^* from Ω .

```

1 for  $t = 1, 2, \dots$  do
2   Generate an independent  $U_{-t} \sim \text{Uniform}[0, 1]$  using the random seed  $U$ ;
3   if  $\Phi_F(U_{-t}) \neq \perp$  then
4     return  $\omega^* = F(\cdot, U_{-1}) \circ F(\cdot, U_{-2}) \circ \dots \circ F(\cdot, U_{-t+1}) \circ \Phi_F(U_{-t})$ 
```

Lemma 5.1. Let $F : \Omega \times [0, 1] \rightarrow \Omega$ and $\Phi_F : [0, 1] \rightarrow \Omega \cup \{\perp\}$ satisfy:

- (a) For independent $\omega \sim \mu_\Omega$ and $U \sim \text{Uniform}[0, 1]$, the random variable $F(\omega, U)$ is distributed according to μ_Ω .

- (b) If $\Phi_F(U) \neq \perp$, then $F(\cdot, U)$ is constant over Ω , and $\Phi_F(U)$ equals its unique image.
- (c) $\Pr_{U \sim \text{Uniform}[0,1]} [\Phi_F(U) \neq \perp] \geq 1/2$.
- (d) The predicate $\Phi_F(U)$ can be evaluated in expected time T_1 , and $F(\omega, U)$ in expected time T_2 .

Then Algorithm 1 terminates in expected time $O(T_1 + T_2)$ and outputs an exact sample from μ_Ω , with exponentially decaying tail bounds on its runtime.

Proof. Let $\omega_0 \sim \mu_\Omega$ be a initial configuration drawn independently to U , and let

$$\omega_t = F(\cdot, U_{-1}) \circ F(\cdot, U_{-2}) \circ \cdots \circ F(\cdot, U_{-t})(\omega_0)$$

denote the configuration obtained after applying t grand couplings from the past. By condition (a), the distribution of ω_t remains uniform over Ω for all t .

Define ω_t^* to be the output of Algorithm 1 truncated at step t (and $\omega_t^* = \perp$ if algorithm does not halt). Let μ_{ω_t} and $\mu_{\omega_t^*}$ denote the distributions of ω_t and ω_t^* , respectively. Under conditions (b) and (c),

$$d_{\text{TV}}(\mu_{\omega_t}, \mu_{\omega_t^*}) = \Pr_U[\omega_t \neq \omega_t^*] = \Pr_U[\omega_t^* = \perp] = \Pr_U[\forall i \leq t : \Phi_F(U_{-i}) = \perp] \leq 2^{-t},$$

by the independence of the seeds $U_{-1}, U_{-2}, \dots, U_{-t}$.

Taking the limit as $t \rightarrow \infty$ yields

$$d_{\text{TV}}(\mu_\Omega, \mu_{\omega^*}) = \lim_{t \rightarrow \infty} d_{\text{TV}}(\mu_{\omega_t}, \mu_{\omega_t^*}) = 0,$$

so the output ω^* is distributed exactly according to μ_Ω .

Since each iteration halts independently with probability at least $p \geq 1/2$, the number of iterations follows a geometric distribution with mean $O(1)$. Therefore, the total expected runtime is $O(T_1 + T_2)$, and the probability that the algorithm exceeds t iterations decays exponentially. $\square$

The above lemma provides a general CFTP framework applicable to any update operator satisfying conditions (a)–(d). We next instantiate this framework for the Glauber dynamics on q -colorings.

Remark 5.2. In our setting, the randomized update operator F in Lemma 5.1 is realized as a finite composition of single-site grand couplings for the Glauber dynamics (cf. Definition 3.2). Because the uniform measure μ_Ω is stationary for the Glauber dynamics, condition (a) is automatically satisfied. Moreover, with the bounding chains, the corresponding coalescence predicate Φ_F inherently satisfies condition (b). Consequently, to prove Theorem 1.2, it suffices to construct an explicit family of such grand couplings that ensure the coalescence probability in (c) and the efficiency guarantees in (d).

5.2. Global grand couplings from local updates. We now explicitly construct the global grand coupling $F^* = F(\cdot, U)$ on the configuration space Ω and its associated coalescence predicate $\Phi_F(U)$, by composing local grand couplings corresponding to single-site updates.

5.2.1. Representing global coupling via composition of local couplings. Recall from Section 4.3 and Definition 3.2 that for each vertex $v \in V$, we have defined several types of local grand couplings $f_v : \mathcal{P}_v \times [0, 1] \rightarrow [q]$ corresponding to single-site updates in Glauber dynamics, including SEEDING, COMPRESS, and DISJOINT. We now describe how these local updates are composed into a global coupling according to a carefully designed schedule.

Throughout this construction, the algorithm has access to a global random seed $U \sim \text{Uniform}[0, 1]$, where each bit of U is sampled i.i.d. uniformly so that all randomness is derived from this single source.

In addition, the construction maintains two global data structures:

- A **bounding list** $\mathcal{L} = (\mathcal{L}(v))_{v \in V}$, where each $\mathcal{L}(v) \subseteq [q]$ records the current set of possible color at vertex v ;
- A **composition list** F^* , which stores the sequence of local updates used to construct the overall mapping $F(\cdot, U)$. Each entry of F^* takes the form (v, f_v, U_v) , where v is the vertex being updated, f_v is the local grand coupling, and $U_v \in [0, 1]$ is the independent random seed used in that update (generated from the global seed U).

The following subroutine $\text{UPDATE}(v, f)$ modifies the bounding list $\mathcal{L}$ at v and appends a new entry to the composition list F^* , given a vertex $v \in V$ and a local grand coupling rule $f = (f_u)_{u \in V}$ specifying at each vertex u a local update function $f_u : \mathcal{P}_u \times [0, 1] \rightarrow [q]$, as in Definition 3.2.

Algorithm 2: $\text{UPDATE}(v, f)$

Input: Vertex $v \in V$; local grand coupling rule $f = (f_u)_{u \in V}$.

Global Variable: Bounding list $\mathcal{L}$; composition list F^* ; global random seed $U \sim \text{Uniform}[0, 1]$.

- 1 Generate an independent $U' \sim \text{Uniform}[0, 1]$ using the random seed U ;
 - 2 $\mathcal{L}(v) \leftarrow \{f_v(p, U') : p \in \mathcal{P}_v\}$;
 - 3 Append tuple (v, f, U') to the end of F^* ;
-

Each call to UPDATE modifies only the bounding list $\mathcal{L}(v)$ at vertex v , and appends one new record (v, f_v, U_v) representing the local update $f_v(\cdot, U_v)$ to F^* . Intuitively, each such tuple in F^* specifies a single atomic update applied to the current bounding chain.

The resulting global coupling $F^* : \Omega \rightarrow \Omega$ is the composition of all local updates in the order they appear in the list. Given an initial configuration $\omega \in \Omega$, its image under F^* can be evaluated as follows.

Algorithm 3: Evaluation of $F^*(\omega)$ for initial configuration ω

1 a

Input: Configuration $\omega \in \Omega$; composition list $F^* = [(v_1, f_1, U_1), \dots, (v_T, f_T, U_T)]$.

Output: Updated configuration $F^*(\omega)$.

2 **for** $i = 1$ **to** T **do**

3 $\omega_{v_i} \leftarrow f_{v_i}(p_{v_i}^\omega, U_i)$, where $p_{v_i}^\omega = \text{Uniform}([q] \setminus \omega_{\Gamma(v_i)})$ is the Glauber dynamics update;

4 **return** ω ;

5.2.2. *Neighborhood maintenance via CLEANUP.* We define an auxiliary routine $\text{CLEANUP}(v, P)$ that applies a sequence of COMPRESS-type updates around v relative to a reference color set $A \subseteq [q]$ built from the bounding lists of a preserved set $P \subseteq V$. Intuitively, this aligns the neighborhoods not currently being updated so that subsequent UPDATE calls meet the required invariants.

Algorithm 4: $\text{CLEANUP}(v, P)$

Input: Vertex v to be cleaned; subset $P \subseteq V$ of vertices to preserve.

Global Variable: Bounding list $\mathcal{L}$; composition list F^* ; global random seed $U \sim \text{Uniform}[0, 1]$.

1 $A \leftarrow \text{GREEDY}(\mathcal{L}_P = (\mathcal{L}(w))_{w \in P})$;

2 **foreach** $w \in \Gamma(v) \setminus P$ **do**

3 $\text{UPDATE}(w, \text{COMPRESS}^A)$;

Remark 5.3 (The GREEDY subroutine in Algorithm 4). *The reference set A used in CLEANUP is chosen greedily based on the current bounding information $\mathcal{L}_P$: the idea is to include as many colors as possible from the target color sets until reaching size Δ ; and if the total size of the targets is less than Δ , we include all of them and fill the remaining slots with arbitrary colors. This ensures $|A| = \Delta$ while maximizing coverage over the currently active neighborhood.*

The specific priority rule of this greedy selection may vary across different stages of the algorithm, as the construction of A will be adapted to the particular coupling objective of each stage. The detailed forms of these greedy strategies will be provided later in the proofs corresponding to each phase of the algorithm.

Remark 5.4. *It is important to note that while the vertex v in each UPDATE or CLEANUP call must be selected independently of the current configuration, the auxiliary set A may depend adaptively on the current state of the bounding chain. The reason is that choosing A only determines which specific member of the family of grand couplings $\{\text{COMPRESS}^A : A \subseteq [q], |A| = \Delta\}$ is applied; it does not modify the underlying Glauber dynamics itself. In other words, the Glauber update rule remains unchanged, and adaptively choosing A merely specifies which valid grand coupling realization to use at that step. Since grand couplings are by definition allowed to be adaptively chosen, this adaptivity does not introduce bias nor affect correctness.*

5.2.3. *Construction of grand coupling and coalescence predicate.* We now define the global coupling $F(\cdot, U)$ by composing the atomic local updates according to a two-phase schedule. The resulting operator $F : \Omega \rightarrow \Omega$ is represented by the composition list F^* , which records the ordered sequence of local updates together with their corresponding random seeds, as described in Algorithm 3. For convenience, we write $F^* \equiv F(\cdot, U)$ to emphasize that the composition list fully specifies the coupling induced by the global seed U .

At the end, if every bounding list $\mathcal{L}(v)$ for $v \in V$ reduces to a singleton, a coalescence is detected and we set $\Phi_F(U)$ to be the corresponding unique configuration. Otherwise, we set $\Phi_F(U) = \perp$.

Algorithm 5: Construction of $F(\cdot, U)$ and evaluation of $\Phi_F(U)$

Input: Graph $G(V, E)$ with maximum degree Δ ; $q \geq 2$; global random seed $U \sim \text{Uniform}[0, 1]$.

Output: Grand coupling $F(\cdot, U)$, represented as a composition list $F^* \equiv F(\cdot, U)$; outcome of coalescence predicate $\Phi_F(U)$.

Initialization:

- 1 Initialize the global variables $\mathcal{L}(v) \leftarrow [q]$ for all $v \in V$, and $F^* \leftarrow ()$ to be an empty list;
- 2 $S \leftarrow \text{VERTEXPARTITION}(G, \Delta)$; // Provided by Lemma 5.7

Phase I (Seeding Phase):

- 3 **foreach** $v_i \in S$ *in fixed order* **do**
- 4 $\text{CLEANUP}(v_i, \{v_1, \dots, v_{i-1}\})$;
- 5 $\text{UPDATE}(v, \text{SEEDING})$;
- 6 **for** $t = 1$ **to** $T_1 = 5|S| \log |S|$ **do**
- 7 Choose $v \in S$ uniformly at random;
- 8 $\text{CLEANUP}(v, S)$;
- 9 $\text{UPDATE}(v, \text{SEEDING})$

Phase II (Converting and Drifting Phase):

- 10 **foreach** $v_i \in V \setminus S$ *in fixed order* **do**
- 11 $\text{CLEANUP}(v_i, S \cup \{v_1, \dots, v_{i-1}\})$;
- 12 $\text{UPDATE}(v, \text{DISJOINT})$;
- 13 **for** $t = 1$ **to** $T_2 = \frac{2(q-\Delta)n \log n}{q-2.5\Delta}$ **do**
- 14 Choose $v \in V$ uniformly at random;
- 15 $\text{UPDATE}(v, \text{DISJOINT})$;

Check coalescence:

- 16 **if** $|\mathcal{L}(v)| = 1$ *for all* $v \in V$ **then**
 - 17 $\Phi_F(U) \leftarrow$ unique configuration $\omega \in \bigotimes_{v \in V} \mathcal{L}(v)$;
 - 18 **else**
 - 19 $\Phi_F(U) \leftarrow \perp$;
 - 20 **return** $(F^*, \Phi_F(U))$;
-

Lemma 5.5 (Efficiency and coalescence). *Suppose $q > (2.5 + \eta)\Delta$ with $\eta = 2\sqrt{(\log \Delta + 1)/\Delta}$. Then Algorithm 5 returns a $(F^*, \Phi_F(U))$ in time $\tilde{O}(n\Delta^2)$, and it holds $\Pr[\Phi_F(U) \neq \perp] \geq 1/2$.*

The proof of lemma 5.5 is deferred to the next subsections.

Lemma 5.6 (Evaluation cost). *Given the composition list F^* returned by Algorithm 5, the mapping $F^*(\omega) = F(\omega, U)$ can be evaluated by Algorithm 3 for any configuration $\omega \in \Omega$ in time $\tilde{O}(n\Delta^2)$.*

Proof. By Lemmas 4.6 to 4.8, each local update costs $O(\Delta \log \Delta)$ in time. Since algorithm 5 performs at most $O(n\Delta \log n)$ updates across both phases, the total time cost is $O(n\Delta^2 \log n \log \Delta)$. $\square$

Combining lemmas 5.5 and 5.6 with lemma 5.1, we conclude that the perfect sampling algorithm terminates in expected time $\tilde{O}(n\Delta^2)$ and outputs an exact sample from the uniform distribution μ_Ω . This establishes theorem 1.2.

5.3. Vertex partition via algorithmic LLL. As a preprocessing step for the global coupling construction in Algorithm 5, we partition the vertex set V into a *seeding set* S and its complement $V \setminus S$. The goal is to ensure that each vertex has a well-balanced neighborhood across S and $V \setminus S$. The seeding set S is used in Phase I of Algorithm 5 to guarantee that the neighborhoods involved in local updates remain sufficiently sparse, which in turn enables efficient list reduction and coalescence guarantees in subsequent phases.

The following lemma shows that such a partition always exists and can be constructed efficiently.

Lemma 5.7. *Let $G = (V, E)$ be a graph with maximum degree Δ . Then for $\eta = 2\sqrt{(\log \Delta + 1)/\Delta}$, there exists a subset $S \subseteq V$ such that for every vertex $v \in V$,*

$$|\Gamma(v) \cap S| \leq \frac{1}{2}\Delta, \quad |\Gamma(v) \cap (V \setminus S)| \leq \left(\frac{1}{2} + \eta\right)\Delta.$$

Moreover, such a set S can be computed in expected time $O(n\Delta)$.

Proof. We construct S by including each vertex $v \in V$ independently with probability $p_0 = \frac{1}{2} - \frac{\eta}{2}$. For each $u \in V$, let I_u be the indicator for the event $u \in S$, and for any vertex v , define

$$I(v) := \sum_{u \in \Gamma(v)} I_u, \quad \mathbb{E}[I(v)] = p_0 \deg(v) = \left(\frac{1}{2} - \frac{\eta}{2}\right) \deg(v).$$

We say that a vertex v is *violating* if either $|\Gamma(v) \cap S| > \frac{1}{2}\Delta$ or $|\Gamma(v) \cap (V \setminus S)| > \left(\frac{1}{2} + \eta\right)\Delta$. Both violations correspond to $|I(v) - \mathbb{E}[I(v)]| > \frac{\eta}{2}\Delta$. By a Chernoff bound, for all v ,

$$\Pr[|I(v) - \mathbb{E}[I(v)]| > \frac{\eta}{2}\Delta] \leq 2 \exp\left(-\frac{(\eta\Delta)^2}{2 \deg(v)}\right) \leq \frac{2}{e^2 \Delta^2}$$

since $\deg(v) \leq \Delta$.

Let $p = 2/(e^2 \Delta^2)$ be the failure probability per vertex, and note that each bad event (i.e., vertex v violating the bounds) depends only on the random choices of vertices in $\Gamma(v)$, and thus at most $d := \Delta$ other events. The symmetric Lovász Local Lemma (LLL) condition holds because

$$ep(d+1) \leq \frac{2e(\Delta+1)}{e^2 \Delta^2} < 1.$$

Therefore, with positive probability, no vertex violates the condition, and such a partition exists.

To obtain such a partition efficiently, we apply the Moser-Tardos resampling algorithm [MT10]. We associate with each vertex $v \in V$ a random bit $s_v \in \{0, 1\}$ indicating its membership in S (initialized to 1 independently with probability p_0), and a counter $c_v = |\Gamma(v) \cap S|$.

We maintain a queue Q of currently violating vertices. Initially, all vertices are placed into Q . While Q is nonempty, we:

- Dequeue a vertex v and check whether c_v violates the desired bounds.
- If so, we resample the bits $\{s_u : u \in \Gamma(v)\}$.
- For each $w \in \Gamma(v)$ where $u \in \Gamma(v)$, update c_w , and enqueue w into Q .

Each resampling step affects at most $O(\Delta^2)$ entries. By the analysis of Moser and Tardos, the expected number of resamplings is $O(n/\Delta)$, so the total expected runtime is $O(n\Delta)$. $\square$

5.4. Phase I: Seeding. Phase I transforms the initial bounding lists into size-two lists on the seeding set S obtained from Lemma 5.7. It consists of two stages:

- i. *Initialization stage:* a deterministic pass over S that ensures $|\mathcal{L}(v)| \in \{2, 3\}$ for every $v \in S$;
- ii. *Randomized drifting stage:* repeated applications of the SEEDING coupling to uniformly random vertices in S , which drive all lists in S to satisfy $|\mathcal{L}(v)| = 2$ with constant probability.

We first establish the deterministic initialization property and then analyze the convergence of the drifting stage.

Lemma 5.8. *Assume $q \geq (2.5 + \eta)\Delta$, and let S be the seeding set constructed in Lemma 5.7. After the initialization stage of Phase I (Lines 2–5 in Algorithm 5), every vertex $v \in S$ satisfies $|\mathcal{L}(v)| \in \{2, 3\}$.*

Proof. We prove by induction on the fixed update order of vertices in S . Assume that before processing v_i , every $v \in P = \{v_1, \dots, v_{i-1}\}$ satisfies $|\mathcal{L}(v)| \in \{2, 3\}$. The procedure $\text{CLEANUP}(v_i, P)$ chooses a $A \subseteq [q]$ of size Δ from

$$\mathcal{S}_P := \bigcup_{u \in \Gamma(v_i) \cup P} \mathcal{L}(u),$$

greedily: include as many colors as possible from $\mathcal{S}_P$, and if $|\mathcal{S}_P| < \Delta$, fill the remaining slots arbitrarily from $[q] \setminus \mathcal{S}_P$.

Next, COMPRESS^A is applied to every $w \in \Gamma(v_i) \setminus P$. By Lemma 4.7, after this step each such neighbor satisfies $\mathcal{L}(w) = A \cup \{c_w\}$ for some $c_w \in [q] \setminus A$. Hence, at this moment

$$\mathcal{S} := \bigcup_{u \in \Gamma(v_i)} \mathcal{L}(u) = (A \cup \mathcal{S}_P) \cup \{c_w : w \in \Gamma(v_i) \setminus P\}.$$

Therefore,

$$|\mathcal{S}| \leq \max(|A|, 3|\Gamma(v_i) \cap P|) + |\Gamma(v_i) \setminus P| \leq 2\Delta,$$

where the first inequality uses $|A| = \Delta$ and the inductive assumption $|\mathcal{L}(v)| \leq 3$ for all $v \in P$, and the second inequality follows from $|\Gamma(v_i)| \leq \Delta$ and $|\Gamma(v_i) \cap P| \leq \frac{1}{2}\Delta$ by Lemma 5.7.

Thus the precondition of Lemma 4.6, namely $|\mathcal{S}| < 2\Delta$, holds. Applying the SEEDING grand coupling then gives $|\mathcal{L}(v_i)| \in \{2, 3\}$, preserving the induction hypothesis. The claim follows for all $v_i \in S$. $\square$

This completes the deterministic initialization of Phase I: after processing all vertices in S , each $v \in S$ satisfies $|\mathcal{L}(v)| \in \{2, 3\}$. In the subsequent *drifting stage*, we repeatedly apply the SEEDING coupling to uniformly random vertices of S , which further reduces each list to size 2 with constant probability.

Lemma 5.9. Assume $q \geq (2.5 + \eta)\Delta$, and let S be the seeding set constructed in Lemma 5.7. Let $m = |S|$ and $T_1 = 5m \log m$. Starting from a configuration in which $|\mathcal{L}(v)| \in \{2, 3\}$ for all $v \in S$ (as ensured by Lemma 5.8), after T_1 iterations of the drifting step in Phase I (Lines 5–9 in Algorithm 5), we have $|\mathcal{L}(v)| = 2$ for all $v \in S$ with probability at least $3/4$.

Proof. Let $\mathcal{L}_t$ denote the bounding list configuration after t randomized drifting step, and define

$$V_t := \{v \in S : |\mathcal{L}_t(v)| = 3\},$$

be the set of vertices whose lists have not yet been reduced to size 2. We introduce the potential function $\Phi_t = |V_t|$ which measures the number of unresolved vertices at time t .

For each vertex $v \in S$, let

$$\mathcal{S}_t(v) := \bigcup_{u \in \Gamma(v)} \mathcal{L}_t(u)$$

denote the *slack color set* of v at iteration t , i.e., the union of bounding lists over its neighborhood at that time. By property (a) of Lemma 4.6, the probability that vertex v remains in the $|\mathcal{L}_t(v)| = 3$ state after one update is

$$p_{t,v} = \frac{(|\mathcal{S}_t(v)| + \Delta - q)(|\mathcal{S}_t(v)| - 1)}{(q - \Delta)\Delta} \leq \frac{2|\mathcal{S}_t(v)|}{q - \Delta} - 2$$

by assuming $|\mathcal{S}_t(v)| \leq 2\Delta$. The expected change in the potential at each iteration can be expressed as

$$\mathbb{E}[\Phi_{t+1} - \Phi_t \mid V_t] \leq \frac{1}{m} \left(\sum_{v \in S \setminus V_t} p_{t,v} - \sum_{v \in V_t} (1 - p_{t,v}) \right) \leq \frac{1}{m} \left(2 \left(\frac{\sum_{v \in S} |\mathcal{S}_t(v)|}{q - \Delta} - m \right) - \Phi_t \right),$$

Next, we bound the neighborhood slack sizes $|\mathcal{S}_t(v)|$ that appear in $p_{t,v}$. By the greedy construction of the reference set A in Lemma 5.8, the CLEANUP operation ensures that

$$|\mathcal{S}_t(v)| \leq \max(\Delta, 2|\Gamma(v) \cap S| + |\Gamma(v) \cap V_t|) + |\Gamma(v) \setminus S| \leq (1.5 + \eta)\Delta + |\Gamma(v) \cap V_t|$$

by applying the degree bounds from Lemma 5.7, namely $|\Gamma(v) \cap S| \leq \frac{1}{2}\Delta$ and $|\Gamma(v) \setminus S| \leq (\frac{1}{2} + \eta)\Delta$.

Summing over all $v \in S$ yields the aggregate bound

$$\sum_{v \in S} |\mathcal{S}_t(v)| \leq (1.5 + \eta)m\Delta + \sum_{v \in S} |\Gamma(v) \cap V_t| \leq (1.5 + \eta)m\Delta + \frac{1}{2}\Phi_t\Delta.$$

Plugging into the drift expression and applying $q \geq (2.5 + \eta)\Delta$:

$$\mathbb{E}[\Phi_{t+1} - \Phi_t \mid V_t] \leq \frac{1}{m} \left(2m \left(\frac{(1.5 + \eta)\Delta}{(2.5 + \eta)\Delta - \Delta} - 1 \right) + \frac{\Phi_t \Delta}{q - \Delta} - \Phi_t \right) = -\frac{\Phi_t}{m} \left(1 - \frac{\Delta}{q - \Delta} \right).$$

Let $\delta := \frac{1 - \frac{\Delta}{q - \Delta}}{m}$. Then we have $\mathbb{E}[\Phi_{t+1} \mid \Phi_t] \leq \Phi_t(1 - \delta)$, i.e., the potential decreases in expectation by a multiplicative factor. Since $\delta T = \left(1 - \frac{\Delta}{q - \Delta}\right) \cdot 5 \log m \geq \log(4m)$ for all $q \geq (2.5 + \eta)\Delta$, applying the Multiplicative Drift Theorem and Markov's inequality yields:

$$\Pr[\Phi_T \geq 1] \leq \mathbb{E}[\Phi_T] \leq m e^{-\delta T} \leq \frac{1}{4}.$$

Thus, with probability at least $3/4$, we have $\Phi_T = 0$, meaning all vertices in S have label size 2. $\square$

The two-stage analysis above completes Phase I: after the initialization pass (lemma 5.8) and the drifting stage (lemma 5.9), all vertices in the seeding set S are 2-bounded with probability at least $3/4$. This configuration that is 2-bounded on seeding vertices serves as the starting point for Phase II.

5.5. Phase II: Converting and Drifting. Phase II begins with the $|\mathcal{L}(v)| = 2$ for all vertex in seeding set S by Phase I and completes the coalescence of the bounding chain. It consists of two stages:

- i *Conversion stage*: a deterministic pass over $V \setminus S$ that ensures $|\mathcal{L}(v)| \in \{1, 2\}$ for every $v \in V \setminus S$;
- ii *Drifting stage*: randomized updates on all vertices that drive every list to size 1 with constant probability.

We begin with the conversion stage.

Lemma 5.10. Assume $q \geq (2.5 + \eta)\Delta$, and let S be the seeding set constructed in Lemma 5.7. After the conversion stage of Phase II (Lines 9–12 in Algorithm 5), every vertex $v \in V \setminus S$ satisfies $|\mathcal{L}(v)| \in \{1, 2\}$.

Proof. We prove by induction on the fixed update order of vertices in $V \setminus S$. Assume that before processing v_i , every $v \in P = S \cup \{v_1, \dots, v_{i-1}\}$ satisfies $|\mathcal{L}(v)| \leq 2$. The procedure $\text{CLEANUP}(v_i, P)$ constructs a reference set $A \subseteq [q]$ of size Δ as follows. Define

$$S_P := \bigcup_{u \in \Gamma(v_i) \cup P} \mathcal{L}(u), \quad \mathcal{D}_P := \bigcup_{\substack{u \in \Gamma(v_i) \cup P: \\ \forall w \in \Gamma(v_i), w \neq u, \mathcal{L}(u) \cup \mathcal{L}(w) = \emptyset}} \mathcal{L}(u),$$

where S_P collects all colors currently active in the neighborhood of v_i and the processed vertices P , while $\mathcal{D}_P$ collects those colors belonging to *disjoint* bounding lists in this neighborhood. GREEDY try to build A including as many colors as possible from $S_P \setminus \mathcal{D}_P$, and then $\mathcal{D}_P$.

Let x denote the number of bounding lists $\mathcal{L}(u)$ for preserved neighbors $u \in P$ that are not contained in $\mathcal{D}_P$, and let y be the number of those that are contained in $\mathcal{D}_P$. Further, let $z = |\Gamma(v_i) \setminus P|$ be the number of unprocessed neighbors of v_i . By definition, $x + y + z = |\Gamma(v_i)| \leq \Delta$. From Lemma 5.7, we also have $z = |\Gamma(v_i)| - x - y \leq (\frac{1}{2} + \eta)\Delta$.

Next, COMPRESS^A is applied to each $w \in \Gamma(v_i) \setminus P$. By Lemma 4.7, after this step every such neighbor satisfies $\mathcal{L}(w) = A \cup \{c_w\}$ for some $c_w \in [q] \setminus A$. Accordingly, the disjoint and slack color sets at v_i can be written as

$$\mathcal{D} := \bigcup_{\substack{u \in \Gamma(v_i): \\ \forall w \in \Gamma(v_i), w \neq u, \mathcal{L}(u) \cup \mathcal{L}(w) = \emptyset}} \mathcal{L}(u) = \mathcal{D}_P \setminus A \setminus \{c_w : w \in \Gamma(v_i) \setminus P\},$$

and

$$\mathcal{S} := \bigcup_{u \in \Gamma(v_i)} \mathcal{L}(u) = (A \cup (S_P \setminus \mathcal{D}_P) \cup \mathcal{D}_P) \cup (\{c_w : w \in \Gamma(v_i) \setminus P\} \setminus A \setminus \mathcal{D}_P).$$

Let $u = |\{c_w : w \in \Gamma(v_i) \setminus P\} \cup (\mathcal{D}_P \setminus A)|$ be the number of distinct colors newly introduced outside A interact to $\mathcal{D}$ during CLEANUP . By definition of x, y, z and the structure of $\mathcal{S}$ and $\mathcal{D}$ above, we have

$$|\mathcal{S}| \leq \max(\Delta, 1.5x + 2y) + z - u, \quad |\mathcal{D}| \geq \max(0, 2y - \max(0, \Delta - 1.5x)) - u,$$

since every bounding list $\mathcal{L}(v)$ not contained in $\mathcal{D}$ contributes at most 1.5 colors to $\mathcal{S}$ by [JSS21].

Combining the two inequalities yields

$$|\mathcal{S}| - \frac{1}{2}|\mathcal{D}| \leq \max(\Delta, 1.5x + 2y) - \max(0, y - \max(0, \frac{1}{2}\Delta - 0.75x)) + z - \frac{1}{2}u.$$

We now upper bound the right-hand side of the above inequality under the constraints

$$x + y + z \leq \Delta, \quad z \leq (\frac{1}{2} + \eta)\Delta, \quad u \geq 0.$$

Since the expression is decreasing in u , it suffices to consider $u = 0$. For fixed z , the feasible region for (x, y) is the triangle $\{x, y \geq 0, x + y \leq \Delta - z\}$, on which the objective

$$f(x, y, z) = \max(\Delta, 1.5x + 2y) - \max(0, y - \max(0, \frac{1}{2}\Delta - 0.75x)) + z$$

is piecewise linear. When $1.5x \geq \Delta$, $f(x, y, z) = 1.5x + y + z \leq 1.5\Delta$. When $1.5 + 2y \leq \Delta$, $f(x, y, z) = \Delta + z \leq (1.5 + \eta)\Delta$. Otherwise $f(x, y, z) = 2.25x + y + z - 0.5\Delta \leq 1.5\Delta$.

Substituting this extremal configuration gives

$$|\mathcal{S}| - \frac{1}{2}|\mathcal{D}| \leq (1.5 + \eta)\Delta.$$

This establishes the desired bound and verifies that the precondition of Lemma 4.8 holds for every vertex v_i during the conversion stage. $\square$

Having completed the deterministic conversion stage, we now analyze the subsequent randomized drifting process. Starting from the globally 2-bounded configuration produced above, we repeatedly apply the DISJOINT coupling on uniformly random vertices until full coalescence.

Lemma 5.11 ([JSS21, Theorem 5.1]). *Assume the process starts from a globally 2-bounded configuration and that $q > (2.5 + \eta)\Delta$, where $\eta = 2\sqrt{(\log \Delta + 1)/\Delta}$. Then, after $T_2 = \frac{2(q-\Delta)n \log n}{q-2.5\Delta}$ iterations of the DISJOINT coupling, the bounding chain collapses to a singleton with probability at least 3/4.*

Remark 5.12. *The proof of Lemma 5.11 follows the same potential-based argument used in the drifting stage of Phase I. The potential function $\Phi_t = |\{v : |\mathcal{L}_t(v)| = 2\}|$ decreases geometrically under the contraction bound provided by Lemma 4.8, leading to exponential convergence in expectation.*

Combining Lemma 5.10 and Lemma 5.11 establishes that, for $q \geq (2.5 + \eta)\Delta$, Phase II succeeds with a probability of 3/4, thereby completing the proof of coalescence.

5.6. Wrapping up. We now combine the results established in the preceding subsections. Assume $q \geq (2.5 + \eta)\Delta$. Lemmas 5.8 and 5.9 show that Phase I (Seeding phase) produces a configuration that is 2-bounded over the seeding set S with probability at least 3/4. Then, Lemmas 5.10 and 5.11 ensure that, starting from such a configuration, Phase II (Converting and Drifting phase) reduces all bounding lists to singletons, also with probability at least 3/4.

By a union bound, a successful coalescence occurs with probability at least 1/2. Together with Lemma 5.6, the total expected computation time is $O(n\Delta^2 \log \Delta)$, which completes the proof of Lemma 5.5.

Combining Lemmas 5.1 and 5.5, we obtain the correctness and efficiency guarantees of the CFTP algorithm, thereby completing the proof of Theorem 1.2.

6. CONCLUSION

We developed a general framework for constructing and analyzing grand couplings based on bounding chains, and applied it to the problem of perfect sampling of proper q -colorings via the coupling-from-the-past (CFTP) method. This framework allows us to formulate the design of bounding chains as an optimization problem, unifying prior constructions and enabling systematic analysis.

Using this approach, we established that $q = 2.5\Delta$ is the tight threshold for bounding-chain-based CFTP algorithms. Our upper bound is achieved through an explicit grand coupling that minimizes the expected size of bounding lists, while the matching lower bound proves that no contractive bounding chain can succeed below this threshold. Together, these results reveal that the bounding chain paradigm is inherently limited by worst-case local configurations that block contraction.

This insight suggests that surpassing the 2.5Δ barrier will likely require moving beyond the bounding-chain framework itself, or at least beyond worst-case contraction analysis.

REFERENCES

- [AJ22] Konrad Anand and Mark Jerrum. Perfect sampling in infinite spin systems via strong spatial mixing. *SIAM Journal on Computing*, 51(4):1280–1295, 2022.
- [BC20] Siddharth Bhandari and Sayantan Chakraborty. Improved bounds for perfect sampling of k -colorings in graphs. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 631–642, 2020.
- [BD97] Russ Bubley and Martin Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In *FOCS*, pages 223–231, 1997.
- [CDM⁺19] Sitan Chen, Michelle Delcourt, Ankur Moitra, Guillem Perarnau, and Luke Postle. Improved bounds for randomly sampling colorings via linear programming. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2216–2234, 2019.
- [CFYZ21] Xiaoyu Chen, Weiming Feng, Yitong Yin, and Xinyuan Zhang. Rapid mixing of Glauber dynamics via spectral independence for all degrees. In *FOCS*, pages 137–148, 2021.
- [CLMM23] Zongchen Chen, Kuikui Liu, Nitya Mani, and Ankur Moitra. Strong spatial mixing for colorings on trees and its algorithmic applications. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 810–845, 2023.
- [CLV21] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of Glauber dynamics: entropy factorization via high-dimensional expansion. In *STOC*, pages 1537–1550. ACM, 2021.
- [CPP02] Henry Cohn, Robin Pemantle, and James Propp. Generating a random sink-free orientation in quadratic time. *The Electronic Journal of Combinatorics [electronic only]*, 9(1):Research paper R10, 13 p.–Research paper R10, 13 p., 2002.
- [CV25] Charlie Carlson and Eric Vigoda. Flip dynamics for sampling colorings: Improving $(11/6 - \varepsilon)$ using a simple metric. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2194–2212. SIAM, 2025.
- [DFHV13] Martin Dyer, Alan M. Frieze, Thomas P. Hayes, and Eric Vigoda. Randomly coloring constant degree graphs. *Random Struct. Algorithms*, 43(2):181–200, 2013.
- [FGW⁺23] Weiming Feng, Heng Guo, Chunyang Wang, Jiaheng Wang, and Yitong Yin. Towards derandomising Markov chain Monte Carlo. In *FOCS*, pages 1963–1990. IEEE, 2023.
- [FGY22] Weiming Feng, Heng Guo, and Yitong Yin. Perfect sampling from spatial mixing. *Random Structures & Algorithms*, 61(4):678–709, 2022.
- [FGYZ22] Weiming Feng, Heng Guo, Yitong Yin, and Chihao Zhang. Rapid mixing from spectral independence beyond the boolean domain. *ACM Transactions on Algorithms*, 18(3):1–32, 2022.
- [FH00] James Allen Fill and Mark Huber. The randomness recycler: A new technique for perfect sampling. In *Proceedings 41st annual symposium on foundations of computer science*, pages 503–511. IEEE, 2000.
- [FVY21] Weiming Feng, Nisheeth K Vishnoi, and Yitong Yin. Dynamic sampling from graphical models. *SIAM Journal on Computing*, 50(2):350–381, 2021.
- [GJL19] Heng Guo, Mark Jerrum, and Jingcheng Liu. Uniform sampling through the lovász local lemma. *Journal of the ACM (JACM)*, 66(3):1–31, 2019.
- [GK12] David Gamarnik and Dmitriy Katz. Correlation decay and deterministic fptas for counting colorings of a graph. *Journal of Discrete Algorithms*, 12:29–47, 2012.
- [GLMP24] Andreas Göbel, Jingcheng Liu, Pasin Manurangsi, and Marcus Pappik. Perfect sampling from rapidly mixing markov chains, 2024. Accessed: 2025-06-28.
- [GŠV15] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability for antiferromagnetic spin systems in the tree nonuniqueness region. *J. ACM*, 62(6):Art. 50, 60, 2015.
- [Hay06] Thomas P. Hayes. A simple condition implying rapid mixing of single-site dynamics on spin systems. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 39–46, 2006.
- [HLP52] Godfrey Harold Hardy, John Edensor Littlewood, and George Pólya. *Inequalities*. Cambridge university press, 1952.

- [HN98] O. Häggström and K. Nelander. Exact sampling from anti-monotone systems. *Statistica Neerlandica*, 52(3):360–380, 1998.
- [HSW21] Kun He, Xiaoming Sun, and Kewen Wu. Perfect sampling for (atomic) lovász local lemma. arXiv:2107.03932, 2021.
- [Hub98] Mark Huber. Exact sampling and approximate counting techniques. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 31–40, 1998.
- [Hub04] Mark L. Huber. Perfect sampling using bounding chains. *The Annals of Applied Probability*, 14(2):735–753, 2004.
- [Hub16] M.L. Huber. *Perfect Simulation*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC Press, 2016.
- [HV03] Thomas P. Hayes and Eric Vigoda. A non-markovian coupling for randomly sampling colorings. In *FOCS*, pages 618–627, 2003.
- [HV06] Thomas P. Hayes and Eric Vigoda. Coupling with the stationary distribution and improved sampling for colorings and independent sets. *Ann. Appl. Probab.*, 16(3):1297–1318, 2006.
- [HWY22] Kun He, Chunyang Wang, and Yitong Yin. Sampling lovász local lemma for general constraint satisfaction solutions in near-linear time. In *FOCS*, pages 147–158. IEEE, 2022.
- [Jer95] Mark Jerrum. A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Structures & Algorithms*, 7(2):157–165, 1995.
- [Jon02] Johan Jonasson. Uniqueness of uniform random colorings of regular trees. *Statistics & Probability Letters*, 57(3):243–248, 2002.
- [JPV22] Vishesh Jain, Huy Tuan Pham, and Thuy-Duong Vuong. Spectral independence, coupling, and the spectral gap of the glaufer dynamics. *Information Processing Letters*, 177:106268, 2022.
- [JSS21] Vishesh Jain, Ashwin Sah, and Mehtaab Sawhney. Perfectly sampling $k > (8/3 + o(1))\delta$ -colorings in graphs. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1589–1600, 2021.
- [JVV86] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.*, 43(2-3):169–188, 1986.
- [LSS19] Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava. A deterministic algorithm for counting colorings with $2\text{-}\delta$ colors. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1380–1404. IEEE, 2019.
- [LWY26] Hongyang Liu, Chunyang Wang, and Yitong Yin. Local gibbs sampling beyond local uniformity. To appear in *SODA ’26*, 2026.
- [LWYY25] Jingcheng Liu, Chunyang Wang, Yitong Yin, and Yixiao Yu. Phase transitions via complex extensions of markov chains. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC ’25, page 903–914, 2025.
- [LY13] Pinyan Lu and Yitong Yin. Improved fptas for multi-spin systems. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 639–654. Springer, 2013.
- [MT10] Robin A Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *Journal of the ACM (JACM)*, 57(2):11, 2010.
- [PW96] James Gary Propp and David Bruce Wilson. Exact sampling with coupled markov chains and applications to statistical mechanics. *Random Structures & Algorithms*, 9(1-2):223–252, 1996.
- [QWZ22] Guoliang Qiu, Yanheng Wang, and Chihao Zhang. A perfect sampler for hypergraph independent sets, 2022.
- [SS97] Jesús Salas and Alan D Sokal. Absence of phase transition for antiferromagnetic potts models via the dobrushin uniqueness theorem. *Journal of Statistical Physics*, 86:551–579, 1997.
- [Vig99] Eric Vigoda. Improved bounds for sampling colorings. In *40th Annual Symposium on Foundations of Computer Science*, pages 51–59. IEEE, 1999.
- [Wil96] David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC ’96, page 296–303, 1996.