The Randomized Coloring Procedure with Symmetry-Breaking

Sriram Pemmaraju[1] and Aravind Srinivasan[2*]

[1] Dept. of Computer Science,
The University of Iowa, Iowa City, IA 52242-1419, USA
sriram@cs.uiowa.edu

[2] Dept. of Computer Science and Institute for Advanced Computer Studies,
University of Maryland, College Park, MD 20742, USA
srin@cs.umd.edu

Abstract. A basic *randomized coloring procedure* has been used in probabilistic proofs to obtain remarkably strong results on graph coloring. These results include the asymptotic version of the List Coloring Conjecture due to Kahn, the extensions of Brooks' Theorem to sparse graphs due to Kim and Johansson, and Luby's fast parallel and distributed algorithms for graph coloring. The most challenging aspect of a typical probabilistic proof is showing adequate concentration bounds for key random variables. In this paper, we present a simple symmetry-breaking augmentation to the randomized coloring procedure that works well in conjunction with *Azuma's Martingale Inequality* to easily yield the requisite concentration bounds. We use this approach to obtain a number of results in two areas: *frugal coloring* and *weighted equitable coloring.* A *β-frugal coloring* of a graph G is a proper vertex-coloring of G in which no color appears more than β times in any neighborhood. Let $G = (V, E)$ be a vertex-weighted graph with weight function $w : V \rightarrow [0, 1]$ and let $W = \sum_{v \in V} w(v)$. A *weighted equitable coloring* of G is a proper k-coloring such that the total weight of every color class is "large", i.e., "not much smaller" than W/k; this notion is useful in obtaining tail bounds for sums of dependent random variables.

1 Introduction and Summary of Results

The *randomized coloring procedure* refers to a simple randomized algorithm for coloring graphs that has been used in probabilistic proofs over the past two decades, to obtain remarkably strong results on graph coloring. Some of these results are existential, whereas some lead to polynomial-time algorithms. These results include the asymptotic version of the List Coloring Conjecture due to Kahn [16], and Kim [18] and Johansson's [14] extensions of Brooks' Theorem for graphs with lower-bounded girth. The randomized coloring procedure has also been used by Luby [22] to obtain fast parallel and distributed algorithms for graph coloring. In its simplest form, the procedure is:

> Each vertex v picks a *tentative color* uniformly at random from its color palette. With high probability the tentative coloring may not be proper; it is repaired by uncoloring each vertex v that receives the same color as a neighbor.

The randomized coloring procedure allows us to claim the existence of a proper, partial coloring of the graph, having certain desired properties. This coloring can be extended

* Supported in part by NSF Award CCR-0208005, NSF ITR Award CNS-0426683, and NSF Award CNS-0626964. Part of this work was done while on sabbatical at the Network Dynamics and Simulation Science Laboratory of the Virginia Bioinformatics Institute, Virginia Tech.

to a complete proper coloring in a variety of ways (for example, greedily). Variants of this procedure involve vertices making a *non-uniform* color choice or each vertex having an *activation probability*. For example, in Johansson's extension of Brooks' Theorem for triangle-free graphs [14], it is important that some colors are used more than others and to ensure this, a non-uniform probability distribution is used for the choice of colors. In Luby's parallel algorithm for graph coloring [22] each node has an activation probability of 1/2; the algorithm starts by flipping an unbiased coin for each vertex and only those vertices that are activated in this step take further part in the coloring procedure.

In the *iterative* or *incremental* version of this procedure, the partial coloring obtained after one application of the procedure is not completed deterministically; instead the coloring is incrementally extended by repeating the procedure. Specifically, after one application of the procedure we argue that (i) sufficient progress has been made and (ii) the partial coloring has the desired properties with positive (though, typically very small) probability. A "good" partial coloring is then fixed and we extend this coloring by performing the next iteration of the randomized coloring procedure. This incremental version of the procedure is extremely powerful and the results [14, 16, 18, 22] mentioned above, are obtained thus. The incremental randomized coloring procedure is a special case of the general technique variously referred to as the "semi-random method," the "pseudo-random method," or the "Rödl Nibble."

In this paper, we present a simple symmetry-breaking augmentation to the randomized coloring procedure. This symmetry-breaking approach, when used in conjunction with *Azuma's Martingale Inequality*, allows us to use the incremental randomized coloring procedure to prove a number of new results (mentioned below). In its simplest form, our approach starts by picking a permutation π of the vertex set V of the given graph G. Depending on the application, the permutation π can be arbitrary or random or dictated by the structure of G. As in the standard randomized coloring procedure, in the first step, vertices pick tentative colors. Then, in the uncoloring step, a vertex v is uncolored *only if there is a neighbor of lower rank in* π *that has received the same tentative color*. Thus the difference between the symmetry-breaking approach and the standard randomized procedure is only in which neighbors are examined in the uncoloring step.

The most challenging aspect of a typical probabilistic-method proof is showing the adequate concentration of key random variables. For example, for the first result of this paper, we need to show that after an iteration of the randomized coloring, a constant fraction of the neighbors of each vertex get colored. Let the random variable P_v denote the number of neighbors of vertex v that get permanently colored in one iteration of the coloring procedure. It can be shown that $E[P_v] = \alpha \cdot degree(v)$ for some constant $0 < \alpha < 1$; we need to show that P_v is sharply concentrated around $E[P_v]$. In establishing the concentration of P_v, we need to take into account the fact that even though vertices independently chose tentative colors, the permanent acquisition of colors by vertices may be highly correlated. Recall that a vertex u permanently acquires a color x if u has tentatively chosen x and no neighbor of u has. To get around such "dependence" problems, concentration inequalities such as Azuma's Inequality that do not require independence, have been used widely. However, obtaining sharp enough concentration bounds using Azuma's Inequality is not always easy. Consider the following version of the inequality [23]:

Lemma 1. ***[Azuma's Inequality]*** *Let X be a random variable determined by n trials $T_1, T_2, \ldots, T_n$, such that for each i, and any two possible sequences of outcomes $t_1, t_2, \ldots, t_{i-1}, t_i$ and $t_1, t_2, \ldots, t_{i-1}, t'_i$:*

$$\Big| E[X|T_1 = t_1, \ldots, T_i = t_i] - E[X|T_1 = t_1, \ldots, T_i = t'_i] \Big| \leq c_i \tag{1}$$

then

$$\Pr\Big[|X - E[X]| > t\Big] \leq 2e^{-t^2/(2\sum c_i^2)}. \tag{2}$$

The difficulty in using Azuma's Inequality arises from the need to show that $\sum_i c_i^2$ is small. For example, if each c_i is shown to be bounded above by a small constant, then the resulting bound in (2) is $e^{-\varepsilon t^2/n}$ for a positive constant ε. The "n" here is problematic, as it may be considerably larger than $E[X]$. The more desirable concentration bound is $e^{-\Omega(\varepsilon t^2/E[X])}$. For our random variable P_v, n may be as large as Δ^2 (Δ being the maximum degree of the graph), whereas $E[P_v]$ is linear in Δ. This is because the permanent acquisition of colors by vertices in the neighborhood of v depends on the tentative choices of colors by vertices in the distance-2 neighborhood of v. The symmetry-breaking approach provides critical help in showing that $\sum_i c_i^2$ is small, usually $O(E[X])$. For example, to show that P_v is concentrated around $E[P_v]$, we consider $U_v = \langle u_1, u_2, \ldots, u_\ell \rangle$, the sequence of vertices at distance at most two from v, arranged in *increasing rank according to π*. Letting T_i denote the tentative choice of a color by u_i, we observe that changing T_i *does not affect any of the lower ranked vertices; only the higher ranked vertices.* However, (1) only considers the *expected* effect of changing T_i on higher ranked vertices, rather than the worst-case effect. These observations play a critical role in proving a better bound on $\sum_i c_i^2$; see, e.g., the proofs of Lemmas 4 and 6.

Results and Notation. We apply randomized coloring with symmetry-breaking to obtain a number of results in two areas: *frugal coloring* and *weighted equitable coloring.* Our results include existential bounds, polynomial-time algorithms, and polylog-time distributed algorithms. We state our specific results next. All logarithms here are natural logarithms, unless specified otherwise. We use e to denote the base of the natural logarithm, and for any positive integer h, $[h]$ to denote the set $\{1, 2, \ldots, h\}$. For a vertex v, $N(v)$ denotes the set of neighbors of v in some graph G that is clear from the context; $deg(v)$ denotes $|N(v)|$. We make use of the following version of the Lovász Local Lemma.

Lemma 2. ***[Lovász Local Lemma]*** *Consider a set $\mathcal{E}$ of events such that for each $A \in \mathcal{E}$, $\Pr[A] \leq p < 1$, and A is mutually independent of a set of all but at most d other events from $\mathcal{E}$. If $4pd \leq 1$, then with positive probability, none of the events in $\mathcal{E}$ occur.*

Frugal coloring. A *β-frugal coloring* of a graph G is a proper vertex-coloring of G in which no color appears more than β times in any neighborhood. Frugal coloring is a useful subroutine in *total coloring* [10] and can also be used as a subroutine for efficient channel-allocation schemes in multi-channel, multi-radio wireless networks [11]. We obtain four results (F1)-(F4) for frugal coloring; let Δ denote the maximum degree and n denote the number of vertices of the graph.

(F1) Every graph has a $(\Delta+1)$-coloring that is $O(\frac{\log^2 \Delta}{\log\log \Delta})$-frugal. Such a coloring can be computed in polynomial time. Furthermore, there is a distributed algorithm that can compute an $O(\log \Delta \cdot \frac{\log n}{\log\log n})$-frugal, $(\Delta+1)$-coloring in $O(\log n)$ communication rounds.

Hind, Molloy, and Reed [10] show that any graph has a $O(\log^5 \Delta)$-frugal, $(\Delta+1)$ coloring, and also show a lower bound of $\Omega(\log \Delta / \log\log \Delta)$ on the frugality of any $(\Delta+1)$-coloring. Thus our result improves on the Hind-Molloy-Reed upper bound and is "$\log \Delta$" away from being optimal[3]. Result (F1) is obtained via $O(\log \Delta)$ iterations of the randomized coloring procedure with symmetry-breaking and a proof that shows that in each iteration, each color is used $O(\log \Delta / \log\log \Delta)$ times in each neighborhood. In order to tightly control the coloring procedure, we also assign activation probabilities to vertices before each iteration; those vertices which are not activated, sleep through the iteration.

(F2) Let g be the maximum number of nodes at distance within $C_0 \log \Delta$ from any vertex in the given graph G, where C_0 is an absolute constant. Then G has a $(\Delta+1)$-coloring that is t-frugal, where $t = O\left(\frac{\log g}{\log(\log g / \log \Delta)}\right)$.

This result subsumes the bound of (F1), since $g \leq \Delta^{O(\log \Delta)}$. Note that if $g \leq \Delta^{O(1)}$, then t is just $O(\log \Delta)$. This result is particularly applicable to "growth-bounded" graphs [8, 21] which are graphs for which the number of nodes within distance r from any node grows much smaller than Δ^r (typically, as a polynomial in r). Thus, even for graphs with *exponential* growth (the number of nodes at a distance r can be $\mathrm{poly}(\Delta, 2^r)$), we still obtain an $O(\log \Delta)$ bound on the frugality. For this result, we again use iterative randomized coloring, but carefully exploit the fact that each iteration is *local* in the graph, and that there are only $O(\log \Delta)$ iterations: this is the reason for our definition of g going only up to distance $O(\log \Delta)$.

(F3) Let G be a *d-inductive* graph. Then G has a $(d+1)$-coloring that is t-frugal, where $t = O(\frac{\Delta}{d} \cdot \log^{1+\epsilon} \Delta)$ for any constant $\epsilon > 0$ whenever $d < 2\Delta / \log \Delta$, and $t = O(\frac{\Delta}{d} \cdot \frac{\log^2 \Delta}{\log(d \log \Delta / \Delta)})$ when $d \geq 2\Delta / \log \Delta$.

Recall that a graph is said to be *d-inductive* if there is an ordering $v_1, v_2, \ldots, v_n$ of its vertex set such that every vertex v_i has at most d neighbors in $\{v_1, v_2, \ldots, v_{i-1}\}$. For instance, planar graphs are 5-inductive. Any such ordering $v_1, v_2, \ldots, v_n$ is called a *d-inductive ordering*. A greedy algorithm that considers vertices of a d-inductive graph G in d-inductive order and assigns to each vertex the smallest available color, succeeds in producing a proper $(d+1)$-vertex coloring of G. Note that no $(d+1)$-coloring can have frugality better than $\Delta/(d+1)$ and therefore our result provides upper bounds that are within $O(\mathrm{polylog}(\Delta))$ of the $\Delta/(d+1)$ lower bound. Also note that since any graph with maximum vertex degree Δ is Δ-inductive, our result (F1) on arbitrary graphs can be viewed as a special case of this result.

(F4) Let G be an n-vertex graph with maximum degree Δ and girth at least 5. Then G has an $O(\Delta / \log \Delta)$-coloring that is $O(\frac{\log^2 \Delta}{\log\log \Delta})$-frugal. Such a coloring can be computed in polynomial time. Furthermore, an $O(\Delta / \log \Delta)$-coloring that is $O(\log \Delta \cdot \frac{\log n}{\log\log n})$-frugal can be constructed by a distributed algorithm in $O(\log n)$ rounds of communication.

---

[3] In recent personal communication, Molloy and Reed have mentioned that they have proved the existence of an $O(\log \Delta / \log\log \Delta)$-frugal $(\Delta+1)$-coloring; their paper is currently under preparation.

This result generalizes the result of Kim [18] and the result of Grable and Panconesi [7]. In 1948 Brooks showed that any connected graph with maximum degree Δ, with the exception of odd cycles and the $(\Delta + 1)$-clique, has a Δ-coloring. In 1968 Vizing asked if this bound could be improved for "sparse graphs" of certain kinds — graphs with large girth, for example. This question remained unanswered until the work of [14] showed that any triangle-free graph has chromatic number $O(\Delta/\log \Delta)$. Building on this and a result due to [18] for graphs with girth at least 5, the paper [7] presents a simple randomized distributed algorithm that produces an $O(\Delta/\log \Delta)$-coloring of a triangle-free graph with high probability, in $O(\log n)$ communication rounds. This algorithm requires that $\Delta \geq \log^{1+\delta} n$ for any constant $\delta > 0$, whereas the existential result of [14] holds for any Δ. We extend the analysis of [7, 18] to obtain (F4). We do not require our symmetry-breaking approach for the polynomial-time algorithm guaranteed by (F4).

Weighted Equitable Coloring. An *equitable coloring* of a graph is a proper vertex coloring such that the sizes of any two color classes differ by at most 1. If a k-coloring of an n-vertex graph G is equitable, then the size of every color class is in $\{\lceil n/k \rceil, \lfloor n/k \rfloor\}$. In a pivotal result, Hajnal and Szemerédi [9] showed that every graph G with maximum degree at most Δ has an equitable k-coloring for every $k \geq \Delta + 1$. Recently, Kierstead and Kostochka [17] have presented a short proof of this result, along with a polynomial-time algorithm for computing such a coloring. Equitable colorings naturally arise in scheduling, partitioning, and load balancing [1, 3, 12, 20, 25, 26]. Pemmaraju [24] and Janson and Ruciński [13] have used equitable colorings to derive large-deviation bounds for sums of random variables that exhibit limited dependence. Let $\mathcal{X} = \{X_1, X_2, \ldots, X_n\}$ be a collection of bounded random variables with $X_i \in [0, 1]$ for all i, let $S = \sum_i X_i$, and let $\mu = E[S]$. It is well-known that if the X_i's are independent then the following bounds on the tail probabilities of S, due to Chernoff [5], hold:

$$\Pr[S \geq (1+\varepsilon)\mu] \leq F^+(\mu, \varepsilon) \doteq \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{(1+\varepsilon)}}\right)^{\mu}; \quad \Pr[S \leq (1-\varepsilon)\mu] \leq F^-(\mu, \varepsilon) \doteq e^{-\mu\varepsilon^2/2}. \tag{3}$$

Researchers have attempted to extend the Chernoff bounds to situations where the X_i's exhibit limited dependence. We will model the "limited dependence" as usual by a *dependency graph* G of $\mathcal{X}$: G is an undirected graph with vertices $\{1, 2, \ldots, n\}$ such that if $\{i_1, i_2, \ldots, i_\ell\}$ *is any independent set in G, then the random variables* $X_{i_1}, X_{i_2}, \ldots, X_{i_\ell}$ *are mutually independent.* If the X_i's have the *same* mean and G can be equitably k-colored, then the following tail-bounds hold [24]: $\Pr[S \geq (1+\varepsilon)\mu] \leq \frac{4k}{e} F^+(\mu, \varepsilon)^{1/k}$, and $\Pr[S \leq (1-\varepsilon)\mu] \leq \frac{4k}{e} F^-(\mu, \varepsilon)^{1/k}$, where $F^+(\mu, \varepsilon)$ and $F^-(\mu, \varepsilon)$ are from (3). Observe that if k, the palette size of the equitable coloring, is small then these bounds are quite good relative to the Chernoff bounds.

The above bounds apply only in the case where the X_i's have the same mean. To deal with the general case of arbitrary $X_i \in [0, 1]$ using a similar approach, one needs to consider *weighted equitable colorings*. Let $G = (V = [n], E)$ be a vertex-weighted graph with weight function $w : V \to [0, 1]$ such that $w(i) = E[X_i]$, and let $W = \mu = E[S] = \sum_{v \in V} w(v)$. Informally, a weighted equitable coloring of G is a proper k-coloring such that the total weight of **every** color class is "large", i.e., "not much smaller" than W/k. (A classical equitable coloring is the special case with unit weight for all vertices.) Call a k-coloring of G with λ being the minimum weight of all color classes, a (k, λ)-*coloring*. Using the approach of [13, 24], such a coloring can be used to show that

$$\Pr[S \geq (1+\varepsilon)\mu] \leq k \cdot F^{+}(\lambda, \varepsilon); \quad \Pr[S \leq (1-\varepsilon)\mu] \leq k \cdot F^{-}(\lambda, \varepsilon). \tag{4}$$

Hence, given k, we aim for as large a λ as possible. We prove two results (E1), (E2) for $k = \Delta + 1$ colors, where Δ denotes the maximum degree of G. We prove result (E3) for d-inductive graphs with $k = d + 1$.

(E1) A $(\Delta+1, \lambda)$-coloring exists, with $\lambda \geq \left(1 - \frac{1}{e}\right)\frac{W}{\Delta+1} - O(\sqrt{W \log(\Delta+1)})$.

(E2) There is a constant $c > 0$ such that a $(\Delta + 1, \lambda)$-coloring exists, with $\lambda \geq \left\lfloor \frac{cW}{\Delta \log(\Delta+1)} \right\rfloor$.

When $W > \Delta^{2+\Omega(1)}$, the lower bound in (E1) simplifies to $(1 - 1/e - o(1))\frac{W}{\Delta+1}$, which is to within a constant of the best possible. Result (E2) holds for all values of W relative to Δ and in particular holds for "small" W as well. As shown by (4), these results yield bounds on the tail probabilities of sums of *arbitrary* bounded random variables. (E1) is obtained via a single iteration of randomized coloring with symmetry-breaking; (E2) is obtained by combining [9] and (E1) with a partitioning approach. We note that for a slightly smaller choice of c, we can also obtain (E2) by combining [9] and a second-moment analysis with partitioning; however, we are not aware of any other approach that yields our result (E1).

(E3) Every d-inductive graph has a $(d{+}1, \lambda)$-coloring with $\lambda \geq \frac{1}{e} \cdot \frac{W}{d+1} - O\left(\sqrt{\frac{W}{d+1}\Delta \log(d+1)}\right)$.

When $W > \Delta \cdot d^{1+\Omega(1)}$, the lower bound in (E3) simplifies to $(\frac{1}{e} - o(1))\frac{W}{d+1}$. Note that since a graph with maximum degree Δ is a Δ-inductive graph, result (E3) can be viewed as a generalization of result (E1), with the small change that the leading constant is $1/e$ instead of $1 - 1/e$. This result extends known results on classical equitable colorings of d-inductive graphs [4, 19]. Bollobas and Guy [4] consider the equitable coloring of 1-inductive graphs (i.e., forests) whereas Kostochka et al. [19] consider d-inductive graphs for arbitrary d. Specifically, Kostochka et al. [19] show that every d-inductive graph has an equitable coloring with at most $16d$ colors (provided $\Delta < n/15$). Here, in result (E3), we relax the notion of "equitability" while requiring that exactly $d{+}1$ colors be used. Like (E1), result (E3) is also obtained via one iteration of the randomized coloring procedure, but with π being an arbitrary d-inductive vertex ordering.

Organization. Due to space constraints, we only present proofs of results (F1) and (E1) in this paper; these appear in the next two sections. For result (F1) we only present the existential proof, postponing the algorithmic results to the full version of the paper. We highlight the use of symmetry-breaking and Azuma's Inequality in our proofs.

2 Frugal coloring for arbitrary graphs

In this section we prove result (F1). This follows by repeated application of the following result, that describes what happens in one iteration of randomized coloring procedure.

Theorem 1. *Let $G = (V, E)$ be a graph with maximum vertex degree Δ. Suppose that associated with each vertex $v \in V$, there is a palette $P(v)$ of colors, where $|P(v)| \geq$*

$deg(v)+1$. Furthermore, suppose $|P(v)| \geq \Delta/4$ for all vertices v in G. Then, for some subset $C \subseteq V$, there is a list coloring of the vertices in C such that:

(a) $G[C]$ is properly colored.
(b) For every vertex $v \in V$ and for every color x, there are at most $9 \cdot \frac{\log \Delta}{\log\log \Delta}$ neighbors of v colored x.
(c) For every vertex $v \in V$, the number of neighbors of v not in C is at most $\Delta(1-\frac{1}{e^5})+27\sqrt{\Delta \log \Delta}$.
(d) For every vertex $v \in V$, the number of neighbors of v in C is at most $\frac{\Delta}{e^5}+27\sqrt{\Delta \log \Delta}$.

Before we prove this theorem, we show how repeated applications of it yield result (F1), proving the existence of an $O\left(\frac{\log^2 \Delta}{\log\log \Delta}\right)$-frugal, $(\Delta+1)$-coloring of a graph G with maximum degree Δ. Start by associating the palette of colors $[\Delta+1]$ to each vertex. Letting $P_0(v)$ denote the initial palette of a vertex v, we have $P_0(v)=[\Delta+1]$ for all $v \in V$. Let $G_0 = G$. For each $i \geq 0$, we apply Theorem 1 to obtain a partial coloring of G_i. Let G_{i+1} denote the subgraph of G_i induced by vertices that are not colored in this partial coloring of G_i. The palette of colors $P_{i+1}(v)$ associated with a vertex v in G_{i+1} is obtained by deleting from $P_i(v)$ all colors used by neighbors of v in the partial coloring of G_i. Let $deg_i(v)$ denote the degree of vertex v in G_i. Let Δ_i denote the maximum vertex degree of G_i and let p_i denote the minimum palette size in G_i. Thus, $\Delta_0 = \Delta$ and $p_0 = \Delta+1$. Note that initially the requirements of Theorem 1 are satisfied. Suppose that for some $i \geq 0$, the requirements of the theorem are satisfied. That is, (i) $|P_i(v)| \geq deg_i(v)+1$ for all vertices v and (ii) $p_i \geq \Delta_i/4$. Since the palette of a vertex loses at most as many colors as neighbors that are colored, it is still true that $|P_{i+1}(v)| \geq deg_{i+1}(v)+1$ for all v in G_{i+1}. Theorem 1(c) implies that $\Delta_{i+1} \leq \Delta_i(1-\frac{1}{e^5})+27\sqrt{\Delta_i \log \Delta_i}$ and Theorem 1(d) implies that $p_{i+1} \geq p_i - \frac{\Delta_i}{e^5} - 27\sqrt{\Delta_i \log \Delta_i}$. Thus the worst case behavior of Δ_i and p_i is captured by the recurrences:

$$\Delta_{i+1} = \Delta_i\left(1-\frac{1}{e^5}\right)+27\sqrt{\Delta_i \log \Delta_i}; \qquad p_{i+1} = p_i - \frac{\Delta_i}{e^5} - 27\sqrt{\Delta_i \log \Delta_i}. \tag{5}$$

The above recurrences can be solved to obtain the following bounds on Δ_i and p_i.

Lemma 3. *Let $\alpha = (1-1/e^5)$. For all i for which $\Delta_i \geq 10^9$, $\Delta_i \leq 2\Delta_0 \cdot \alpha^i$ and $p_i \geq \frac{\Delta_0}{2}\alpha^i$.*

This implies that, provided p_{i+1} and Δ_{i+1} are large enough, it is the case that $p_{i+1} \geq \Delta_{i+1}/4$, thereby permitting the next application of the above theorem. To get a $(\Delta+1)$-coloring of the desired frugality, we repeatedly obtain partial colorings by applying the Theorem 1 until $\Delta_i < 10^9$. Given the rate of decay of Δ_i, letting $\alpha = (1-1/e^5)$, we see that at most $\log_{1/\alpha}\left(\frac{2\Delta_0}{10^9}\right) = O(\log \Delta)$ applications of the theorem are needed. Since the palette at every vertex has at least one more color than the number of neighbors of the vertex, a greedy list coloring algorithm will succeed in completing the coloring of the graph. Clearly, what we have constructed is a $(\Delta+1)$-proper vertex coloring of G. In each round i, each color appears in a neighborhood at most $9\log \Delta_i/\log\log \Delta_i = O(\log \Delta/\log\log \Delta)$ times. The final round adds only $O(1)$ copies of any color to any neighborhood, yielding a $(\Delta+1)$-coloring that is $O(\log^2 \Delta/\log\log \Delta)$-frugal. We have thus proved result (F1).

Proof of Theorem 1. We start by describing a randomized coloring procedure that will produce, with positive probability, a partial coloring of G with the four desired properties. Let π be an arbitrary permutation of V. This establishes a ranking of the vertices. For any vertex v and color $x \in P(v)$, let $L(v, x)$ be the set of neighbors u of v such that u has a lower rank than v in π and u contains x in its palette $P(u)$. Each vertex v computes the quantity

$$q_v = \frac{1}{|P(v)|} \sum_{x \in P(v)} \prod_{u \in L(v,x)} \left(1 - \frac{1}{|P(u)|}\right).$$

This is the probability that no lower ranked neighbor of v will tentatively pick the color picked by v. Vertex v will use the value of q_v to determine its "sleep probability." For each vertex $v \in V$, independently pick a color $x \in P_v$ uniformly at random. We say that x is the *tentative color* of v. After picking a tentative color, v either goes off to sleep for the rest of this round or stays awake and attempts to make its tentative color permanent. Specifically, vertex v stays awake with probability $a_v = \frac{1}{q_v \cdot e^5}$ and it dozes off with probability $(1 - a_v)$. Later we will show that q_v is never smaller than $1/e^5$ and therefore $a_v \leq 1$. If v dozes off, then it remains uncolored at the end of the procedure. Note that v dozes off only after picking a tentative color and even though it may fall asleep, this choice of tentative color by v may have an influence on whether a neighbor gets permanently colored or not. If v stays awake, then it checks if there is a neighbor $u \in N(v)$ of smaller rank in π, that has the same tentative color as v. If no such u exists, then v is *permanently* colored x. The vertex subset C consists of all vertices that are permanently colored at the end of the procedure. The rest of the vertices are said to be *uncolored.*

For any vertex $v \in V$ and color $x \in C$, let $T_{v,x}$ be the indicator random variable that equals 1 if x is picked as v's tentative color. Note that $\Pr[T_{v,x} = 1] = 1/|P(v)|$ and therefore $E[T_{v,x}] = 1/|P(v)|$. For every $v \in V$ and $x \in C$, let $N_{v,x}$ be the random variable that equals the number of neighbors of v permanently colored x. Since a vertex has to be tentatively colored x before it can be permanently colored x, $N_{v,x} \leq \sum_{u \in N(v)} T_{u,x}$. By linearity of expectation, $E[\sum_{u \in N(v)} T_{u,x}] = \sum_{u \in N(v)} 1/|P(u)| \leq \frac{\Delta}{\Delta/4} = 4$. Note that the lower bound on the palette sizes plays a critical role here. Since the $T_{u,x}$ are mutually independent, we can use the Chernoff bound to show that $\sum_{u \in N(v)} T_{u,x}$ exceeds $9 \cdot \frac{\log \Delta}{\log \log \Delta}$ with probability less than $\frac{1}{\Delta^6}$. Therefore, for any vertex v and any color x,

$$\Pr\left[\text{Number of neighbors of } v \text{ colored } x \text{ exceeds } 9 \cdot \frac{\log \Delta}{\log \log \Delta}\right] < \frac{1}{\Delta^6}. \qquad (6)$$

For any vertex $v \in V$, let R_v be the indicator random variable that equals 1 if v is permanently colored at the end of the coloring procedure. Note that v is permanently colored if v stayed awake and if no lower-ranked neighbor picked the same tentative color as vertex v did. Therefore, the probability that R_v equals 1, is $\Pr[R_v = 1] = a_v \cdot \frac{1}{|P(v)|} \cdot \sum_{x \in P(v)} \prod_{u \in L(v,x)} \left(1 - \frac{1}{|P(u)|}\right) = a_v \cdot q_v$. Recall that a_v was chosen to be $\frac{1}{q_v \cdot e}$ and therefore $\Pr[R_v = 1] = 1/e^5$. However, for the "go to sleep" step of the coloring procedure to be well-defined, we need to show that $a_v \leq 1$. We do this by showing a lower bound of $1/e^5$ on q_v. Recall that q_v is the probability that no neighbor will choose the same tentative color as v. This probability is minimized when vertex v's palette is as small as possible, v has as many neighbors as possible, each of these neighbors have palettes that are as small as possible, and finally each of these palettes is identical to v's palette. Therefore,

$$q_v = \frac{1}{|P(v)|} \cdot \sum_{x \in P(v)} \prod_{u \in L(v,x)} \left(1 - \frac{1}{|P(u)|}\right) \geq \frac{1}{|P(v)|} \cdot \sum_{x \in P(v)} \left(1 - \frac{1}{\Delta/4}\right)^{\Delta} = \left(1 - \frac{4}{\Delta}\right)^{\Delta} \geq \frac{1}{e^5}.$$

Since $q_v \geq 1/e^5$, it follows that $a_v = 1/(q_v e^5) \leq 1$.

Let P_v denote the number of neighbors of v that are permanently colored by the procedure. Note that $P_v = \sum_{u \in N(v)} R_u$. Then by linearity of expectation, $E[P_v] = \sum_{u \in N(v)} E[R_u] = deg(v)/e^5$. Since the random variables R_u are not mutually independent for $u \in N(v)$, the Chernoff bound cannot be applied to show the concentration of P_v about its expectation. Instead, we apply Azuma's inequality in Lemma 4 to show the following concentration bound for P_v. The vertex-ordering imposed by the permutation π, will play a crucial role in this lemma.

Lemma 4.

$$\Pr\left[\left|P_v - \frac{deg(v)}{e^5}\right| > 27\sqrt{\Delta \log \Delta}\right] < \frac{2}{\Delta^{4.5}}. \tag{7}$$

Proof. Let $U_v = \langle u_1, u_2, \ldots, u_m \rangle$ be the sequence of vertices at distance at most two from v, arranged in increasing rank according to π. Note that $m \leq \Delta^2$. Let S_i indicate whether vertex u_i has decided to go to sleep or not and let C_i denote the tentative color choice of vertex u_i. Let $T_i = (S_i, C_i)$. Clearly, P_v is completely determined by the trials $T_1, T_2, \ldots, T_m$. Referring to Azuma's Inequality, let D denote the absolute difference in conditional expectation $\Big|E[P_v | T_1 = t_1, \ldots, T_i = t_i] - E[P_v | T_1 = t_1, \ldots, T_i = t'_i]\Big|$. Let $t_i = (s_i, c_i)$ and $t'_i = (s'_i, c'_i)$. Provided u_i is a neighbor of v, the difference between t_i and t'_i may contribute at most 1 to the above difference. Any other contribution to this difference is due to vertices u_j such that (i) $j > i$, (ii) u_j is a neighbor of both u_i and v, and (iii) u_j picks as a tentative color either c_i or c'_i. To make this more precise, let S_i be the set $\{u_j \mid j > i \text{ and } u_j \text{ is adjacent to both } u_i \text{ and } v\}$. Thus, S_i is the set of vertices satisfying conditions (i) and (ii) above. For any $u_j \in S_i$, the probability that u_j picks c_i or c'_i as its tentative color is $2/|P(u_j)|$. Note that u_j has to pick one or the other color to many any contribution to D. The quantity $2/|P(u_j)|$ is bounded above by $8/\Delta$, since $|P(u_j)| \geq \Delta/4$. Therefore, the expected contribution of u_j to the quantity D is at most $8/\Delta$. So, the expected contribution of all of S_i to D is bounded above by $8|S_i|/\Delta$. Referring to Azuma's Inequality, we can therefore take $c_i = 1 + 8|S_i|/\Delta$ for each $i : u_i \in N(v)$ and $c_i = 8|S_i|/\Delta$ for all other i. Now note that since S_i is a subset of the set of neighbors of v, $|S_i| \leq \Delta$ for every i. Also, $\sum_{i=1}^{m} |S_i|$ is bounded above by Δ^2. This is because for every vertex $u_j \in S_i$, there is a unique corresponding edge $\{u_i, u_j\}$ in the graph G. Noting that there are at most Δ^2 edges incident on neighbors of v, we get the upper bound on $\sum_{i=1}^{m} |S_i|$. The two inequalities, $|S_i| \leq \Delta$ and $\sum_{i=1}^{m} |S_i| \leq \Delta^2$ together imply that $\sum_{i=1}^{m} |S_i|^2 \leq \Delta^3$. Hence,

$$\sum_{i=1}^{m} c_i^2 = \sum_{i : u_i \in N(v)} \left(1 + \frac{8|S_i|}{\Delta}\right)^2 + \sum_{i : u_i \notin N(v)} \left(\frac{8|S_i|}{\Delta}\right)^2$$

$$\leq \frac{64}{\Delta^2} \cdot \sum_{i=1}^{m} |S_i|^2 + \frac{16}{\Delta} \cdot \sum_{i=1}^{m} |S_i| + \Delta \leq 81\Delta$$

Finally, plugging the value $t = 27 \cdot \sqrt{\Delta \log \Delta}$ and $\sum_{i=1}^{m} c_i^2 \leq 81\Delta$ into Azuma's Inequality (see (2)), we get the desired result.

Let B_v denote the "bad event" that for some color x, vertex v has more than $9\frac{\log \Delta}{\log \log \Delta}$ neighbors colored x *or* $|P_v - \frac{deg(v)}{e^5}| > 27\sqrt{\Delta \log \Delta}$. Using Inequalities (6), (7), and the union bound we get that $\Pr[B_v] < (\Delta + 1)/\Delta^6 + 2/\Delta^{4.5} \leq 4/\Delta^{4.5}$. If vertices u and v are more than four hops away from each other in G, then B_u and B_v are mutually independent. Therefore, B_v is independent of all except at most Δ^4 other bad events. Applying the Lovász Local Lemma (see Lemma 2) yields the theorem.

Remark on constants. In this section we have explicitly specified constants so that a careful reader may verify our calculations completely. However, we have not attempted to optimize these constants. For example, it is possible to show that Δ_i and p_i decay at a rate of $(1 - \frac{1}{e^{1+\epsilon}})$ (rather than $(1 - \frac{1}{e^5})$) for any $\epsilon > 0$. The choice of ϵ here affects various constants in the proof including the constant in the asymptotic notation used to specify the frugality of the coloring.

3 Weighted Almost-Equitable Colorings

In this section we prove result (E1). Let $G = (V = [n], E)$ be a vertex-weighted graph with weight function $w : V \to [0, 1]$. Let π be a permutation of $(1, 2, \ldots, n)$ picked uniformly at random from the set of all permutations of $(1, 2, \ldots, n)$. Run one round of our random coloring procedure described in Section 2, but without any nodes falling asleep. For any color x and for any vertex v let $P(v, x)$ denote the event that v is (permanently) colored x, at the end of one round of the random coloring procedure.

Lemma 5. $\Pr[P(v, x)] \geq \left(1 - \frac{1}{e}\right) \frac{1}{\Delta+1}$.

Proof. Let $L(v)$ be the subset of neighbors of vertex v that are ranked before v by π. Then, $\Pr[P(v, x)] = \frac{1}{\Delta+1} \sum_{j=0}^{deg(v)} \frac{1}{deg(v)+1} \cdot \left(1 - \frac{1}{(\Delta+1)}\right)^j$. In this expression, the term "$\frac{1}{(\Delta+1)}$" is the probability that color x is tentatively picked by v, the summation is over all possible sizes of $L(v)$, the term "$\frac{1}{deg(v)+1}$" is the probability that exactly j of the neighbors of v are ranked below v in π, and the term "$(1 - \frac{1}{(\Delta+1)})^j$" is the probability that none of these j neighbors that are ranked below v, pick x as their tentative color. This expression can be simplified to yield

$$\Pr[P(v, x)] = \frac{1}{deg(v) + 1} \left(1 - \left(1 - \frac{1}{(\Delta + 1)}\right)^{deg(v)+1}\right).$$

Noting that the r.h.s. above achieves the minimum value at $deg(v) = \Delta$, we obtain

$$\Pr[P(v, x)] \geq \frac{1}{\Delta + 1} \left(1 - \left(1 - \frac{1}{(\Delta + 1)}\right)^{\Delta+1}\right) \geq \frac{1}{\Delta + 1} \cdot \left(1 - \frac{1}{e}\right).$$

The above lemma implies that the expected weight of any color class is at least $(1 - 1/e) \cdot W/(\Delta + 1)$. We now show that the weight of any color class is concentrated around its expectation and in particular the weight of a color class is much smaller than $(1 - 1/e) \cdot W/(\Delta + 1)$ only with small probability. As usual, we employ Azuma's Martingale inequality to prove this result; again the symmetry-breaking approach plays a critical role.

Lemma 6. *Let $x \in [\Delta + 1]$ and let W_x denote the total weight of vertices colored x. Then for any fixed $c > 0$,*

$$\Pr\left[W_x < \left(1 - \frac{1}{e}\right) \cdot \frac{W}{(\Delta + 1)} - c\sqrt{W \log(\Delta + 1)}\right] \leq \frac{1}{(\Delta + 1)^{c^2/18}}.$$

Proof. Let T_i be the tentative choice of a color by the vertex of rank i in π. W_x is completely determined by the outcomes of the trials $T_1, T_2, \ldots, T_n$. Now consider the difference in conditional expectations from Azuma's Martingale Inequality:

$$c_i = \Big|E[W_x \mid T_1 = t_1, T_2 = t_2, \ldots, T_i = t_i] - E[W_x \mid T_1 = t_1, T_2 = t_2, \ldots, T_i = t'_i]\Big|.$$

Let v be the vertex of rank i in π. The above difference is at most

$$c_i \leq w(v) + \frac{2}{\Delta + 1} \sum_{u \in U(v)} w(u),$$

where $U(v)$ is the set of neighbors of v ranked higher than v in π. The first term "$w(v)$" in the above bound is due to the change in the tentative color of v from t_i to t'_i. The second term is the *expected* change in W_x; note that this occurs only at vertices in $U(v)$. Using the fact that every vertex weight is in $[0, 1]$, we get that $c_i \leq 1 + 2\Delta/(\Delta + 1) < 3$. Also,

$$\sum_{i=1}^{n} c_i \leq \sum_{v \in V} w(v) + \frac{2}{\Delta + 1} \sum_{v \in V} \sum_{u \in U(v)} w(u) \leq W + \frac{2\Delta}{\Delta + 1} \cdot \sum_{v \in V} w(v) < 3W.$$

Finally, $\sum_{i=1}^{n} c_i^2 < 3 \sum_{i=1}^{n} c_i < 9W$. Therefore, the bound in Azuma's Inequality simplifies to

$$\exp\left(\frac{-t^2}{2\sum_{i=1}^{n} c_i^2}\right) \leq \exp\left(\frac{-c^2 W \log(\Delta + 1)}{18W}\right) = \frac{1}{(\Delta + 1)^{c^2/18}}.$$

From this we get, using Azuma's Inequality,

$$\Pr\left[W_x < \left(1 - \frac{1}{e}\right) \cdot \frac{W}{\Delta + 1} - c \cdot \sqrt{W \log(\Delta + 1)}\right] < \frac{1}{(\Delta + 1)^{c^2/18}}.$$

Since we have $\Delta + 1$ color classes, using the union bound we get that

$$\Pr\left[\exists x \in [\Delta + 1] : W_x < \left(1 - \frac{1}{e}\right) \cdot \frac{W}{\Delta + 1} - c \cdot \sqrt{W \log(\Delta + 1)}\right] \leq \frac{(\Delta + 1)}{(\Delta + 1)^{c^2/18}}.$$

Choosing $c \geq 5$ guarantees that with positive probability, for every color $x \in [\Delta + 1]$, the weight W_x of the vertices colored x is at least $\left(1 - \frac{1}{e}\right) \cdot \frac{W}{\Delta+1} - c \cdot \sqrt{W \log(\Delta + 1)}$. This proves result (E1).

Acknowledgment. We thank Mike Molloy for his helpful comments and encouragement, and the referees for their comments that helped improve the paper.

References

1. B. Baker, E. Coffman, Mutual exclusion scheduling. *Theor. Comput. Sci.* **162** (1996), 225–243.
2. J. Beck. An algorithmic approach to the Lovász Local Lemma. *Random Structures and Algorithms*, 2: 343–365, 1991.
3. J. Blazewicz, K. Ecker, E. Pesch, G. Schmidt, J. Weglarz, *Scheduling computer and manufacturing processes. 2nd ed.*, Berlin: Springer. 485 p. (2001).
4. B. Bollobás and R.K. Guy. Equitable and proportional coloring of trees, *Journal of Combinatorial Theory, Series B*, 34: 177–186, 1983.
5. H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *American Statistical Association Journal*, 58:13–30, 1963.
6. P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and Finite Sets*, A. Hajnal et. al., eds., Colloq. Math. Soc. J. Bolyai 11, North Holland, Amsterdam, 1975, pp. 609–627.
7. D. A. Grable and A. Panconesi. Fast distributed algorithms for Brooks-Vizing colorings. *Journal of Algorithms*, 37: 85–120, 2000.
8. A. Gupta, R. Krauthgamer and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proc. IEEE Symposium on Foundations of Computer Science*, 2003.
9. A. Hajnal and E. Szemerédi, Proof of a conjecture of Erdős, in: *Combinatorial Theory and its Applications, Vol. II* (P. Erdős, A. Rényi and V. T. Sós, Editors), (North-Holland, 1970), 601–603.
10. H. Hind, M. Molloy, and B. Reed. Total colouring with $\Delta + \text{polylog}(\Delta)$ colours. *SIAM Journal on Computing*, 28:816–821, 1998.
11. Y. T. Hou, V. S. A. Kumar, M. V. Marathe, and A. Srinivasan. Personal communication, 2006.
12. S. Irani, V. Leung, Scheduling with conflicts, and applications to traffic signal control, *Proceedings of the 7th Annual ACM-SIAM symposium on discrete algorithms, held in Atlanta, GA, 1996*, Philadelphia, PA: SIAM, 85–94 (1996).
13. S. Janson and A. Ruciński, The infamous upper tail, *Random Structures and Algorithms*, **20** (2002), 317–342.
14. A. Johansson. Asymptotic choice number for triangle free graphs. In *DIMACS Technical Report*, 91-5, 1996.
15. Ö. Johansson. Simple distributed $\Delta + 1$-coloring of graphs. *Information Processing Letters*, 70: 229–232, 1999.
16. J. Kahn. Asymptotically good list-colorings. *J. Combinatorial Theory, Series A*, 73:1–59, 1996.
17. H. A. Kierstead and A. V. Kostochka, A Short Proof of the Hajnal-Szemeredi Theorem on Equitable Coloring. To appear, *Combinatorics, Probability, and Computing.*
18. J. H. Kim. On Brooks' Theorem for Sparse Graphs. *Combinatorics, Probability, and Computing*, 4: 97–132, 1995.
19. A. V. Kostochka, K. Nakprasit, and S. V. Pemmaraju. On Equitable Coloring of d-Degenerate Graphs. *SIAM J. Discret. Math.*, 19 (1): 83–95, 2005.
20. J. Krarup, D. de Werra, Chromatic optimisation: Limitations, objectives, uses, references, *Eur. J. Oper. Res.*, **11** (1982), 1–19.
21. R. Krauthgamer and J. R. Lee. Navigating nets: simple algorithms for proximity search. In *Proc. ACM-SIAM Symposium on Discrete Algorithms*, 2004.
22. M. Luby. Removing randomness in parallel computation without a processor penalty. *Journal of Computer and System Sciences*, 47(2): 250–286, 1993.
23. M. Molloy and B. Reed. *Graph Colouring and the Probabilistic Method*, Springer, 2000.
24. S. V. Pemmaraju, Equitable colorings extend Chernoff-Hoeffding bounds, *Proceedings of the 5th International Workshop on Randomization and Approximation Techniques in Computer Science (APPROX-RANDOM 2001)*, 2001, 285–296.
25. B. F. Smith, P. E. Bjorstad, and W. D. Gropp, *Domain decomposition. Parallel multilevel methods for elliptic partial differential equations*, Cambridge: Cambridge University Press, 224 p. (1996).
26. A. Tucker, Perfect graphs and an application to optimizing municipal services, *SIAM Review*, 15(1973), 585–590.