

Acyclic Edge Coloring through the Lovász Local Lemma[☆]

Ioannis Giotis^{a,c}, Lefteris Kirousis^{a,c,*}, Kostas I. Psaromiligkos^a,
Dimitrios M. Thilikos^{a,b,c}

^a*Department of Mathematics, National & Kapodistrian University of Athens, Greece*

^b*ALGCo project-team, CNRS, LIRMM, Montpellier, France*

^c*Computer Technology Institute & Press “Diophantus”, Patras, Greece*

Abstract

We give a probabilistic analysis of a Moser-type algorithm for the Lovász Local Lemma (LLL), adjusted to search for acyclic edge colorings of a graph. We thus improve the best known upper bound to acyclic chromatic index, also obtained by analyzing a similar algorithm, but through the entropic method (basically counting argument). Specifically we show that a graph with maximum degree Δ has an acyclic proper edge coloring with at most $\lceil 3.74(\Delta - 1) \rceil + 1$ colors, whereas, previously, the best bound was $4(\Delta - 1)$. The main contribution of this work is that it comprises a probabilistic analysis of a Moser-type algorithm applied to events pertaining to dependent variables.

Keywords: Acyclic edge coloring, Algorithmic proof of the Lovász Local Lemma

2010 MSC: 05D40, 05C15

1. Introduction and the basic algorithm

Let $G = (V, E)$ be a (simple) graph with l vertices and m edges. The chromatic index of G is the least number of colors needed to *properly* color its edges, i.e., to color them so that no adjacent edges get the same color. If Δ is the maximum degree of G , it is known that its chromatic index is either Δ or $\Delta + 1$ (Vizing [1]).

A cycle of G of length s is a sequence $v_i, i = 0, \dots, s-1$ of distinct vertices so that $\forall i = 0, \dots, s-1$, v_i and $v_{i+1 \pmod s}$ are connected by an edge. The *acyclic* chromatic index of G (ACI) is defined as the least number of colors needed to properly color the edges of G so that no cycle is *bichromatic*, i.e., so that there is no cycle whose edges are properly colored with only two colors. Notice that in any properly colored graph, any cycle of odd length is necessarily at least

[☆]Research co-financed by the European Union (European Social Fund ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF) Research Funding Program: ARISTEIA II.

*Corresponding author

Email addresses: igiotis@cs.upc.edu (Ioannis Giotis), lkirousis@math.uoa.gr (Lefteris Kirousis), kostaspsa@gmail.com (Kostas I. Psaromiligkos), sedthilk@thilikos.info (Dimitrios M. Thilikos)

trichromatic, i.e., its edges have three or more colors. It has been conjectured (J. Fiamčík [2] and Alon et al. [3]) that the acyclic chromatic index of any graph with maximum degree Δ is at most $\Delta + 2$. A number of successively tighter upper bounds to the acyclic chromatic index have been provided in the literature. Most recently, Esperet and Parreau [4] proved that the acyclic chromatic index is at most $4(\Delta - 1)$. Their proof makes use of the technique of Moser and Tardos [5] that constructively proves the Lóvasz Local Lemma (a technique which became known as the “entropy compression method” [6]). An approach using the entropy compression method was also used for the vertex analogue of the edge chromatic number by Gonçalves et al. [7].

In this work, we modify the technique used by Esperet and Parreau [4] in that for a Moser-type edge coloring algorithm, we use as tool of analysis the approach we described in Section 2 of [8]. Namely, instead of an essentially counting argument as used in the entropy compression method, we give a probabilistic analysis that yields an upper bound of $\lceil 3.74(\Delta - 1) \rceil + 1$ for the acyclic chromatic index, improving over $4(\Delta - 1)$ in [4] (in contrast to the paper by Moser and Tardos [5], a probabilistic analysis was used in the original paper of Moser [9]; see the elegant exposition by Spencer in [10]). The present paper can be read independently of [8].

An interesting aspect of this application is that the edge colors to which the “undesirable” events of LLL refer to are not probabilistically independent. This dependence introduces certain conceptual difficulties, some of which are, we believe, non-trivial (see the proof of Lemma 5 and the preceding remarks). The randomized algorithm that we deal with allows the freedom, when coloring an edge, to make a selection, uniformly at random, from a guaranteed minimum number of available colors. However, the “guarantee” of a minimum number of available colors is valid only if any conditioning refers only to colors previously assigned. To handle the probabilistic analysis of such a randomized algorithm without introducing posterior probabilities, which would render the analysis unmanageable, we put all events referring to colors that edges have in chronological order according to the instant these edges got their current color. We consider this approach of handling dependent events in constructive proofs of LLL, rather than just the improvement of the coefficient of the upper bound from 4 to 3.74, as the essential aspect of the contribution of this work (nevertheless, see the discussion in Section 3 for possible further numerical improvement).

We also get improved numerical results with respect to graphs with bounded girth, some specific values of which are sampled in Figure 1.

Girth	Number of colors	Previously known [4]
-	$3.731(\Delta - 1) + 1$	$4(\Delta - 1)$
7	$3.326(\Delta - 1) + 1$	$3.737(\Delta - 1)$
53	$2.494(\Delta - 1) + 1$	$3.135(\Delta - 1)$
219	$2.323(\Delta - 1) + 1$	$3.043(\Delta - 1)$

Figure 1: Our results

Below, to facilitate notation, we call a proper edge-coloring *s-acyclic* if it contains no bichromatic cycle of length s or less. We call the corresponding graph parameter the *s-acyclic chromatic index*.

We start by mentioning the following fact, proved in Esperet and Parreau [4]:

Lemma 1 (Esperet and Parreau [4]). *At any step of any successive coloring of the edges of a graph, there are at most $2(\Delta - 1)$ colors that should be avoided in order to produce a 4-acyclic coloring.*

Proof Sketch. Notice that for each edge e , one has to avoid the colors of all edges adjacent to e , and moreover for each pair of homochromatic (of the same color) edges e_1, e_2 adjacent to e at different endpoints (which contribute one to the count of colors to be avoided), one has also to avoid the color of the at most one edge e_3 that together with e, e_1, e_2 define a cycle of length 4. So the total count of colors to be avoided adds up to $2(\Delta - 1)$. \square

Assume now that we have $K = \lceil (2 + \gamma)(\Delta - 1) \rceil + 1$ colors, where γ is a nonnegative constant to be computed.

We assume below that the edges of the graph, and its cycles, are ordered according to fixed a priori orderings.

Notice that for each cycle of the graph there are two consecutive traversals of its edges. It does not matter which we use, but for concreteness when we start the traversal from an edge e , the next edge to be traversed is the least one from the two adjacent to e . We call this traversal “positive”.

Also in an even length cycle, the edges can be partitioned into two subsets of equal cardinality, the elements of each of which have pairwise odd distance. These sets are called equal parity sets. If we color such a cycle with two colors so that no adjacent edges get the same color, then the equal parity sets are the *monochromatic* sets, i.e. the sets whose respective elements get the same color.

Consider now the algorithm given in Figure 2.

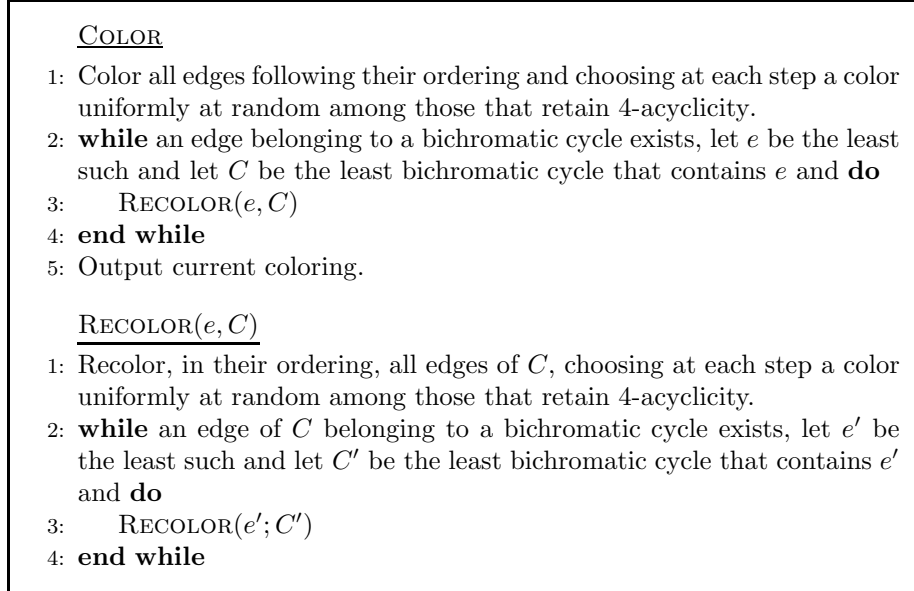


Figure 2: *The coloring algorithm*

A *root* call of RECOLOR is a call invoked in line 3 of COLOR. A *phase* is the collection of steps involved in the execution of a call of RECOLOR (root or

not) or the collection of steps executed in line 1 of COLOR. The latter phase is called the initial phase (phases are nested). The time complexity of COLOR is expressed in terms of the number of invocations (root or not) of RECOLOR, i.e. the number of phases.

Notice that in this algorithm, each non-root invocation of RECOLOR chooses a cycle sharing an edge with the cycle of the calling RECOLOR. So this algorithm is in the spirit of the algorithm in the original paper by Moser [9]; in the subsequent algorithm by Moser and Tardos [5], the event chosen at the end of each phase does not necessarily share a variable with the event of the previous phase.

The following follows easily from line 2 of COLOR.

Lemma 2. *COLOR outputs an acyclic edge coloring if it ever stops.*

The main result of this paper, Theorem 2, states that if there are at least $3.74(\Delta - 1) + 1$ available colors, then the probability that COLOR lasts for at least n phases is < 1 for sufficiently large n (actually, Theorem 2 states that this probability is subexponential), therefore there is an acyclic edge coloring.

Lemma 3. *Consider an arbitrary call of RECOLOR(e, C). Let \mathcal{E} be the set of edges that do not belong to a bichromatic cycle at the beginning of this call together with the set of edges of C . Then, if the call terminates, the edges in \mathcal{E} do not belong to any bichromatic cycle at the end of RECOLOR(e, C).*

Proof. Consider an edge $e' \in \mathcal{E}$. Assume first that e' does not belong to C . Then, if e' is in a bichromatic cycle C' after the execution of RECOLOR(e, C), that cycle must have become bichromatic during the execution. Therefore there is a cycle C^- and an edge of it e^- such that the process RECOLOR($e^-; C^-$) was called as a recursive call of RECOLOR(e, C) and C' and C^- share an edge and C' became bichromatic during RECOLOR($e^-; C^-$). But the call RECOLOR($e^-; C^-$) will not terminate until all edges in C^- do not belong to a bichromatic cycle. Assuming RECOLOR(e, C) terminates, RECOLOR($e^-; C^-$) must have also terminated. So C' cannot be bichromatic at the end of RECOLOR($e^-; C^-$). The same argument can be reapplied every time e changes color and becomes an element of a bichromatic cycle during RECOLOR(e, C).

If, on the other hand, $e' \in C$, then by definition e' does not belong to any bichromatic cycle at the end of RECOLOR(e, C). \square

As an immediate corollary we get:

Corollary 1. *The cycles of root calls of RECOLOR have pairwise distinct sets of edges. Therefore the number of root phases is at most m , the number of edges of the graph.*

We define labeled forests to be rooted forests whose nodes are labeled with pairs (e, C) , where e is an edge, called the edge-label, and C is an even cycle of half-length ≥ 3 that contains e , called the cycle-label. We consider such labeled forests as ordered by ordering the set of roots and each set of siblings (children of the same parent) according to the order of their respective edge-label. We also define:

Definition 1. *A labeled forest is called feasible if*

- i. Pairwise, the cycle-labels of the roots of the trees do not share an edge,
- ii. pairwise the cycle-labels of the children of every node do not share an edge,
- iii. if C is the cycle-label of a node u , the edge-label of any child of u belongs to C .

Feasible forests are intended to represent the structure of recursive calls to RECOLOR. For technical reasons, we add to a feasible forest some new leaves to which we assign just an edge-label (their cycle-label can be taken to be an empty cycle): First we add new trees comprised of a root only, so that the set of edge-labels of all the roots of the trees of forest becomes equal to the set of edges of the graph; second, we hang from each original leaf v of \mathcal{F} as many new leaves as the edges of the cycle-label of v , and we label them so that the set of their edge-labels coincides with the set of edges of the cycle-label of v . So we assume in the sequel that there are exactly m roots whose edge-labels comprise the set of all edges, and that any internal node with a cycle-label C of half-length k has exactly $2k$ children whose edge-labels comprise the set of all edges in C . Traversing the internal nodes of each tree of a feasible forest in pre-order (depth-first), and visiting the trees in the order of their roots' edge-labels, we obtain the forest's label-sequence

$$\mathcal{L} = (e_1, C_1), \dots, (e_n, C_n)$$

(labels of leaves are not included in the label-sequence).

Definition 2. *The n -witness forest (or just the witness forest, when n is clear from the context) of an execution of COLOR with at least n phases is the feasible forest that results by creating one node per each of the n invocations of RECOLOR, labeling it by its argument, and structuring the trees of the forest as the calls of RECOLOR appear in the recursion stack of each root phase, i.e. a node labeled with (e_2, C_2) is a progeny of a node labeled with (e_1, C_1) if RECOLOR(e_2, C_2) is called while RECOLOR(e_1, C_1) is executed (additional leaves as described in the previous paragraph are also added).*

2. A Bound for the Acyclic Chromatic Index

Towards finding an upper bound that COLOR lasts for at least n phases, we consider the algorithm COLORVAL in Figure 3 that takes as input an arbitrary sequence

$$\mathcal{S} = (e_1^1, e_1^2, k_1), \dots, (e_n^1, e_n^2, k_n),$$

such that for all $s = 1, \dots, n$, e_s^1 and e_s^2 are adjacent edges contained in some cycle of half-length $k_s \geq 3$, and appear in this order in the cycle's positive traversal, when we start from e_s^1 . We call such sequences *admissible sequences*.

Observe that COLORVAL on input an arbitrary admissible sequence \mathcal{S} , with $|\mathcal{S}| = n$, always produces as output a uniquely defined sequence of cycles $\mathcal{C} = C_1, \dots, C_n$ (but, alas, may end up with success = **false**). The notion of “phases” for COLORVAL is defined in analogy to COLOR, i.e. phases are comprised of the number of steps that are involved with the execution of line 1 (initial phase) or with the execution of a loop of line 3 of COLORVAL (phase s).

COLORVAL (\mathcal{S} admissible sequence $(e_s^1, e_s^2, k_s), s = 1, \dots, n$)

- 1: Color all edges of the graph, following their ordering, choosing at each step uniformly at random a color that retains 4-acyclicity.
- 2: success := **true**
- 3: **for** $s = 1, \dots, n$ **do**
- 4: **if** there is a cycle C_s that:
 - (a) has length $2k_s$ and contains e_s^1, e_s^2 in this order in its positive traversal,
 - (b) is bichromatic under the current coloring
- 5: **then** let C_s be the unique such cycle /*the uniqueness is due to the fact that no distinct homochromatic edges can be coincident to the same vertex*/
- 6: **else** success := **false**; let C_s be an arbitrary (e.g. the first) cycle that contains e_s^1, e_s^2 in this order in its positive traversal, starting from e_s^1
- 7: **end if**
- 8: Recolor all edges of C_s , in their ordering, choosing at each step a color uniformly at random among those that retain 4-acyclicity
- 9: **end for**

Figure 3: *The coloring validation algorithm*

Given a feasible forest \mathcal{F} whose label-sequence is $(e_s, C_s), s = 1, \dots, n$ the corresponding admissible sequence $\mathcal{S}_{\mathcal{F}}$ is obtained by letting e_s^1 be e_s , e_s^2 be the edge of C_s following e_s in C_s 's positive traversal and k_s be the half-length of C_s , $s = 1, \dots, n$. Given an admissible sequence \mathcal{S} let $\mathfrak{F}_{\mathcal{S}}$ be the class of all feasible forests \mathcal{F} such that $\mathcal{S}_{\mathcal{F}} = \mathcal{S}$.

Proposition 1. *Given an admissible sequence \mathcal{S}*

$$\sum_{\mathcal{F} \in \mathfrak{F}_{\mathcal{S}}} \Pr[\text{COLOR executes with witness forest } \mathcal{F}] \leq \Pr[\text{COLORVAL is successful on input } \mathcal{S}]. \quad (1)$$

Proof. We prove first that the probability P for at least one $\mathcal{F} \in \mathfrak{F}_{\mathcal{S}}$ being a witness forest of COLOR is bounded from above by the probability that COLORVAL is successful on input \mathcal{S} . For this it is sufficient to notice that if the random choices made by an execution of COLOR that produces an arbitrary $\mathcal{F} \in \mathfrak{F}_{\mathcal{S}}$ are made by COLORVAL on input $\mathcal{S}_{\mathcal{F}} = \mathcal{S}$, then COLORVAL is successful. The result now follows by observing the the events that \mathcal{F} is a witness forest of COLOR for various $\mathcal{F} \in \mathfrak{F}_{\mathcal{S}}$ are mutually exclusive, therefore P can be written as the sum in the lhs of (1). \square

Given an admissible sequence $\mathcal{S} = (e_1^1, e_1^2, k_1), \dots, (e_n^1, e_n^2, k_n)$, an arbitrary sequence $\mathcal{C} = C_1, \dots, C_n$ of cycles such that $C_s, s = 1, \dots, n$, has length $2k_s$ and contains e_s^1, e_s^2 , in this order in its positive traversal, starting from e_s^1 , is called a

cycle-sequence *associated* with the admissible \mathcal{S} (a cycle-sequence \mathcal{C} associated with \mathcal{S} need not necessarily be the cycle-sequence produced by COLORVAL).

Given a cycle-sequence $\mathcal{C} = (C_s)_s$ associated with an admissible sequence $\mathcal{S} = (e_s^1, e_s^2, k_s)_s$, let the other edges of C_s in its positive traversal be $e_s^3, \dots, e_s^{2k_s}$ and let $v_s^i, i = 1, \dots, 2k_s$, be the vertices of C_s with $v_s^i, v_s^{i+1 \bmod 2k_s}$ being the endpoints of e_s^i . Finally let $2\hat{k} = \sum_{s=1}^n (2k_s - 2)$.

For linguistic convenience, we call the edges e_s^1, e_s^2 the *first and second pivotal edges* of C_s , respectively (these are common to all \mathcal{C}). Let $e_s^{l_s^1}$ (respectively, $e_s^{l_s^2}$) be the edge of C_s that gets the color it has at phase s of COLORVAL earliest among the edges of C_s that have equal parity with e_s^1 (e_s^2 , respectively). We call $e_s^{l_s^1}, e_s^{l_s^2}$ the *early edges* of C_s .

We call *instants* of an execution of COLORVAL the discrete successive values $t = 1, 2, \dots$ of a time parameter at which COLORVAL assigns a color to an edge. Reassigning the same color is assumed to take place at a new instant.

Call the color an edge has immediately before the s 'th repetition of loop 3 the edge's color at phase s . Also call an edge's color at phase 1 the edge's *initial* color. Given an edge e and a phase number s , let $\text{time}(s; e)$ be the latest instant before phase s starts that e was assigned the color it has at phase s , and let $\chi(s; e)$ be this color.

Remark 1. Notice that in general it is possible to have two distinct phases s', s and an edge e so that $\text{time}(s'; e) = \text{time}(s; e)$, because an edge may not be assigned a color for the duration of several phases; however if $C_{s'}$ is the cycle produced by an execution of COLORVAL at phase s' , $s' < s$, and e belongs to $C_{s'}$, then $\text{time}(s'; e) < \text{time}(s; e)$, because the cycles of the output cycle-sequence are recolored at each repetition of loop 3.

Let c_s^1, c_s^2 be the colors that the early edges $e_s^{l_s^1}, e_s^{l_s^2}$, respectively, of a cycle C_s have at phase s i.e. $c_s^j = \chi(s; e_s^{l_s^j}), j = 1, 2$.

Remark 2. To avoid a possible misinterpretation, let us stress here that we do not assume that c_s^1, c_s^2 are a priori fixed and that the executions considered are conditioned on the early edges of C_s taking the colors c_s^1, c_s^2 , respectively. Rather, given an arbitrary execution, c_s^1, c_s^2 denote the colors the early edges will have at phase s . Much like the Monty Hall problem [11], where the assumption that Monty Hall opens box A does not mean that the choice of the contestant is conditioned on Monty Hall always opening box A. "Box A" is just a "name" for whatever box Monty Hall will open. Its actual value depends on the box selected by the contestant.

Given a phase number s a color c and an edge e (belonging to C_s or not), the event of *assignment of color c* for phase s (CA-event, in short), notationally $\text{CA}(s, c; e)$, is the event that occurs if the color assigned to e by COLORVAL at instant $t = \text{time}(s; e)$ is c . If e belongs to the cycle C_s , i.e. if $e = e_s^i$ for some i , then the CA-event $\text{CA}(s, c; e)$ for $c = c_s^{((i+1) \bmod 2)+1}$ is called the *correct color assignment event* (CCA-event, in short), notationally $\text{CCA}(C_s; e)$. In other words, a CCA-event occurs when the edge of its argument takes the color of the corresponding early edge of the cycle of its argument. We refer to the color $c_s^{((i+1) \bmod 2)+1}$ for the edge e_s^i as e_s^i 's "correct" color for phase s .

The chronological order of the edges e_s^i of a cycle sequence \mathcal{C} is the order induced by the ordering of $\text{time}(s; e_s^i)$ (recall Remark 1).

Definition 3. An anterior conditional for an event $\text{CA}(s, c; e)$ is any conjunction of events $\text{CA}(s', c'; e')$, or negations of them, such that $\text{time}(s'; e') < \text{time}(s; e)$.

Lemma 4. The probability of any CA-event, not referring to a cycle's early edge and its correct color, given any anterior conditional, is at most

$$\frac{1}{\gamma(\Delta - 1) + 1}.$$

Proof. It follows immediately from the fact that COLORVAL assigns a color chosen uniformly at random from the set of colors that do not presently destroy 4-acyclicity and that this set, by Lemma 1, has cardinality at least $\gamma(\Delta - 1) + 1$, independently of past assignments (recall that the total number of available colors is $K = \lceil (2 + \gamma)(\Delta - 1) \rceil + 1$). \square

Remark 3. [Identifying the pivotal edges with their corresponding early edges] Observe that by definition, the CCA-events corresponding to the early edges of cycles occur with probability 1. To keep the notation simple, we will consider CCA-events $\text{CCA}(C_s; e_s^i)$ for all $i \geq 3$ but with the notational convention that for $j = 1, 2$, $\text{CCA}(C_s; e_s^{l^j})$ stands for $\text{CCA}(C_s; e_s^j)$, i.e. the CCA-event of an early edge is, by convention, just the CCA-event of the corresponding pivotal edge.

Lemma 5. Given an admissible sequence $\mathcal{S} = (e_s^1, e_s^2, k_s), s = 1, \dots, n$, then for the probability of success of COLORVAL on input \mathcal{S} we have that:

$$\Pr[\text{COLORVAL is successful on } \mathcal{S}] \leq \left(\frac{1}{\gamma(\Delta - 1) + 1} \right)^n \prod_{s=1}^n \left(1 - \left(1 - \frac{1}{\gamma(\Delta - 1) + 1} \right)^{\Delta-1} \right)^{2k_s-3}. \quad (2)$$

Proof. Given a cycle-sequence \mathcal{C} associated with \mathcal{S} , the edges incident on v_s^i other than e_s^{i-1} are called these edges stemming out v_s^i . Let these edges in chronological order be $o^1, \dots, o^{\delta_s^i}$, where δ_s^i denotes their cardinality (the o^l 's themselves depend on s and i , however for brevity we omit the corresponding indices). Obviously $\delta_s^i \leq \Delta - 1$, where Δ is the max degree of the underlying graph G . When e_s^{i-1} is the previous to last edge of C_s in its positive traversal, we set $\delta_s^i = 1$. Also, let the *parity* of a non-pivotal edge e_s^i of C_s be 1 or 2 depending on whether e_s^i is at odd distance from e_s^1 or e_s^2 .

Let $\mathcal{E}_C^1, \dots, \mathcal{E}_C^{2\hat{k}}$ denote the CCA-events $\text{CCA}(C_s; e_s^i), 1 \leq s \leq n, i = 3, \dots, 2k_s$ of \mathcal{C} in *chronological order* (by Remark 1, we may assume that these events correspond to distinct time instants). Let \mathcal{A}_t denote the conditional that the chronologically first $t - 1$ CCA-events of \mathcal{C} hold (\mathcal{A}_t depends on the execution).

The *geometric order* of the edges of \mathcal{C} is the one where the edges in each cycle are ordered as they are traversed in its positive traversal, whereas the cycles are ordered by their index (corresponding phase number). Assume that the CCA-event for the k 'th such edge ($k = 1, \dots, 2\hat{k}$) in the geometric order is $\mathcal{E}_C^{t_k}$. In other words, $\mathcal{E}_C^{t_k}, t = 1, \dots, 2\hat{k}$ is a reordering, in the sense of the geometric order, of the chronological order $\mathcal{E}_C^t, t = 1, \dots, 2\hat{k}$ (the reordering depends on the execution).

We now define a random variable, denoted by \mathcal{C} , over the random color-choices of COLORVAL. \mathcal{C} is the unique cycle-sequence such that for every cycle C_s of \mathcal{C} , every non-pivotal edge e_s^i is either the (unique) edge stemming out of v_s^i that, at the beginning of phase s , has the same color as the pivotal edge of C_s of the same parity, if there is such one, or is the first edge stemming out of v_s^i in some predetermined order of all edges of the underlying graph, otherwise. This is a well defined function over the space of random color-choices of COLORVAL because we cannot have two homochromatic edges incident onto the same vertex. It is admittedly somewhat confusing that until now the notation \mathcal{C} referred to an arbitrary but fixed cycle-sequence, whereas in the sequel it denotes a random variable depending on the random choices of COLORVAL, yet we believe that we thus avoid overloading the notation. Once \mathcal{C} is seen as an well-defined random variable, all the parameters defined up to now in terms of a fixed arbitrary \mathcal{C} , like e_s^i (the i -th edge in the positive traversal of the s -th cycle of \mathcal{C}), δ_s^i (the number of edges stemming out of the v_s^i), o^l (the chronologically l -th edge stemming out of v_s^i) etc. become well-defined random variables, depending only on the random choices of COLORVAL and its input \mathcal{S} . Also the events previously described remain meaningful with \mathcal{C} being the random variable defined above instead of arbitrary and fixed.

Now, observe that:

$$\begin{aligned} \Pr[\text{COLORVAL is successful on } \mathcal{S}] &= \Pr[\mathcal{C} \text{ is bichromatic}] \\ &= \prod_{t=1}^{2\hat{k}} \Pr[\mathcal{E}_{\mathcal{C}}^t \mid \mathcal{A}_t] = \prod_{k=1}^{2\hat{k}} \Pr[\mathcal{E}_{\mathcal{C}}^{t_k} \mid \mathcal{A}_{t_k}]. \quad (3) \end{aligned}$$

Indeed, we obviously have that the event “COLORVAL is successful on \mathcal{S} ” is implied by the event “ \mathcal{C} is bichromatic” (\mathcal{C} is said to be bichromatic if for all s , C_s is bichromatic at the beginning of phase s). Also because no two homochromatic edges can be incident to the same vertex, we have that the latter event is implied by the former, so we get the first equality of Equation 3. The second equality follows by Remark 3 and the third one is trivial.

In order to find an upper bound for $\Pr[\mathcal{E}_{\mathcal{C}}^{t_k} \mid \mathcal{A}_{t_k}]$, first assume that for some $s = 1, \dots, n, i = 3, \dots, 2k_s$, the event $\mathcal{E}_{\mathcal{C}}^{t_k}$ is the event $\text{CCA}(C_s; e_s^i)$. For brevity let δ_k denote δ_s^i and c_k the correct color for e_s^i .

Let now E^l be the event “ $(e_s^i = o^l) \wedge \text{CA}(s; c_k, o^l)$ ”, $l = 1, \dots, \delta_k$. Intuitively, E^l means that e_s^i is the chronologically l -th edge stemming out of v_s^i and it gets the correct color. In case e_s^i is a non-pivotal early edge of a cycle C_s (i.e. the earliest edge among the same parity edges of C_s), then E^l stands for “ $(e_s^i = o^l) \wedge \text{CA}(s; c_k, e)$ ”, where e is the pivotal edge e_s^1 or e_s^2 of the same parity as e_s^i and c_k is e_s^i 's color (recall Remark 3).

We have:

$$\Pr[\mathcal{E}_{\mathcal{C}}^{t_k} \mid \mathcal{A}_{t_k}] = \Pr[\text{CA}(s, c_k; e_s^i) \mid \mathcal{A}_{t_k}] = \Pr\left[\bigvee_{l=1}^{\delta_k} E^l \mid \mathcal{A}_{t_k}\right], \quad (4)$$

because always (deterministically) e_s^i is one of the edges that stem out of v_s^i . Intuitively, $\Pr\left[\bigvee_{l=1}^{\delta_k} E^l \mid \mathcal{A}_{t_k}\right]$ denotes the probability that one of the edges stemming out of e_s^i gets the correct color.

Observe that:

$$\begin{aligned} \Pr \left[\bigvee_{l=1}^{\delta_k} E^l \mid \mathcal{A}_{t_k} \right] &= 1 - \Pr \left[\bigwedge_{l=1}^{\delta_k} (\neg E^l) \mid \mathcal{A}_{t_k} \right] \\ &= 1 - \prod_{l=1}^{\delta_k} \Pr \left[\neg E^l \mid \left(\bigwedge_{m=1}^{l-1} \neg E^m \right) \wedge \mathcal{A}_{t_k} \right]. \end{aligned} \quad (5)$$

Intuitively, the expression $\Pr \left[\bigwedge_{l=1}^{\delta_k} (\neg E^l) \mid \mathcal{A}_{t_k} \right]$ denotes the probability that none of the possible instantiations of e_s^i (i.e. the edges that stem out of the v_s^i) gets the correct color.

Now observe that:

- (i) for an l 's for which $e_s^i \neq o^l$, it holds that $\Pr \left[\neg E^l \mid \left(\bigwedge_{m=1}^l \neg E^m \right) \wedge \mathcal{A}_{t_k} \right] = 1$, whereas
- (ii) for an l for which $e_s^i = o^l$, it holds that

$$\Pr \left[\neg E^l \mid \left(\bigwedge_{m=1}^{l-1} \neg E^m \right) \wedge \mathcal{A}_{t_k} \right] \geq 1 - \frac{1}{\gamma(\Delta-1)+1},$$

because when $e_s^i = o^l$, the event E^l is equivalent to $\text{CA}(s, c_k; o^l)$ and the conditional

$$\left(\bigwedge_{m=1}^{l-1} \neg E^m \right) \wedge \mathcal{A}_{t_k}$$

is equivalent to \mathcal{A}_{t_k} , which is anterior to $\text{CA}(s, c_k; o^l)$, because $e_s^i = o^l$.

Therefore for all l , we have that

$$\Pr \left[\neg E^l \mid \left(\bigwedge_{m=1}^{l-1} \neg E^m \right) \wedge \mathcal{A}_{t_k} \right] \geq 1 - \frac{1}{\gamma(\Delta-1)+1},$$

Therefore:

$$\Pr \left[\bigwedge_{l=1}^{\delta_k} (\neg E^l) \mid \mathcal{A}_{t_k} \right] \geq \left(1 - \frac{1}{\gamma(\Delta-1)+1} \right)^{d_k},$$

where $d_k = 1$ if k corresponds to the last (in the geometric order) edge e_s^i of C_s , and $d_k = \Delta - 1$ otherwise (d_k is *not* random).

Therefore, from (4) and (5) we get that:

$$\Pr [\mathcal{E}_C^{t_k} \mid \mathcal{A}_{t_k}] \leq 1 - \left(1 - \frac{1}{\gamma(\Delta-1)+1} \right)^{d_k}.$$

Observe that among the $k = 1, \dots, 2\hat{k}$ there are n values of k for which $d_k = 1$ (the k 's for which e_s^i is the closing edge of C_s), whereas for the remaining $\sum_{s=1}^n (2k_s - 3)$ values of k , we have $d_k = \Delta - 1$. So we can conclude from (3) that that $\Pr [\text{COLORVAL is successful on } \mathcal{S}]$ is bounded from above by:

$$\left(\frac{1}{\gamma(\Delta-1)+1} \right)^n \prod_{s=1}^n \left(1 - \left(1 - \frac{1}{\gamma(\Delta-1)+1} \right)^{\Delta-1} \right)^{2k_s-3}. \quad \square$$

Remark 4. The fact that no two homochromatic edges can be incident onto the same vertex is essential for the correctness of the above proof. For example in the case of vertex coloring, where we can have more than one homochromatic vertices neighboring with the same vertex, the proof does not go through. Indeed, even if we define the random variable \mathcal{C} to be the unique cycle-sequence such that for every cycle C_s of \mathcal{C} , every non-pivotal edge e_s^i is either the first edge stemming out v_s^i that has the same color as the pivotal edge of C_s of the same parity, if there is one, or is the first edge stemming out of v_s^i , if there is none, still Equation 3 can only take the form:

$$\begin{aligned} \Pr[\text{COLORVAL is successful on } \mathcal{S}] &\geq \Pr[\mathcal{C} \text{ becomes bichromatic}] \\ &= \prod_{t=1}^{2\hat{k}} \Pr[\mathcal{E}_{\mathcal{C}}^t \mid \mathcal{A}_t] = \prod_{k=1}^{2\hat{k}} \Pr[\mathcal{E}_{\mathcal{C}}^{t_k} \mid \mathcal{A}_{t_k}]. \end{aligned} \quad (6)$$

Corollary 2. If \mathcal{S} is as in Lemma 5 then

$$\Pr[\text{COLORVAL is successful on } \mathcal{S}] \leq \left(\frac{1}{\Delta-1}\right)^n \prod_{i=1}^n \left(\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k_i-3}\right). \quad (7)$$

Proof. Use the inequality

$$1 - x > e^{-\frac{x}{1-x}}, \forall x < 1, x \neq 0$$

(see e.g. [12, inequality (4.5.7)]) for $x = \frac{1}{\gamma(\Delta-1)+1}$, which implies, after elementary operations, that $\left(1 - \frac{1}{\gamma(\Delta-1)+1}\right)^{\Delta-1}$ is at least $e^{-\frac{1}{\gamma}}$. \square

We now turn to the estimation of the probability, call it \hat{P}_n , that COLOR lasts for at least n phases.

Theorem 1. The expression P_n defined as

$$P_n = \sum_{\substack{n_1 + \dots + n_m = n \\ n_1, \dots, n_m \geq 0}} Q_{n_1} \cdots Q_{n_m}, \quad (8)$$

where Q_n is defined by the recursion:

$$Q_n = \sum_{k \geq 3} \left(\left(\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k-3} \right) \cdot \sum_{\substack{n_1 + \dots + n_{2k} = n-1 \\ n_1, \dots, n_{2k} \geq 0}} Q_{n_1} \cdots Q_{n_{2k}} \right), \quad (9)$$

with $Q_0 = 1$, is an upper bound of \hat{P}_n .

Proof. Let f be an unlabeled ordered forest with m trees (recall, m is the number of edges of G) and with n internal nodes, each, considered in their pre-order, having an even degree $2k_s, s = 1, \dots, n$ ($k_s \geq 3$). Let p_f be the probability that the witness forest of an execution of COLOR on G has underlying unlabeled forest f . Obviously

$$\hat{P}_n = \sum_f p_f, \quad (10)$$

where the sum above ranges over unlabeled forests as described.

With each internal node of f with half-degree k_s , we associate weight

$$\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k_s-3},$$

and we associate weight 1 with the leaves of f . We define:

$$\|f\| = \prod_{i=1}^n \left(\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k_i-3} \right). \quad (11)$$

We now prove:

Claim. $p_f \leq \|f\|$.

Proof of Claim. Let $\mathcal{L} = (e_1^1, C_1), \dots, (e_n^1, C_n)$ be a sequence which if considered as label-sequence of f leads to an n -witness forest for COLOR. We will examine what the possibilities for \mathcal{L} are and thus bound the probability that at least one such label-sequence, together with f , comprises a witness of COLOR.

First observe that the lengths of the cycles C_s should coincide with the positive degrees $2k_s$ of f , therefore the lengths of the C_s of \mathcal{L} are uniquely determined. Then observe that the edge-label e_1^1 of \mathcal{L} , which is to be assigned to the first root of f that is not a leaf, is uniquely determined to be the k 'th edge of G , where $k-1$ is the number of leaf-roots that precede r (recall the definition of a witness forest, the way it is defined by COLOR and the fact that we introduced isolated roots for feasible forests in order to cover all edges of G). There are at most $(\Delta-1)$ choices for e_1^1 , the second pivotal edge of C_1 . This introduces a factor of $(\Delta-1)$ for the sought after upper bound of p_f . Once e_1^1, e_1^2 of \mathcal{L} are determined, the event that f , together with such an \mathcal{L} , gives a witness for COLOR necessitates that some C_1 with pivotal edges e_1^1, e_1^2 and length $2k_1$ is bichromatic. For each such possibility for C_1 , the second edge-label e_2^1 of \mathcal{L} is determined to be the first edge of C_1 that does not correspond to a leaf of f (recall again the definition of a witness forest, the way it is defined by COLOR and the fact that we introduced isolated leaves for feasible forests in order to cover all edges of each cycle-label of each internal node). Again there are at most $(\Delta-1)$ possibilities for e_2^1 , thus introducing a factor of $(\Delta-1)$ to the sought after upper bound. And again, it becomes necessary that some C_2 with pivotal edges e_2^1, e_2^2 and length $2k_2$ is bichromatic. We continue in this fashion, following the structure of f : when a leaf is reached, we go back to the last internal node (in a depth-first fashion), then the next edge of this internal node's cycle-label should be the next edge-label in \mathcal{L} ; when a tree of the forest f is exhausted, we go to the next root, then the next edge of G should be the next edge-label \mathcal{L} . Now from Proposition 1, which bounds the probability that at least one forest with labels $(e_s^1, C_s)_s$ is a witness for COLOR, and Corollary 2, which bounds the probability that COLORVAL is successful on input $(e_s^1, e_s^2, k_s)_s$, and taking into account the n factors $(\Delta-1)$, which correspond to the possible choices of e_s^2 , we conclude that p_f is bounded from above by the product of the weights:

$$\|f\| = \prod_{i=1}^n \left(\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k_i-3} \right).$$

This concludes the proof of the Claim and we now return to the proof of Theorem 1.

Consider *unlabeled* ordered forests with m ordered trees with n internal nodes in total, whose out-degrees are even and ≥ 6 ; assign weight w_{2k} to all internal nodes of out-degree $2k$; also assign weight 1 to all leaves and let the weight of such a tree be the product of the weight assigned to all its nodes. Then by distributivity of multiplication over addition it follows that the sum f_n of the weights of all such weighted unlabeled ordered forests is given by:

$$f_n = \sum_{\substack{n_1 + \dots + n_m = n \\ n_1, \dots, n_m \geq 0}} t_{n_1} \cdots t_{n_m}, \quad (12)$$

where t_n , the sum of the weights of weighted unlabeled ordered rooted *trees* with n internal nodes, is given by the recurrence:

$$t_n = \sum_{k \geq 3} w_{2k} \left(\sum_{\substack{n_1 + \dots + n_{2k} = n-1 \\ n_1, \dots, n_{2k} \geq 0}} t_{n_1} \cdots t_{n_{2k}} \right), \quad (13)$$

with $t_0 = 1$. By the Claim above we get that:

$$\sum_f p_f \leq \sum_f \|f\| = \sum_f \left[\prod_{i=1}^n \left(\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}} \right)^{2k_i-3} \right) \right]. \quad (14)$$

The required follows from (10) and (14) making use of the recurrence given by (12) and (13). \square

2.1. Asymptotic analysis of the recurrence

We will asymptotically analyze the coefficients of the OGF $Q(z)$ of Q_n . Towards this end, multiply both sides of equality (9) by z^n and sum for $n = 1, \dots, \infty$ to get

$$Q(z) - 1 = \sum_{k \geq 3} \left[\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}} \right)^{2k-3} z Q(z)^{2k} \right], \quad (15)$$

with $Q(0) = 1$. Setting $W(z) = Q(z) - 1$ we get

$$W(z) = \sum_{k \geq 3} \left[\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}} \right)^{2k-3} z (W(z) + 1)^{2k} \right], \quad (16)$$

with $W(0) = 0$. For notational convenience, set $W = W(z)$. Then from (16) we get:

$$W = z \frac{1}{\gamma} \cdot \frac{\left(1 - e^{-\frac{1}{\gamma}} \right)^3 (W + 1)^6}{1 - \left(1 - e^{-\frac{1}{\gamma}} \right)^2 (W + 1)^2}. \quad (17)$$

Set now

$$\phi(x) = \frac{1}{\gamma} \cdot \frac{\left(1 - e^{-\frac{1}{\gamma}} \right)^3 (x + 1)^6}{1 - \left(1 - e^{-\frac{1}{\gamma}} \right)^2 (x + 1)^2}, \quad (18)$$

to get from (17):

$$W = z\phi(W). \quad (19)$$

Let now $R = \frac{1}{\left(1 - e^{-\frac{1}{\gamma}}\right)} - 1$ be the radius of convergence of the series representing ϕ at 0. Let also τ be the (necessarily unique) solution in the interval $(0, R)$ of the characteristic equation (in τ):

$$\phi(\tau) - \tau\phi'(\tau) = 0. \quad (20)$$

Finally, let

$$\rho = \frac{\tau}{\phi(\tau)}. \quad (21)$$

By [13, Proposition IV.5] (it is trivial to check that the hypotheses in that Theorem are satisfied for $\gamma > 0$), we get $[z^n]Q \asymp (1/\rho)^n$, i.e. $\limsup ([z^n]Q)^{1/n} = 1/\rho$ (see [13, IV.3.2]).

Now by a simple search (through Maple, for the code see [14]) we found that for $\gamma = 1.73095$, the unique positive solution of (20) in the radius of convergence is $\tau = 0.1747094762$, and this value of τ gives $1/\rho = 0.9999789027 < 1$. Therefore by making use of (8), we get:

Theorem 2. *Assuming Δ , the maximum degree of the graph G , is constant, and given the availability of at least $3.74(\Delta - 1) + 1$ colors, there exists an integer N , which depends linearly on m , the number of edges of G , and a constant $\rho > 1$ such that if $n/\log n \geq N$ then the probability that COLOR executes at least n calls of RECOLOR is $< (1/\rho)^n$; therefore the graph has an acyclic edge coloring.*

Now if the graph has girth $2r - 1$ for $r \geq 4$, the previous arguments carry over with minimal changes. Namely, equation (9) becomes:

$$Q_n = \sum_{k \geq r} \left[\frac{1}{\gamma} \left(1 - e^{-\frac{1}{\gamma}}\right)^{2k-3} \cdot \sum_{\substack{n_1 + \dots + n_{2k} = n-1 \\ n_1, \dots, n_{2k} \geq 0}} Q_{n_1} \cdots Q_{n_{2k}} \right], \quad Q_0 = 1. \quad (22)$$

Also in (15) and (16), the starting point of the summation is changed from 3 to r . Moreover, equation (17) becomes:

$$W = z \frac{1}{\gamma} \cdot \frac{\left(1 - e^{-\frac{1}{\gamma}}\right)^{2r-3} (W+1)^{2r}}{1 - \left(1 - e^{-\frac{1}{\gamma}}\right)^2 (W+1)^2}, \quad (23)$$

and equation (18) becomes:

$$\phi(x) = \frac{1}{\gamma} \cdot \frac{\left(1 - e^{-\frac{1}{\gamma}}\right)^{2r-3} (x+1)^{2r}}{1 - \left(1 - e^{-\frac{1}{\gamma}}\right)^2 (x+1)^2}. \quad (24)$$

Working as before, we get numerical results depicted in Figure 4 with sample specific values explicitly given in Figure 1.




Figure 4: γ as a function of girth

3. Discussion

There are several conceivable possibilities for improvement of the $\lceil 3.74(\Delta - 1) \rceil + 1$ bound. For example, in COLOR we can recolor not all edges of C_s but all but two consecutive of them (because it is only those that determine bichromaticity). Also in Theorem 1, we can conceivably improve the bound provided by claiming that for each execution, there is only one k_s so that the cycle C_s of half-length k_s is bichromatic. Nevertheless we opted not to consider these possible improvements, as our aim was only to present a probabilistic analysis of a Moser-type algorithm with dependent variables.

Acknowledgements

We are grateful to Juanjo Ru   for showing to us how to deal with the asymptotics of the coefficients of inverse generating functions. We are grateful to Dieter Mitsche on one hand and to Dimitris Achlioptas and Fotis Iliopoulos, on the other, for pointing out errors in previous versions of this paper; the second and fourth authors are also indebted to the latter two for initiating them to this line of research. The contributions of the undergraduate students Rafail Kartsioukas, Georgios Kontogeorgiou, Stavros Messaris and Zoi Terzopoulou at various stages of this work were crucial. We sincerely thank them all. We are also much indebted to Gwena  l Joret for pointing out the incompleteness of the proof of [15, Lemma 5] and for the illuminating for us exchange of messages that followed. John Livieratos' remarks at the last phases of this work were very helpful.

References

References

- [1] V. G. Vizing, Critical graphs with a given chromatic class, *Diskret. Analiz* 5 (1) (1965) 9–17.
- [2] J. Fiam  ik, The acyclic chromatic class of a graph (in Russian), *Math. Slovaca* 28 (1978) 139–145.
- [3] N. Alon, B. Sudakov, A. Zaks, Acyclic edge colorings of graphs, *Journal of Graph Theory* 37 (3) (2001) 157–167.
- [4] L. Esperet, A. Parreau, Acyclic edge-coloring using entropy compression, *Eur. J. Comb.* 34 (6) (2013) 1019–1027.

- [5] R. A. Moser, G. Tardos, A constructive proof of the general Lovász Local Lemma, *Journal of the ACM (JACM)* 57 (2) (2010) 11.
- [6] T. Tao, Mosers entropy compression argument Available: <https://terrytao.wordpress.com/2009/08/05/mosers-entropy-compression-argument/>.
- [7] D. Gonçalves, M. Montassier, A. Pinlou, Entropy compression method applied to graph colorings, [arXiv:1406.4380](https://arxiv.org/abs/1406.4380) (2014).
- [8] I. Giotis, L. Kirousis, K. I. Psaromiligkos, D. M. Thilikos, On the algorithmic Lovász local lemma and acyclic edge coloring, in: *Proceedings of the twelfth workshop on analytic algorithmics and combinatorics*, Society for Industrial and Applied Mathematics, 2015, <http://epubs.siam.org/doi/pdf/10.1137/1.9781611973761.2>.
- [9] R. A. Moser, A constructive proof of the Lovász Local Lemma, in: *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, ACM, 2009, pp. 343–350.
- [10] J. Spencer, Robin Moser makes Lovász Local Lemma algorithmic! Available: <http://cs.nyu.edu/spencer/moserlovasz1.pdf>.
- [11] S. Selvin, A problem in probability (letters to the editor), *American Statistician* 29 (1) (1975) 67–67.
- [12] Nist digital library of mathematical functions, <http://dlmf.nist.gov/4.5.7>, Release 1.0.9 of 2014-08-29, online companion to [16].
- [13] P. Flajolet, R. Sedgewick, *Analytic Combinatorics*, 1st Edition, Cambridge University Press, New York, NY, USA, 2009.
- [14] I. Giotis, L. Kirousis, K. I. Psaromiligkos, D. M. Thilikos, Maple code, available: <http://users.uoa.gr/~lkirousis/lll/bichromatic.mw.zip>.
- [15] I. Giotis, L. Kirousis, K. I. Psaromiligkos, D. M. Thilikos, Acyclic edge coloring through the lovász local lemma, *Theoretical Computer Science* 665 (2017) 40–50.
- [16] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, C. W. Clark (Eds.), *NIST Handbook of Mathematical Functions*, Cambridge University Press, New York, NY, 2010, print companion to [12].