Improved Streaming Edge Coloring

Shiri Chechik
Tel Aviv University, Israel

Hongyi Chen
State Key Laboratory for Novel Software Technology, New Cornerstone Science Laboratory, Nanjing University, China

Tianyi Zhang
ETH Zürich, Switzerland

Abstract

Given a graph, an edge coloring assigns colors to edges so that no pairs of adjacent edges share the same color. We are interested in edge coloring algorithms under the W-streaming model. In this model, the algorithm does not have enough memory to hold the entire graph, so the edges of the input graph are read from a data stream one by one in an unknown order, and the algorithm needs to print a valid edge coloring in an output stream. The performance of the algorithm is measured by the amount of space and the number of different colors it uses.

This streaming edge coloring problem has been studied by several works in recent years. When the input graph contains n vertices and has maximum vertex degree Δ, it is known that in the W-streaming model, an $O(\Delta^2)$-edge coloring can be computed deterministically with $\tilde{O}(n)$ space [Ansari, Saneian, and Zarrabi-Zadeh, 2022], or an $O(\Delta^{1.5})$-edge coloring can be computed by a $\tilde{O}(n)$-space randomized algorithm [Behnezhad, Saneian, 2024] [Chechik, Mukhtar, Zhang, 2024].

In this paper, we achieve polynomial improvement over previous results. Specifically, we show how to improve the number of colors to $\tilde{O}(\Delta^{4/3+\epsilon})$ using space $\tilde{O}(n)$ deterministically, for any constant $\epsilon > 0$. This is the first deterministic result that bypasses the quadratic bound on the number of colors while using near-linear space.

2012 ACM Subject Classification Theory of computation → Streaming, sublinear and near linear time algorithms; Theory of computation → Graph algorithms analysis

Keywords and phrases edge coloring, streaming

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.48

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2504.16470` [18]

Funding *Shiri Chechik*: European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 803118 UncertainENV).
Tianyi Zhang: Starting grant "A New Paradigm for Flow and Cut Algorithms" (no. TMSGI2_218022) of the Swiss National Science Foundation.

1 Introduction

Let $G = (V, E)$ be an undirected graph on n vertices with maximum vertex degree Δ. An edge coloring of G is an assignment of colors to edges in E such that no pairs of adjacent edges share the same color, and the basic objective is to understand the smallest possible number of colors that are needed in any edge coloring, which is called the edge-chromatic number of G. It is clear that the total number of colors should be at least Δ, and a simple greedy algorithm can always find an edge coloring using $2\Delta - 1$ colors. By the celebrated Vizing's theorem [43] and Shannon's theorem [40], $(\Delta + 1)$-edge coloring and $\lfloor 3\Delta/2 \rfloor$-edge coloring always exist in simple and multi-graphs respectively, and these two upper bounds are tight in some hard cases.

The edge coloring problem has been studied widely from the algorithmic perspective. There have been efficient algorithms for finding good edge coloring in various computational models, including sequential [2, 30, 41, 3, 9, 12, 4], dynamic [6, 10, 26, 21, 11, 20], online [22, 13, 38, 36, 14, 15, 27], and distributed [37, 28, 29, 31, 5, 16, 8, 21, 24] models. In this paper, we are particularly interested in computing edge coloring under the streaming model where we assume the input graph does not fit in the memory of the algorithm and can only be accessed via one pass over a stream of all edges in the graph. Since the output of edge coloring is as large as the graph size, the algorithm cannot store it in its memory. To address this limitation, the streaming model is augmented with an output stream in which the algorithm can write its answers during execution, and this augmented streaming model is thus called the W-streaming model.

Edge coloring in the W-streaming model was first studied in [7], and improved by follow-up works [17, 1] which led to a deterministic edge coloring algorithm using $O(\Delta^2)$ colors and $O(n)$ space, or $O(\Delta^2/s)$ colors and $O(ns)$ space as a general trade-off. In [32], the authors improved the trade-off to $O(\Delta^2/s)$ colors and [1]$\tilde{O}(n\sqrt{s})$ memory, yet it does not break the quadratic bound in the most natural $\tilde{O}(n)$ memory regime; this algorithm is randomized but can be derandomized in exponential time. The quadratic upper bound of colors was subsequently bypassed in [39, 19] where it was shown that an $O(\Delta^{1.5})$-edge coloring can be computed by a randomized algorithm using $\tilde{O}(n)$ space. Furthermore, in [39] the authors obtained a general trade-off of $O(\Delta^{1.5}/s+\Delta)$ colors and $\tilde{O}(ns)$ space which is an improvement over [32].

1.1 Our Results

In this paper, we focus on the basic $\tilde{O}(n)$-memory setting and improve the recent $\Delta^{1.5}$ randomized upper bound to $\Delta^{4/3+\epsilon}$.

▶ **Theorem 1.** *Given a simple graph $G = (V, E)$ on n vertices with maximum vertex degree Δ, for any constant $\epsilon > 0$, there is a randomized W-streaming algorithm that outputs a proper edge coloring of G using $O\left((\log \Delta)^{O(1/\epsilon)} n\right)$ space and $O\left((\log \Delta)^{O(1/\epsilon)} \Delta^{4/3+\epsilon}\right)$ different colors; both upper bounds hold in expectation.*

Furthermore, we also show that our algorithm can be derandomized using bipartite expanders based on error correcting codes at the cost of slightly worse bounds, as stated below. The deterministic algorithm and proof can be found in the full version [18].

▶ **Theorem 2.** *Given a simple graph $G = (V, E)$ on n vertices with maximum vertex degree Δ, for any constant $\epsilon > 0$, there is a deterministic W-streaming algorithm that outputs a proper edge coloring of G using $O\left((\log \Delta)^{O(1/\epsilon)} \cdot (1/\epsilon)^{O(1/\epsilon^3)} \cdot \Delta^{4/3+\epsilon}\right)$ colors and $O\left(n \cdot (\log \Delta)^{O(1/\epsilon^4)}\right)$ space.*

1.2 Technical Overview

Previous Approaches

Using a deterministic general-to-bipartite reduction from [32], we can assume the input graph $G = (L \cup R, E)$ is bipartite. Also, it suffices to color only a constant fraction of all edges in G, because we can recurse on the rest $1 - \Omega(1)$ fraction of G which only incurs an extra factor of $O(\log \Delta)$ on the total number of different colors.

[1] $\tilde{O}(f)$ hides $\log f$ factors.

Let us begin by recapping the randomized $\tilde{O}(\Delta^{1.5})$-edge coloring from [39]. Since the algorithm has $\tilde{O}(n)$ bits of memory, we can assume that input graph is read from stream in batches of size $\Theta(n)$. If the subgraph formed by every batch has maximum degree at most $\Delta^{1/2}$, then we can allocate $O(\Delta^{1/2})$ new colors for each batch, using $O(\Delta^{1.5})$ colors in total. Thus, the main challenge arises when some batches contain subgraphs with maximum degree exceeding $\Delta^{1/2}$.

To simplify the problem, let us assume that in each batch of $O(n)$ edges, every vertex in R has degree $d > \Delta^{1/2}$. To assign colors to edges, organize a table of colors of size $\Delta \times (\Delta/d)$, represented as a matrix $C[i, j]$ with indices $1 \leq i \leq \Delta$ and $1 \leq j \leq \Delta/d$. Then, for every vertex $u \in L$, draw a random shift r_u uniformly at random from $[\Delta]$. During the algorithm, each vertex $u \in L$ keeps a counter c_u of its degree in the stream so far, and each vertex $v \in R$ keeps a counter b_v of the number of batches in which it has appeared so far. Then, to color a single batch, for edge (u, v), the algorithm tentatively assigns the color $C[r_u + c_u, b_v]$ (indices are under modulo Δ).

Clearly, edges incident on the same vertex $u \in L$ receive different colors, because the counter c_u is incremented for each edge (u, v); also edges incident on the same vertex $v \in R$ but arriving in different batches are different, because the values of counter b_v are different. Randomization ensures that edges incident on $v \in R$ within the same batch receive mostly different colors from the same column in $C[*, b_v]$. Consider all the d neighbors of v in a single batch $u_1, u_2, \ldots, u_d$. Since all the random shifts r_{u_i} are independent and all the counters c_{u_i} are deterministic (they only depend on the input stream), with high probability, most row indices $(r_{u_i} + c_{u_i}) \bmod \Delta$ will be different.

Bypassing the $\Delta^{1.5}$ Bound

To better understand the bottleneck in the approach of [39], consider the following case. If every batch forms a regular subgraph with uniform degree d, then we can reduce the size of the color table from $\Delta \times (\Delta/d)$ to $(\Delta/d) \times (\Delta/d)$, since each vertex $u \in L$ could only appear in Δ/d different batches. So the main difficult case is when the batches are subgraphs of **unbalanced degrees**. As an extreme example, consider when vertices in L have degree 1, while the vertices in R have degree d. For the rest, our main focus will be on this extreme case, and show how to obtain a $\Delta^{1+\epsilon}$-edge coloring which is almost optimal.

The flavor of our approach is similar to [19]. Divide all Δ batches into Δ/d phases, each phase consisting of d consecutive batches. Let D be a parameter which upper bounds the number of batches in which any vertex $v \in R$ could appear during a single phase. Then, the maximum degree of a single phase would be bounded by Dd, so in principle we should be able to color all edges within this phase using $O(Dd)$ different colors. To implement the coloring procedure, at the beginning of each phase, prepare D fresh palettes $C_1, C_2, \ldots, C_D$, each of size d, and assign each batch in this phase a palette C_i where i is chosen from $[D]$ uniformly at random. To keep track of the colors already used around each vertex, we maintain the following data structures.

- Each vertex $v \in R$ keeps a list $U_v \subseteq \{C_1, C_2, \ldots, C_D\}$ of used palettes.
- Each vertex $u \in L$ initializes a random shift $r_u \in [d]$ at the beginning of the algorithm.

When a batch of edges $F \subseteq E$ arrives, we will use the assigned palette C_i to color this batch. For every edge $(u, v) \in F$, if $C_i \notin U_v$, then tentatively assign the color $(r_u + c_i) \bmod d$ in C_i to edge (u, v), where c_i denotes the number of times that palette C_i has been assigned to previous batches. If $C_i \in U_v$, we mark the edge as uncolored. Since the palettes are assigned to batches randomly, in expectation, a constant fraction of edges will be successfully colored. In the case that multiple edges incident on v are assigned the same color, we retain

only one such edge and mark the remaining edges as uncolored. Since all the random shifts r_u's are uniformly random and independent of the randomness of counters c_i, most tentative colors around any vertex $v \in R$ in this batch should be different. Once the batch F is processed, add C_i to all lists $U_v, v \in R$ such that $\deg_F(v) = d$.

To reason about the space usage, we can argue that the total size of the lists $U_v, \forall v \in R$ does not exceed $O(n)$, because a palette C_i appears in the list U_v only when v receives d edges in a batch which uses palette C_i. Since the total number of edges in a phase is $O(dn)$, the total list size should be bounded by $O(n)$. In this way, we are able to store all the used palettes of vertices in R; this is not new to our approach, and a similar argument was also used in [19]. The new ingredient is the way we store the used colors around vertices in L. Here we have exploited the fact that vertices in L have low degrees in each batch, so their progress in every palette C_i is **synchronized**; that is, we only need to store a single counter c_i which is the number of times that C_i appears, and then the next tentative color of $u \in L$ would be $(r_u + c_i) \bmod d$.

The above scheme would use $O(Dd)$ different colors in a single phase, so $O(\Delta D)$ colors across all Δ/d phases. When $D \leq \Delta^{1/4}$, this bound would be much better than $\Delta^{1.5}$. So, what if $D > \Delta^{1/4}$? To deal with this case, we will apply a two-layer approach. Specifically, let us further group every D consecutive phases as a super-phase, so there are at most $\Delta/Dd < \Delta^{1/4}$ super-phases in total (recall that we were assuming $d > \Delta^{1/2}$). Within a super-phase, we will allocate Δ fresh colors which are divided into Δ/Dd different color packages $P_1, P_2, \ldots, P_{\Delta/Dd}$, where each color package P_i is further divided into D palettes of size d as $P_i = C_{i,1} \cup C_{i,2} \cup \cdots \cup C_{i,D}$. In this way, the total number of colors would be $\Delta^{5/4}$.

Then, like what we did before, for each phase in a super-phase, assign a color package P_i from $P_1, P_2, \ldots, P_{\Delta/Dd}$ uniformly at random. Within each phase, we will stick to its assigned color package P_i and reuse the algorithm within a phase we have described before. Since a color package is shared by multiple phases, each vertex $v \in R$ needs to store a list $U_{v,2}$ which stores all the packages P_i it has used so far, and in any phase where P_i is assigned but P_i is already contained in $U_{v,2}$, we would not color any edges incident on v in this particular phase. By repeating the same argument, we can argue that the total size of all the lists $U_{v,2}, \forall v \in R$ is bounded by $O(n)$ as well.

We can further generalize this two-layer approach to multiple layers and reduce the total number of colors to $\Delta^{1+\epsilon}$. However, this only works with the most unbalanced setting where vertices on L always have constant degrees in each batch of input. In general, when low-degree side has degree d', our algorithm needs $d' \cdot \Delta^{1+\epsilon}$ colors. If d' is large, then we would switch to the color table approach from [39]. Balancing the two cases, we end up with $\Delta^{4/3+\epsilon}$ colors overall.

Derandomization using Bipartite Expanders

Randomization is used both in the unbalanced case and in the regular case. To replace the choices of the random shifts r_u and random color package assignments, we will show one can apply unbalanced bipartite expanders [42, 34, 35] in a black box manner. However, for the random shifts used in the regular case where we apply the color table idea from [39], it would not be enough to use an arbitrary bipartite expander, because the counters c_u's could be different and possibly damage the expansion guarantee; for example, it is not clear to us how to apply the bipartite expander construction based on Parvaresh–Vardy Codes [34]. To fix this issue, it turns out that it would be most convenient to use the bipartite expander construction based on multiplicity codes [35].

2 Preliminaries

Basic Terminologies

For any integer k, let us conventionally define $[k] = \{1, 2, \ldots, k\}$. For any set S and integer k, let $k \odot S$ be the multi-set that contains exactly k copies of each element in S.

Let $G = (V, E)$ denote the simple input graph on n vertices and m edges with maximum degree Δ. For any subset of edges $F \subseteq E$ and any vertex $u \in V$, let $\deg_F(u)$ be the number of edges in F incident on u. Sometimes we will also refer to the subgraph (V, F) simply as F.

Problem Definition

In the edge coloring problem, we need to find an assignment of colors to edges such that adjacent edges have distinct colors, and the objective is to minimize the total number of different colors.

In the W-streaming model introduced by [25, 33], all edges of the graph G arrive one by one in the stream in an arbitrary order, and the algorithm makes one pass over the stream to perform its computation. For the task of edge coloring, since the algorithm has much less space than the total size of the graph, it cannot store all the edge colors in its memory. To output the answer, the algorithm is given a write stream in which it can print all colors.

Next, to set the stage in a convenient way, we will state several reductions for the problem.

General-to-Bipartite Reduction

Let us first simplify the problem by a deterministic reduction from edge coloring in general graphs to bipartite graphs.

▶ **Lemma 3** (Corollary 3.2 in [32]). *Given an algorithm $\mathcal{A}$ for streaming edge coloring on bipartite graphs using $f(\Delta)$ colors and $g(n, \Delta)$ space, there is an algorithm $\mathcal{B}$ using $O(f(\Delta))$ colors and $O\left(g(n, \Delta) \log n + n \log n \log \Delta\right)$ space in general graphs. Furthermore, this reduction is deterministic.*

Reduction to Fixed Degree Pairs

By the above reduction, we focus on edge coloring for bipartite graphs $G = (V, E)$, where $V = L \cup R$. Since the algorithm has space $\Omega(n)$, we can divide the input stream into at most $O(m/n) = O(\Delta)$ batches, each batch containing n edges. For any vertex $u \in V$ and any batch F, let $\deg_F(u)$ be the number of edges in F incident on u. For every pair of integers $0 \le l, r \le \log_2 \Delta$ and every batch F, let $F_{l,r} = \{(u, v) \in F \mid \deg_F(u) \in [2^l, 2^{l+1}), \deg_F(v) \in [2^r, 2^{r+1})\}$, and define $E_{l,r} = \bigcup_F F_{l,r}$ over all batches F in the stream and $m_{l,r} = |E_{l,r}|$.

In the main text, we will devise an algorithm to handle edges in $F_{l,r}$ for any fixed pair of (l, r). The final coloring is obtained by taking the union of the colors over all (l, r), which will blow up the number of colors and space by a factor of $O(\log^2 \Delta)$.

Adapting to an Unknown Δ

So far we have assumed that the maximum vertex degree Δ in G is known in advance. Our algorithm can be adapted to an unknown value Δ in a standard way. Specifically, we can maintain the value Δ_t which is the maximum degree of the subgraph containing the first t edges in the input stream. Whenever $\Delta_t \in (2^{k-1}, 2^k]$, we apply our algorithm with $\Delta = 2^k$ to color all the edges. When k increments at some point, we restart a new instance of the algorithm with a new choice $\Delta = 2^k$ and continue to color the edges with fresh colors. Overall, the total number of colors will not change asymptotically.

3 Randomized $\Delta^{4/3+\epsilon}$ Edge Coloring

This section is devoted to the proof of Theorem 1.

Reduction to Partial Coloring

To find an edge coloring of a graph, it was shown in [19] that it suffices to color a constant fraction of edges in E if we do not care about $O(\log \Delta)$ blow-ups in the number of colors. Roughly speaking, for the uncolored edges, we can view them as another instance of edge coloring, and solve it recursively. The following statement formalizes this reduction.

▶ **Lemma 4** (implicit in [19]). *Suppose there is a randomized streaming algorithm $\mathcal{A}$ with space $g(n, \Delta)$ space, such that for each edge in $e \in E$, it assigns a color from $[f(\Delta)]$ or marks it as $\perp$ (uncolored) and print this information in the output stream, with the guarantee that there are at least δm edges in expectation which receive actual colors, for some value $\delta > 0$. Then, there is a randomized edge coloring algorithm $\mathcal{B}$ which uses at most $O\left(\frac{\log \Delta}{\delta} f(\Delta)\right)$ colors and $O\left(\frac{\log \Delta}{\delta} g(n, \Delta)\right)$ space in expectation.*

Proof sketch. The idea is to recursively apply the streaming algorithm on all edges marked with $\perp$ (uncolored). In expectation, each time we reapply the algorithm, the number of uncolored edges decrease by a factor of $1 - \delta$. So the expected recursion depth would be $O(\frac{\log \Delta}{\delta})$ before all uncolored edges fit in memory. ◀

According to the reductions in the preliminaries, we will focus on partial edge coloring in bipartite graphs with fixed degree pairs. More specifically, our algorithm consists of the following two components.

▶ **Lemma 5.** *Fix a parameter $\epsilon > 0$. Given an graph $G = (V, E)$ on n vertices with maximum vertex degree Δ, for any constant $\epsilon > 0$, and fix an integer pair (l, r), there is a randomized W-streaming algorithm that outputs a partial coloring of edges $F_{l,r}$ such that least $\delta m_{l,r}$ edges receive colors in expectation; here $\delta = 2^{-O(1/\epsilon)}$ is also a constant. The algorithm uses $O\left((\log \Delta)^{O(1/\epsilon)} \cdot \Delta^{1+\epsilon} \cdot 2^l\right)$ colors and $O\left((\log \Delta)^{O(1/\epsilon)} \cdot n\right)$ space.*

▶ **Lemma 6.** *Given an graph $G = (V, E)$ on n vertices with maximum vertex degree Δ, and fix an integer pair (l, r), there is a randomized W-streaming algorithm that outputs a partial coloring of edges $F_{l,r}$ such that least $m_{l,r}/2$ edges receive colors in expectation. The algorithm uses $O\left(\Delta + \Delta^2/2^{l+r}\right)$ colors and $O(n)$ space.*

Proof of Theorem 1. Basically, Lemma 5 deals with the case when $\min\{2^l, 2^r\}$ is small, and Lemma 6 deals with the case when $\min\{2^l, 2^r\}$ is large. For any (l, r), if $\min\{2^l, 2^r\} \le \Delta^{1/3}$, then by Lemma 5, the number of colors is at most $O\left((\log \Delta)^{O(1/\epsilon)} \cdot \Delta^{4/3+\epsilon}\right)$, and otherwise by Lemma 6 the number of colors would be $O\left(\Delta^{4/3}\right)$. Either way, the total number of colors over all (l, r) is bounded by $O\left((\log \Delta)^{O(1/\epsilon)} \Delta^{4/3+\epsilon}\right)$. ◀

3.1 Proof of Lemma 5

3.1.1 Data Structures

Before presenting the details of the data structures we will use in the main algorithm, let us start with a brief technical overview. The data structures consist of three components.

- **Forest structures on batches.** This part organizes edge input batches into parameterized forests in a way similar to B-trees. The height of the forests will be $O(1/\epsilon)$, and the branch size of each level will be an integer power of 2 in the range $[\Delta^\epsilon, \Delta]$, and so the total number of different forests will be bounded by $(\log \Delta)^{O(1/\epsilon)}$.
- **Color allocation on forests.** Each forest will be associated with a separate color set of size roughly $\Delta^{1+\epsilon}$. On each level of a forest, we will randomly partition the color set of this forest into color subsets (packages) and assign them to the tree nodes on this level. We will also make sure that the color packages on tree nodes are nested; that is, the color package of a node is a subset of the color package of its parent node.
  Each vertex will choose colors for its incident edges according to its own frequencies of accumulating edges in the stream. For example, when a vertex is gathering a large number of incident edges in a short interval of batches, it would use color packages on tree nodes with large branch sizes.
- **Vertex-wise data structures.** To ensure that we never assign the same color to adjacent edges, each vertex needs to remember which colors it has already used around it. To efficiently store all the previously used colors, we will show that the used colors are actually concentrated in color packages, so each vertex only needs to store the tree nodes corresponding to those color packages, instead of storing every used colors individually.

Next, let us turn to the details of the data structures we have outlined above.

Forest Structures on Batches

Without loss of generality, assume $l \le r$, $2^r > \Delta^\epsilon$, and Δ^ϵ is an integer power of two; note that if $2^r \le \Delta^\epsilon$, then the maximum degree inside each batch $F_{l,r}$ is at most Δ^ϵ, so we can use a fresh palette of size $O(\Delta^\epsilon)$, so the total number of colors would be $O(\Delta^{1+\epsilon})$.

As we have done in the preliminaries, partition the entire input stream into at most $m/n \le \Delta$ batches of size n. We will create at most $(\log \Delta)^{O(1/\epsilon)}$ different forest structures, where each leaf represents a batch, and the internal tree nodes represent consecutive batches. Each forest is parameterized by an integer vector $\mathbf{f} = (f_1, f_2, \ldots, f_h)$ such that:

- f_i is an integer power of two;
- $2^{r+1} = f_0 \ge f_1 \ge f_2 \ge \cdots \ge f_h = \Delta^\epsilon$, and $2^{r+1} \cdot \prod_{i=1}^{h-1} f_i \le m/n$.

We can assume the total number of batches in the input stream is an integer multiple of $2^{r+1} \cdot \prod_{i=1}^{h-1} f_i \le m/n$ by padding empty batches. Given such a vector $\mathbf{f} = (f_1, f_2, \ldots, f_h)$, we will define a forest structure $\mathcal{T}_\mathbf{f}$ with $h+1$ levels by a bottom-up procedure; basically we will build a forest on all the batches with branching factors $2^{r+1}, f_1, \ldots, f_h$ bottom-up. More specifically, consider the following inductive procedure.

- All the batches will be leaf nodes on level 0. Partition the sequence of all batches into groups of consecutive $f_0 = 2^{r+1}$ batches. For each group, create a tree node at level-1 connecting to all leaf nodes in that group.
- Given any $1 \le i \le h-1$, assume we have defined all the tree nodes on levels $1 \le j \le i$. List all the tree nodes on level i following the same ordering of the batches, and partition this sequence of level-i nodes into groups of size f_i. For each group, create a tree node at level-$(i+1)$ that connects to all nodes in the group.

According to the definition, in general, tree levels up to k have the same topological structure for all frequency vectors which share the same first $k-1$ coordinates $f_1, f_2, \ldots, f_{k-1}$. For any node $N \in V(\mathcal{T}_\mathbf{f})$, let $\mathcal{T}_\mathbf{f}(N)$ be the subtree rooted at node N. By the above definition, for any $1 \le k \le h$, for any level-k node N, the set of all leaf nodes contained in the subtree $\mathcal{T}_\mathbf{f}(N)$ form a sub-interval of the batch sequence with length $2^{r+1} \cdot \prod_{i=1}^{k-1} f_i$.

Color Allocation on Forests

Next, we will allocate color packages at each tree node of each forest $\mathcal{T}_{\mathbf{f}}$. By construction, each forest structure $\mathcal{T}_{\mathbf{f}}$ is a tree of $h+1$ levels (from level-0 to level-h), and each tree is rooted at a level-h node. Go over each tree $T \subseteq \mathcal{T}_{\mathbf{f}}$ and we will allocate color packages to tree nodes in a top-down manner.

- **Basis.** At the root N of the tree T, allocate a color package $\mathcal{C}^N$ with $25 \cdot 2^{l+r+2} \cdot \prod_{i=1}^{h}(5f_i)$ new colors. Divide package $\mathcal{C}^N$ evenly into $5f_h = 5\Delta^\epsilon$ smaller packages (colors are ordered alphabetically in a package):

  $$\mathcal{C}^N = \mathcal{C}^N_1 \sqcup \mathcal{C}^N_2 \sqcup \cdots \sqcup \mathcal{C}^N_{5\Delta^\epsilon}$$

  Here, symbol $\sqcup$ means disjoint union. By construction of tree T, N has f_{h-1} different children $N_1, N_2, \ldots, N_{f_{h-1}}$. Let sequence $(i_1, i_2, \ldots, i_{5f_{h-1}})$ be a random permutation of $(f_{h-1}/\Delta^\epsilon) \odot [5\Delta^\epsilon]$. For each $1 \le j \le f_{h-1}$, define color package $\mathcal{C}^{N_j} = \mathcal{C}^N_{i_j}$.
- **Induction.** In general, suppose we have defined color packages for tree nodes on levels $k, k+1, \ldots, h$. Go over all tree nodes on level k. Inductively, assume $|\mathcal{C}^N| = 25 \cdot 2^{l+r+2} \cdot \prod_{i=1}^{k}(5f_i)$. Divide color package $\mathcal{C}^N$ into $5f_k$ smaller packages (colors are ordered alphabetically in a package):

  $$\mathcal{C}^N = \mathcal{C}^N_1 \sqcup \mathcal{C}^N_2 \sqcup \cdots \sqcup \mathcal{C}^N_{5f_k}$$

  Let $i_1, i_2, \ldots, i_{5f_{k-1}}$ be a random permutation of $(f_{k-1}/f_k) \odot [5f_k]$. For each such level-$k$ node N, by construction, it has f_{k-1} children $N_1, N_2, \ldots, N_{f_{k-1}}$. Then, for each $1 \le j \le f_{k-1}$, define color package $\mathcal{C}^{N_j} = \mathcal{C}^N_{i_j}$.

By the above construction, each leaf node is allocated a color package of size $25 \cdot 2^{l+r+2}$. We will call each such smallest color package a **palette**. Notice that, by definition, the same palette may appear at multiple leaf nodes (which represent input batches). For any leaf node N (or equivalently, a batch), let $\mathsf{cnt}(\mathcal{C}^N)$ count the total number of times that palette $\mathcal{C}^N$ is also allocated to previous leaf nodes (batches). Note that the counters $\mathsf{cnt}(*)$ do not depend on the input stream and can be computed in advance.

Vertex-Wise Data Structures

To carry out the streaming algorithm, we will also maintain some data structures for vertices in V. At the beginning of the streaming algorithm, for each vertex $u \in L$, draw a random shift $r_u \in [3 \cdot 2^{r+1}]$ uniformly at random; these random shifts r_u's will remain fixed throughout the entire execution of the algorithm.

The main part is to specify the data structures associated with each vertex $v \in R$. For each forest $\mathcal{T}_{\mathbf{f}}$, we will maintain a set of marked nodes $M_{v,\mathbf{f}} \subseteq V(\mathcal{T}_{\mathbf{f}})$ throughout the streaming algorithm.

▶ **Invariant 7.** *We will ensure the following properties regarding the marked nodes $M_{v,\mathbf{f}}$ throughout the execution of the streaming algorithm.*

(1) *No two nodes in $M_{v,\mathbf{f}}$ lie on the same root-to-leaf path in the forest $\mathcal{T}_{\mathbf{f}}$. Furthermore, suppose the current input batch corresponds to a leaf F, and let P be the unique tree path from F to the tree root. Then, any node $N \in M_{v,\mathbf{f}}$ is a child of some node on P.*

(2) *For any node $N \in \mathcal{T}_{\mathbf{f}}$ on level-k such that $M_{v,\mathbf{f}} \cap V(\mathcal{T}_{\mathbf{f}}(N)) \neq \emptyset$, let $F_1, F_2, \ldots, F_s \subseteq E$ be all the input batches which correspond to leaf nodes in subtree $\mathcal{T}_{\mathbf{f}}(N)$. Take the union of the batches $U = F_1 \cup F_2 \cup \cdots \cup F_s$. Then, we have $\deg_U(v) \ge 2^{r-k} \cdot \prod_{i=1}^{k} f_i$.*

(3) *For any previous input batch F' before F such that:*
- *F and F' are in the same tree component in $\mathcal{T}_{\mathbf{f}}$,*
- $\deg_{F'}(v) \in [2^r, 2^{r+1})$,
- *v used some colors in $\mathcal{C}^{F'}$ during the algorithm,*

we guarantee that F' must be contained in some subtree $\mathcal{T}_{\mathbf{f}}(N)$ for some $N \in M_{v,\mathbf{f}}$.

(4) *The choices of $M_{v,\mathbf{f}}$ is independent of the randomness of $\{r_u \mid u \in L\}$ and the randomness of color packages $\mathcal{C}^*$, and they only depend deterministically on the input stream.*

Let us explain the purpose of the above properties. Invariant 7(1)(2) ensures that the algorithm only uses $\tilde{O}(n)$ space in total. Invariant 7(3) allows us to rule out all colors used previously, preventing duplicate assignments to edges incident on the same vertex. Invariant 7(4) will be technically important for the analysis of randomization.

3.1.2 Algorithm Description

Next, let us turn to describe the coloring procedure. At the beginning, all the marked sets $M_{v,\mathbf{f}}$ are empty for any $v \in R$ and any frequency vector $\mathbf{f}$. Upon the arrival of a new input batch F, we will describe the procedures that update the marked sets and assign edge colors.

Preprocessing Marked Sets

Since the current input batch has changed due to the arrival of F, we may have violated Invariant 7(2) as the root-to-leaf tree path may have changed. Therefore, we first need to update all the marked sets with the following procedure named UpdateMarkSet(F).

Go over every vertex $v \in R$ and every frequency vector $\mathbf{f}$. Consider the position of F in the forest $\mathcal{T}_{\mathbf{f}}$, and let P be the root-to-leaf path in $\mathcal{T}_{\mathbf{f}}$ ending at leaf F. First, remove all marked nodes $N \in M_{v,\mathbf{f}}$ which are not in the same tree as P; note that this may happen when F is the first leaf in a new tree component of $\mathcal{T}_{\mathbf{f}}$.

Next, go over every node W lying on the tree path P, for any child node N of W, if (1) $V(\mathcal{T}_{\mathbf{f}}(N)) \not\ni F$ and (2) $V(\mathcal{T}_{\mathbf{f}}(N)) \cap M_{v,\mathbf{f}} \neq \emptyset$, then remove all nodes in $V(\mathcal{T}_{\mathbf{f}}(N)) \cap M_{v,\mathbf{f}}$ from $M_{v,\mathbf{f}}$ and add N to $M_{v,\mathbf{f}}$. In other words, we elevate the positions of all the marked nodes in $M_{v,\mathbf{f}}$ to their ancestors which are children of P. See Figure 1 for an illustration.

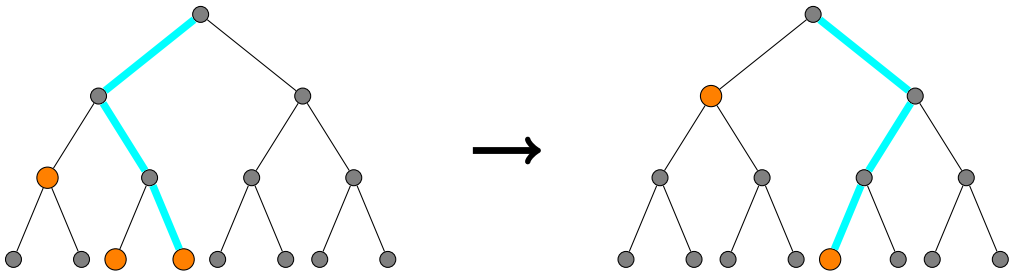

Figure 1 In this picture, it shows an example of a forest $\mathcal{T}_{\mathbf{f}}$ where the orange nodes are the marked ones, and the blue path is the root-to-leaf path ending at the current input batch F. Upon the arrival of a new input batch F, we need to update the root-to-leaf tree path and the marked sets accordingly.

Coloring $F_{l,r}$

To find colors for edges in $F_{l,r}$, we first need to find a proper palette for each vertex $v \in R$ such that $\deg_F(v) \in [2^r, 2^{r+1})$. We will describe an algorithm FREQVEC(F) which finds a proper frequency vector $\mathbf{f}_v = (f_1, f_2, \ldots, f_h)$ for each such $v \in R$ and use the palette $\mathcal{C}^F$ associated with leaf node F in the forest $\mathcal{T}_{\mathbf{f}_v}$. To identify the proper frequency vector $\mathbf{f}_v$, we will figure out each coordinate $f_1, f_2, \ldots, f_h$ one by one inductively.

- **Basis.** Let N be the unique level-1 node containing F (recall that level-1 nodes are defined irrespective of frequency vectors). Find $f_1 \in \{\Delta^\epsilon, 2\Delta^\epsilon, \ldots, 2^{r+1}\}$ which is the smallest integer such that there exists a frequency vector $\mathbf{f}' = (f_1, f_2', f_3', \ldots)$, so that $M_{v,\mathbf{f}'}$ contains less than f_1 children of N. Note that such a vector must exist, because $f_1 = 2^{r+1}$ always satisfies this requirement.

  If $f_1 = \Delta^\epsilon$, return the 1-dimensional vector $\mathbf{f}_v = (f_1)$.
- **Induction.** In general, assume we have specified a prefix $f_1, f_2, \ldots, f_k$ for some $k \geq 1$. Note that all the forests $\mathcal{T}_{\mathbf{f}'}$ share the same topological structures from level 0 up to $k+1$. Let N be the unique level-$(k+1)$ node containing F conditioning on $f_1, f_2, \ldots, f_k$. Check all the frequency vectors $\mathbf{f}'$ that begin with the prefix $f_1, f_2, \ldots, f_k$, and find the smallest possible $f_{k+1} \in \{\Delta^\epsilon, 2\Delta^\epsilon, \ldots, f_k\}$ such that there exists a frequency vector $\mathbf{f}' = (f_1, f_2, \ldots, f_k, f_{k+1}, f_{k+2}', \ldots)$ under the condition that $M_{v,\mathbf{f}'}$ contains less than f_{k+1} children of N. Note that such a vector must exist, because $f_{k+1} = f_k$ always satisfies this requirement.

  If $f_{k+1} = \Delta^\epsilon$, then return the vector $\mathbf{f}_v = (f_1, f_2, \ldots, f_k, f_{k+1})$.
- **Choosing palette $\mathcal{C}_v$.** Once we have finished the above inductive process, we need to choose a palette $\mathcal{C}_v$ for v which contains $25 \cdot 2^{l+r+2}$ different colors. Basically, we will choose the palette associated with leaf node F in forest $\mathcal{T}_{\mathbf{f}_v}$ as $\mathcal{C}_v$, but must ensure that $\mathcal{C}_v$ has not been used previously.

  To make sure of this, let P denote the root-to-leaf tree path ending at leaf node F in $\mathcal{T}_{\mathbf{f}_v}$. Travel down the path from root to leaf and enumerate all the encountered tree nodes. For every such tree node N, if N already contains a sibling $W \in M_{v,\mathbf{f}_v}$ and $\mathcal{C}^W = \mathcal{C}^N$, then abort this procedure and assign $\mathcal{C}_v \leftarrow \emptyset$.

  In the end, if this procedure terminates without abortion, then assign $\mathcal{C}_v \leftarrow \mathcal{C}^F$; here $\mathcal{C}^F$ refers to the palette of size $25 \cdot 2^{l+r+2}$ associated with leaf node F in forest $\mathcal{T}_{\mathbf{f}_v}$.

In this way, we can select a frequency vector $\mathbf{f}_v$ for every vertex $v \in R$ such that $\deg_F(v) \in [2^r, 2^{r+1})$. Let $\mathcal{C}_v$ be the palette assigned to leaf node F by forest $\mathcal{T}_{\mathbf{f}_v}$. Next, we are going to color all edges in $F_{l,r}$ using colors from $\mathcal{C}_v$ for edges incident on $v \in R$. Go over each vertex $u \in L$, and list its neighbors $v_1, v_2, \ldots, v_k, k < 2^{l+1}$ in $F_{l,r}$. For any index $1 \leq i \leq k$, if $\mathcal{C}_{v_i} \neq \emptyset$, then reserve a tentative color to edge (u, v_i), which is the κ_i-th color in palette $\mathcal{C}_{v_i}$ and

$$\kappa_i = 5 \cdot 2^{l+1} \cdot (\mathsf{cnt}(\mathcal{C}_{v_i}) + r_u) + i \mod 25 \cdot 2^{l+r+2}$$

Recall that $\mathsf{cnt}(\mathcal{C}_v)$ counts the number of times that palette $\mathcal{C}_v$ has appeared at previous leaf nodes under $\mathcal{T}_{\mathbf{f}}$. Notice that $k < 2^{l+1}$, these tentative colors are different around u.

On the side of $v \in R$, it checks all the tentative colors on all of its edges in $F_{l,r}$. If a tentative color is used more than once, then it keeps only one of them, and turns other tentative colors back to $\perp$. Finally, for each edge $e \in F_{l,r}$, assign e its own tentative color and print it in the output stream.

Postprocessing Marked Sets

After processing the input batch F, for every vertex $v \in R$ such that $\deg_F(v) \in [2^r, 2^{r+1})$, add node F to $M_{v,\mathbf{f}_v}$; note that we mark F irrespective of whether $\mathcal{C}_v$ is $\emptyset$ or not as will be important for establishing Invariant 7(4). The whole algorithm is summarized in Algorithm 1.

Algorithm 1 ColorLowDeg(F).

```
/* pre-processing marked sets */
for v ∈ R and frequency vector f do
    call UpdateMarkSet(F) (as described in Section 3.1.2) to remove marked
     nodes in previous tree components in T_f,
    and elevate the positions of all the marked nodes in M_{v,f} to their ancestors which
     are children of P;
/* select frequency vector f_v and palette C_v */
for v ∈ R such that deg_F(v) ∈ [2^r, 2^{r+1}) do
    call FreqVec(F) (as described in Section 3.1.2) to determine the frequency
     vector f_v = (f_1, f_2, ...) progressively;
    find palette C_v while ensuring no conflicts with sibling nodes;
/* assign colors for F_{l,r} */
for u ∈ L do
    let v_1, v_2, ..., v_k be all of u's neighbors in F_{l,r};
    assign edge (u, v_i) a tentative color which is the κ_i-th color in palette C_{v_i}, where
     κ_i = 5 · 2^{l+1} · (cnt(C_{v_i}) + r_u) + i  mod 25 · 2^{l+r+2};
for v ∈ R such that deg_F(v) ∈ [2^r, 2^{r+1}) do
    discard tentative colors that appear more than once around v in F_{l,r};
    output the unique tentative colors to the output stream;
/* post-processing marked sets */
for v ∈ R such that deg_F(v) ∈ [2^r, 2^{r+1}) do
    add F to M_{v,f_v};
```

3.1.3 Proof of Correctness

To begin with, let us first bound the total number of different colors that we could possibly use.

▶ **Lemma 8.** *The total number of colors that the algorithm could use is* $O((\log \Delta)^{O(1/\epsilon)} \Delta^{1+\epsilon} \cdot 2^l)$.

Proof. By the design of the forest structures, for each frequency vector $\mathbf{f}$, the number of colors assigned to each tree in the forest $\mathcal{T}_\mathbf{f}$ is bounded by $25 \cdot 2^{l+r+2} \cdot \prod_{i=1}^{h}(5f_i)$. Since each tree in $\mathcal{T}_\mathbf{f}$ covers $2^{r+1} \cdot \prod_{i=1}^{h-1} f_i$ batches, and there are at most $m/n \leq \Delta$ batches, the total number of colors in $\mathcal{T}_\mathbf{f}$ is bounded by

$$25 \cdot 2^{l+r+2} \cdot \prod_{i=1}^{h}(5f_i) \cdot \lceil \frac{\Delta}{2^{r+1} \cdot \prod_{i=1}^{h-1} f_i} \rceil = O(\Delta^{1+\epsilon} \cdot 2^l \cdot 5^{O(1/\epsilon)}).$$

Since there are $(\log \Delta)^{O(1/\epsilon)}$ different choices for $\mathbf{f}$, the total number of colors used by the algorithm is $O((\log \Delta)^{O(1/\epsilon)} \Delta^{1+\epsilon} \cdot 2^l)$. ◀

Let us next state some basic properties of any forest $\mathcal{T}_{\mathbf{f}}$.

▶ **Lemma 9.** *Each palette is used by at most $2^{r+1}/\Delta^{\epsilon}$ different batches.*

Proof. Consider a palette $\mathcal{C}^N$ allocated to a leaf node N. By the construction of color packages:

- At the root, exactly one package contains this palette, and $f_{h-1}/\Delta^{\epsilon}$ children of the root inherit the palette. Therefore, at level $h-1$, the palette $\mathcal{C}^N$ is used in at most $f_{h-1}/\Delta^{\epsilon}$ nodes.
- Suppose at level k, the palette $\mathcal{C}^N$ is used in f_k/Δ^{ϵ} nodes. Each of these nodes has f_{k-1}/f_k children that inherit the palette. Thus, at level $k-1$, the palette is used in at most $f_{k-1}/\Delta^{\epsilon}$ nodes.

By induction, the palette $\mathcal{C}^N$ is used in at most $2^{r+1}/\Delta^{\epsilon}$ different batches at level 0, as required. ◀

▶ **Corollary 10.** *During the algorithm, the values of the counters $\mathsf{cnt}(\mathcal{C})$ never exceed $2^{r+1}/\Delta^{\epsilon}$ for any palette $\mathcal{C}$.*

To bound the total space, it is clear that the bottleneck is storing all the marked sets $M_{v,\mathbf{f}}$, since all the forest structures and color assignments only require space $O\left((\log\Delta)^{O(1/\epsilon)}\Delta\right)$.

▶ **Lemma 11.** *If Invariant 7 is satisfied, then at any point during the execution of the algorithm, for any given frequency vector $\mathbf{f}$, the total size of marked sets $\sum_{v\in R}|M_{v,\mathbf{f}}|$ is bounded by $O(2^h n)$. Consequently, the total space usage is bounded by $O((\log\Delta)^{O(1/\epsilon)}n)$.*

Proof. For any level-k vertex $N \in \mathcal{T}_{\mathbf{f}}$, the subtree $\mathcal{T}_{\mathbf{f}}(N)$ contains exactly $2^{r+1}\cdot\prod_{i=1}^{k-1} f_i$ batches, and thus $2^{r+1}\cdot\prod_{i=1}^{k-1} f_i \cdot n$ edges. Let U denote the union of all batches in $\mathcal{T}_{\mathbf{f}}(N)$. For any vertex v, if $M_{v,\mathbf{f}} \ni N$, by Invariant 7(2), we have $\deg_U(v) \geq 2^{r-k}\cdot\prod_{i=1}^{k} f_i$. Thus, the number of vertices v such that $M_{v,\mathbf{f}} \ni N$ is bounded by:

$$\frac{2^{r+1}\cdot\prod_{i=1}^{k-1} f_i \cdot n}{2^{r-k}\cdot\prod_{i=1}^{k} f_i} = \frac{2^{k+1}\cdot n}{f_k}.$$

Suppose the current working batch is F. By Invariant 7(1), any marked vertex is a child of a node on the root-to-leaf path P. Thus, among all vertices at level k, there are f_k vertices that are candidates for being marked. Taking the summation over all levels, the total size of $\sum_{v\in R}|M_{v,\mathbf{f}}|$ is bounded by

$$\sum_{v\in R}|M_{v,\mathbf{f}}| \leq \sum_{k=0}^{h+1}\frac{2^{k+1}\cdot n}{f_k}\cdot f_k = O(2^h n).$$

Since the depth of the tree is at most $O(1/\epsilon)$, and there are $O((\log\Delta)^{O(1/\epsilon)})$ distinct frequency vectors $\mathbf{f}$, the total space usage is therefore bounded by $O((\log\Delta)^{O(1/\epsilon)}n)$. ◀

Now, our main focus will be on verifying all the properties in Invariant 7.

▶ **Lemma 12.** *Invariant 7 is preserved by the algorithm throughout its execution.*

Proof. Property (1) is preserved in a straightforward manner by the way we maintain all the marked sets $M_{v,\mathbf{f}}$ in the pre- and post-processing step for each input batch. Property (4) is also guaranteed because our updates to marked sets only depend on the input stream, not on the randomness used by the algorithm.

As for property (3), according to the coloring algorithm, for any previous input batch F' where $\deg_{F'}(v) \in [2^r, 2^{r+1})$, if vertex $v \in R$ used some colors associated with leaf node F' in forest $\mathcal{T}_{\mathbf{f}}$, then the algorithm must have added F' to $M_{v,\mathbf{f}}$ back then. Then, according to the preprocessing rules, F' would always be contained in the subtree of some marked as long as the current leaf node F belongs to the same connected tree component as F' in $\mathcal{T}_{\mathbf{f}}$.

Next, we will focus on property (2). We will prove this by an induction on time. As the basis, property (2) holds trivially because all the marked sets are empty. Consider the arrival of any new input batch F and any vertex $v \in R$ such that $\deg_F(v) \in [2^r, 2^{r+1})$.

First, we argue that the pre-processing step does not harm property (2). This is because the inequality in property (2) is required for all ancestors of any marked node (including itself), so if it held right before the arrival of F, then it should also hold after the pre-processing step as we are only raising the positions of marked nodes to their ancestors.

Next, we consider the coloring step and the post-processing step. According to the coloring procedure, since v will only use colors and mark nodes in forest $\mathcal{T}_{\mathbf{f}_v}$, we will only need to worry about property (2) in forest $\mathcal{T}_{\mathbf{f}_v}$. Before the post-processing step, let P be the root-to-leaf path ending at leaf node F in $\mathcal{T}_{\mathbf{f}_v}$, and let W be the highest (in terms of levels) node on P such that $V(\mathcal{T}_{\mathbf{f}_v}(W)) \cap M_{v,\mathbf{f}_v} = \emptyset$, namely the subtree $\mathcal{T}_{\mathbf{f}_v}(W)$ does not contain any marked nodes; this node W always exists because F itself is a possible candidate.

To prove property (2) after we add F to $M_{v,\mathbf{f}_v}$, we only need to verify the inequality for all nodes on P between W and F. List all these nodes as $F = N_0, N_1, N_2, \ldots, N_k = W$. We will prove the inequality of property (2) for each index $i = 0, 1, \ldots, k$.

- When $i = 0$, we already know that $\deg_F(v) \geq 2^r$, so the inequality holds.
- For any index $1 \leq i \leq k$, consider the procedure that chose the coordinate f_i. Because of the minimality of f_i, we know that for some alternate frequency vector $\mathbf{f}' = (f_1, f_2, \ldots, f_{i-1}, f_i/2, \ldots)$, there are at least $f_i/2$ marked children of N'_i (before the post-processing step), where $N'_i \in V(\mathcal{T}_{\mathbf{f}'})$ sits at the same topological position as N_i (or in other words, $\mathcal{T}_{\mathbf{f}_v}(N_i)$ and $\mathcal{T}_{\mathbf{f}'}(N'_i)$ are isomorphic). Let U be the union of the batches which are all leaf nodes of $\mathcal{T}_{\mathbf{f}'}(N'_i)$. Therefore, by our inductive assumption regarding property (2), we know that:

$$\deg_U(v) \geq (f_i/2) \cdot 2^{r-i+1} \cdot \prod_{j=1}^{i-1} f_j = 2^{r-i} \cdot \prod_{j=1}^{i} f_j$$

This verifies the inequality at node N_i. ◀

Next, let us turn to the color assignment part. First, we need to verify that the algorithm never assigns the same color twice around the neighborhood of a single vertex.

▶ **Lemma 13.** *In the output stream, the algorithm never prints the same color for two adjacent edges.*

Proof. First, let us rule out color conflicts for edges around the same vertex $u \in L$. This is rather straightforward, because according to the algorithm description, within each batch, we only assign tentative colors with distinct color indices around each vertex $u \in L$. For two different batches F, F', if we happen to use the same color palette $\mathcal{C}$ in F, F' around the same vertex u, then the counter values $\mathsf{cnt}(\mathcal{C})$ must be different in these two batches. According to Corollary 10 and that $r_u \in [3 \cdot 2^{r+1}]$, the value of $\mathsf{cnt}(\mathcal{C}_v + r_u)$ is always in the range $[3 \cdot 2^{r+1}, 4 \cdot 2^{r+1}]$. Together with the fact that $\deg_F(u), \deg_{F'}(u) < 2^{l+1}$, the algorithm must use distinct colors in F, F' from $\mathcal{C}$.

Next, let us rule out color conflicts for edges around the same vertex $v \in R$. For any color palette provided by $\mathcal{T}_{\mathbf{f}_v}$ which was used before in some previous batch F', according to Invariant 7(3), there must be a marked node N' whose subtree contains F'. According to our coloring procedure, $\mathcal{C}_v$ is nonempty only when $\mathcal{C}^N$ is disjoint from all the color packages of its marked siblings N', for any node N on the root-to-leaf path to F in $\mathcal{T}_{\mathbf{f}_v}$. Therefore, v could not reuse any palettes previously assigned. Furthermore, since we discard repeated colors within a single palette, the assigned colors must be distinct. ◀

Finally, let us prove that the algorithm successfully colors a good fraction of all edges in the input stream.

▶ **Lemma 14.** *The total number of colored edges in the output stream is at least δm in expectation, where $\delta = 2^{-O(1/\epsilon)}$.*

Proof. Fix any input batch F and any vertex $v \in R$ such that $\deg_F(v) \in [2^r, 2^{r+1})$, it suffices to lower bound the expected number of colored edges in $F_{l,r}$ incident on v.

First, we need to analyze the probability that the color palette $\mathcal{C}_v$ is nonempty, based on the randomness of the distribution of color packages on forest $\mathcal{T}_{\mathbf{f}_v}$. As before, let P denote the root-to-leaf path in $\mathcal{T}_{\mathbf{f}_v}$ ending at F, and let $F = N_0, N_1, \ldots, N_h$ be all the nodes on the tree path. According to the selection of the frequency vector $\mathbf{f}_v$, for any $0 \le k \le h-1$, N_k has less than f_{k+1} marked siblings. By design, the color packages of N_{k+1} are determined by the length-f_k prefix of a random permutation $(f_k/f_{k+1}) \odot [5f_{k+1}]$. Therefore, by the independence guarantee from Invariant 7(4), the probability that $\mathcal{C}^{N_k}$ does not conflict with the color packages of any other marked siblings is at least $(1 - 1/(5f_{k+1}))^{f_{k+1}} \ge 4/5$. By applying this bound over all levels, the probability that $\mathcal{C}_v$ is nonempty is at least $(4/5)^h \ge (4/5)^{1/\epsilon}$.

Next, conditioning on the event that $\mathcal{C}_v \neq \emptyset$, let us analyze the amount of edges in $F_{l,r}$ that are colored around v. Let $u_1, u_2, \ldots, u_k$ be the neighbors of v in graph $(V, F_{l,r})$, where $k < 2^{r+1}$. For any fixed $1 \le i \le k$, as r_{u_i} was chosen uniformly at random from $[3 \cdot 2^{r+1}]$, the probability that $r_{u_i} \neq r_{u_j}$ for all $j \neq i$ is at least $(1 - 1/(3 \cdot 2^{r+1}))^k \ge 2/3$. This ensures that the tentative colors between u_i and v survive in the output stream with probability at least $2/3$.

Overall, the expected number of colored edges in $F_{l,r}$ would be at least $(2/3) \cdot (4/5)^{1/\epsilon} \cdot |F_{l,r}|$, which concludes the proof. ◀

3.2 Proof of Lemma 6

Without loss of generality, assume Δ is a power of 2.

3.2.1 Data Structures

At the beginning, for each vertex $u \in L$, draw a random number $s_u \in [\Delta/2^l]$ uniformly at random. For each vertex $v \in R$, draw a random number $t_v \in [\Delta/2^r]$ uniformly at random.

As in the preliminary steps, the input stream is divided into batches of size n (except for the last one). For each $u \in L$, let $\mathsf{cnt}(u)$ count the number of previous batches F where $\deg_F(u) \in [2^l, 2^{l+1})$, and symmetrically let $\mathsf{cnt}(v)$ count the number of previous batches F where $\deg_F(v) \in [2^r, 2^{r+1})$ for $v \in R$.

Additionally, we will use a palette matrix $\mathcal{C}$ of size $\frac{\Delta}{2^l} \times \frac{\Delta}{2^r}$, where each entry in $\mathcal{C}[i,j]$ corresponds to a distinct palette of size Δ_0, with $\Delta_0 \triangleq \left\lceil 4 \cdot \left(\frac{2^{l+r+1}}{\Delta} + 1 \right) \right\rceil$. All the colors can be represented as integers in the range $\left[\frac{\Delta_0 \cdot \Delta^2}{2^{l+r}} \right]$ naturally.

3.2.2 Algorithm Description

Let us describe the coloring procedure upon the arrival of a new input batch F. For each vertex $u \in L$, propose a tentative row index $x_u = (s_u + \mathsf{cnt}(u)) \mod \frac{\Delta}{2^l}$. For each vertex $v \in R$, propose a tentative column index $y_v = (t_v + \mathsf{cnt}(v)) \mod \frac{\Delta}{2^r}$. For each edge $(u, v) \in F$, we will assign a color from the matrix $\mathcal{C}[x_u, y_v]$ in the following manner.

Let $E_{x,y}$ be the set of edges $(u, v) \in F_{l,r}$ where $(x_u, y_v) = (x, y)$, and let $G_{x,y}$ be the subgraph whose edge set is $E_{x,y}$. To color $G_{x,y}$ with only Δ_0 colors, we need to prune it so that its maximum degree does not exceed Δ_0, which is done in this way: for each edge $(u, v) \in E_{x,y}$, if $\max\{\deg_{E_{x,y}}(u), \deg_{E_{x,y}}(v)\} > \Delta_0$, mark it as $\perp$ (uncolored) and remove it from $G_{x,y}$. Finally, since $G_{x,y}$ is a bipartite graph, we can apply the Δ_0-edge coloring algorithm [23] to color $G_{x,y}$ using the palette $\mathcal{C}[x, y]$ which has size Δ_0.

Finally, for each vertex $u \in L$ such that $\deg_F(u) \in [2^l, 2^{l+1})$, increment the counter $\mathsf{cnt}(u)$ by 1; also, increment the counters for $v \in R$ in a symmetric way. The whole algorithm is summarized in Algorithm 2.

Algorithm 2 ColorHighDeg(F).

define $x_u \leftarrow s_u + \mathsf{cnt}(u) \mod \Delta/2^l, \forall u \in L$;
define $y_v \leftarrow t_v + \mathsf{cnt}(v) \mod \Delta/2^r, \forall v \in R$;
define $E_{x,y} = \{(u, v) \in F_{l,r} \mid (x_u, y_v) = (x, y)\}, \forall (x, y)$;
for *every pair* (x, y) *and edge* $(u, v) \in E_{x,y}$ **do**
    `/* prune` $G_{x,y}$ `to cap its maximum degree */`
    remove edge (u, v) from $E_{x,y}$ if $\max\{\deg_{E_{x,y}}(u), \deg_{E_{x,y}}(v)\} > \Delta_0$;
for *every pair* (x, y) **do**
    use palette $\mathcal{C}[x, y]$ to color $E_{x,y}$;
increment counters $\mathsf{cnt}(*)$;

3.2.3 Proof of Correctness

Proper Coloring

To prove that the colored edges form a proper coloring, we need to show that for any two distinct edges $e_1 = (u, v_1)$ and $e_2 = (u, v_2)$ sharing a common vertex u, their colors are different. Since the algorithm is symmetric for L and R, we can assume $u \in L$. There are several cases below.

- If e_1 and e_2 use different matrix entries in $\mathcal{C}$, their colors are already distinct.
- Suppose both edges use palette $\mathcal{C}[x, y]$. Then, there are two sub-cases below.
  - If e_1 and e_2 belong to different batches F_1 and F_2 with batch counters $\mathsf{cnt}^{(1)}(u)$ and $\mathsf{cnt}^{(2)}(u)$, we have $x \equiv s_u + \mathsf{cnt}^{(1)}(u) \equiv s_u + \mathsf{cnt}^{(2)}(u) \pmod{\Delta/2^l}$. Since $\deg_F(u) \in [2^l, 2^{l+1})$, cnt_u never exceeds $\Delta/2^l$. Thus, $\mathsf{cnt}^{(1)}(u) = \mathsf{cnt}^{(2)}(u)$, which leads to a contradiction.
  - If e_1 and e_2 belong to the same batch, they belong to the same subgraph $G_{x,y}$. The correctness of the offline coloring algorithm guarantees they get distinct colors.

Space Usage

For each vertex u, we maintain its batch counter cnt_u and random shift s_u for $u \in L$ (or t_v for $v \in R$), requiring $O(n)$ space in total. During the batch coloring process, we also store the indices x_u and y_v, which require additional $O(n)$ space. Furthermore, each subgraph

$G_{x,y}$ has at most n edges, so the offline coloring algorithm requires $O(n)$ space. These spaces are reused across different subgraph coloring processes and different batches. Therefore, the overall space complexity is $O(n)$.

Number of Colors

The total number of colors is given by

$$\frac{\Delta}{2^l} \cdot \frac{\Delta}{2^r} \cdot \Delta_0 = O\left(\Delta + \frac{\Delta^2}{2^{l+r}}\right).$$

Next, we show that at least half of the edges get colored in expectation.

▶ **Lemma 15.** *During the algorithm, at least a* $1/2$ *fraction of the edges are colored in expectation.*

Proof. For any edge $(u, v) \in F_{l,r}$ such that $u \in L, v \in R$, we estimate the probability that (u, v) is colored. Define $x = x_u, y = y_v$. It suffices to lower bound the probability that $\deg_{E_{x,y}}(u), \deg_{E_{x,y}}(v) \leq \Delta_0$.

Let $(u, v_1), (u, v_2), \ldots, (u, v_k) \in F_{l,r}, k < 2^{l+1} - 1$ be all edges incident on u other than (u, v). Since G is a simple graph, all the vertices $v_1, v_2, \ldots, v_k$ are distinct and are different from v. Since y_{v_i} is uniformly distributed in $[\Delta/2^r]$, the probability that $y_{v_i} = y$ is at most $2^r/\Delta$. Using Markov's inequality, we have:

$$\Pr[\deg_{E_{x,y}}(u) \geq \Delta_0] \leq \frac{(2^{l+1} - 2) \cdot 2^r/\Delta}{\Delta_0} \leq 1/4$$

Symmetrically, we can argue that $\Pr[\deg_{E_{x,y}}(v) \geq \Delta_0] \leq 1/4$. Hence, the probability that (u, v) remains in $E_{x,y}$ after the pruning procedure would be at least $1/2$. This concludes the proof. ◀

References

1. Mohammad Ansari, Mohammad Saneian, and Hamid Zarrabi-Zadeh. Simple streaming algorithms for edge coloring. In *30th Annual European Symposium on Algorithms (ESA 2022)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2022.
2. Eshrat Arjomandi. An efficient algorithm for colouring the edges of a graph with $\Delta + 1$ colours. *INFOR: Information Systems and Operational Research*, 20(2):82–101, 1982.
3. Sepehr Assadi. Faster Vizing and Near-Vizing Edge Coloring Algorithms. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025.
4. Sepehr Assadi, Soheil Behnezhad, Sayan Bhattacharya, Martín Costa, Shay Solomon, and Tianyi Zhang. Vizing's theorem in near-linear time. *arXiv preprint arXiv:2410.05240*, 2024. doi:10.48550/arXiv.2410.05240.
5. Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. Distributed edge coloring in time polylogarithmic in Δ. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, pages 15–25, 2022. doi:10.1145/3519270.3538440.
6. Leonid Barenboim and Tzalik Maimon. Fully-dynamic graph algorithms with sublinear time inspired by distributed computing. In *International Conference on Computational Science (ICCS)*, volume 108 of *Procedia Computer Science*, pages 89–98. Elsevier, 2017. doi:10.1016/J.PROCS.2017.05.098.
7. Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Marina Knittel, and Hamed Saleh. Streaming and massively parallel algorithms for edge coloring. In *27th Annual European Symposium on Algorithms (ESA 2019)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.
8. Anton Bernshteyn. A fast distributed algorithm for $(\Delta + 1)$-edge-coloring. *J. Comb. Theory, Ser. B*, 152:319–352, 2022.

9 Sayan Bhattacharya, Din Carmon, Martín Costa, Shay Solomon, and Tianyi Zhang. Faster $(\Delta + 1)$-Edge Coloring: Breaking the $m\sqrt{n}$ Time Barrier. In *65th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2024.

10 Sayan Bhattacharya, Deeparnab Chakrabarty, Monika Henzinger, and Danupon Nanongkai. Dynamic Algorithms for Graph Coloring. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1–20. SIAM, 2018. doi:10.1137/1.9781611975031.1.

11 Sayan Bhattacharya, Martín Costa, Nadav Panski, and Shay Solomon. Nibbling at Long Cycles: Dynamic (and Static) Edge Coloring in Optimal Time. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2024. doi:10.1137/1.9781611977912.122.

12 Sayan Bhattacharya, Martín Costa, Shay Solomon, and Tianyi Zhang. Even Faster $(\Delta + 1)$-Edge Coloring via Shorter Multi-Step Vizing Chains. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025. (to appear).

13 Sayan Bhattacharya, Fabrizio Grandoni, and David Wajc. Online edge coloring algorithms via the nibble method. In *Proceedings of theACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2830–2842. SIAM, 2021. doi:10.1137/1.9781611976465.168.

14 Joakim Blikstad, Ola Svensson, Radu Vintan, and David Wajc. Online edge coloring is (nearly) as easy as offline. In *Proceedings of the Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 2024. doi:10.1145/3618260.3649741.

15 Joakim Blikstad, Ola Svensson, Radu Vintan, and David Wajc. Deterministic Online Bipartite Edge Coloring. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025.

16 Yi-Jun Chang, Qizheng He, Wenzheng Li, Seth Pettie, and Jara Uitto. Distributed Edge Coloring and a Special Case of the Constructive Lovász Local Lemma. *ACM Trans. Algorithms*, 16(1):8:1–8:51, 2020. doi:10.1145/3365004.

17 Moses Charikar and Paul Liu. Improved algorithms for edge colouring in the w-streaming model. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 181–183. SIAM, 2021. doi:10.1137/1.9781611976496.20.

18 Shiri Chechik, Hongyi Chen, and Tianyi Zhang. Improved streaming edge coloring, 2025. arXiv:2504.16470.

19 Shiri Chechik, Doron Mukhtar, and Tianyi Zhang. Streaming Edge Coloring with Subquadratic Palette Size. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 40:1–40:12, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2024.40.

20 Aleksander B. G. Christiansen. Deterministic dynamic edge-colouring. *CoRR*, abs/2402.13139, 2024. doi:10.48550/arXiv.2402.13139.

21 Aleksander Bjørn Grodt Christiansen. The Power of Multi-step Vizing Chains. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1013–1026. ACM, 2023. doi:10.1145/3564246.3585105.

22 Ilan Reuven Cohen, Binghui Peng, and David Wajc. Tight bounds for online edge coloring. In *60th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1–25. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00010.

23 Richard Cole, Kirstin Ost, and Stefan Schirra. Edge-coloring bipartite multigraphs in o (e logd) time. *Combinatorica*, 21(1):5–12, 2001. doi:10.1007/S004930170002.

24 Peter Davies. Improved distributed algorithms for the lovász local lemma and edge coloring. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4273–4295. SIAM, 2023. doi:10.1137/1.9781611977554.CH163.

25 Camil Demetrescu, Irene Finocchi, and Andrea Ribichini. Trading off space for passes in graph streaming problems. *ACM Transactions on Algorithms (TALG)*, 6(1):1–17, 2009. doi:10.1145/1644015.1644021.

26 Ran Duan, Haoqing He, and Tianyi Zhang. Dynamic edge coloring with improved approximation. In *30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2019.

27 Aditi Dudeja, Rashmika Goswami, and Michael Saks. Randomized Greedy Online Edge Coloring Succeeds for Dense and Randomly-Ordered Graphs. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025.

28 Michael Elkin, Seth Pettie, and Hsin-Hao Su. $(2\Delta - 1)$-Edge-Coloring is Much Easier than Maximal Matching in the Distributed Setting. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 355–370. SIAM, 2014.

29 Manuela Fischer, Mohsen Ghaffari, and Fabian Kuhn. Deterministic distributed edge-coloring via hypergraph maximal matching. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 180–191. IEEE, 2017. doi:10.1109/FOCS.2017.25.

30 Harold N Gabow, Takao Nishizeki, Oded Kariv, Daneil Leven, and Osamu Terada. Algorithms for edge coloring. *Technical Rport*, 1985.

31 Mohsen Ghaffari, Fabian Kuhn, Yannic Maus, and Jara Uitto. Deterministic distributed edge-coloring with fewer colors. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 418–430, 2018. doi:10.1145/3188745.3188906.

32 Prantar Ghosh and Manuel Stoeckl. Low-memory algorithms for online edge coloring. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.

33 Christian Glazik, Jan Schiemann, and Anand Srivastav. Finding euler tours in one pass in the w-streaming model with o (n log (n)) ram. *arXiv preprint arXiv:1710.04091*, 2017. arXiv:1710.04091.

34 Venkatesan Guruswami, Christopher Umans, and Salil Vadhan. Unbalanced expanders and randomness extractors from parvaresh–vardy codes. *Journal of the ACM (JACM)*, 56(4):1–34, 2009. doi:10.1145/1538902.1538904.

35 Itay Kalev and Amnon Ta-Shma. Unbalanced expanders from multiplicity codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

36 Janardhan Kulkarni, Yang P. Liu, Ashwin Sah, Mehtaab Sawhney, and Jakub Tarnawski. Online edge coloring via tree recurrences and correlation decay. In *54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 104–116. ACM, 2022. doi:10.1145/3519935.3519986.

37 Alessandro Panconesi and Romeo Rizzi. Some simple distributed algorithms for sparse networks. *Distributed computing*, 14(2):97–100, 2001. doi:10.1007/PL00008932.

38 Amin Saberi and David Wajc. The greedy algorithm is not optimal for on-line edge coloring. In *48th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 198 of *LIPIcs*, pages 109:1–109:18, 2021. doi:10.4230/LIPICS.ICALP.2021.109.

39 Mohammad Saneian and Soheil Behnezhad. Streaming edge coloring with asymptotically optimal colors. In *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8-12, 2024, Tallinn, Estonia*, volume 297 of *LIPIcs*, pages 121:1–121:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ICALP.2024.121.

40 Claude E Shannon. A theorem on coloring the lines of a network. *Journal of Mathematics and Physics*, 28(1-4):148–152, 1949.

41 Corwin Sinnamon. Fast and simple edge-coloring algorithms. *arXiv preprint arXiv:1907.03201*, 2019.

42 Amnon Ta-Shma, Christopher Umans, and David Zuckerman. Loss-less condensers, unbalanced expanders, and extractors. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 143–152, 2001. doi:10.1145/380752.380790.

43 Vadim G Vizing. The chromatic class of a multigraph. *Cybernetics*, 1(3):32–41, 1965.