

$(\Delta + 1)$ Vertex Coloring in $O(n)$ Communication

Maxime Flin
Reykjavik University
Reykjavik, Iceland
maximef@ru.is

Parth Mittal
University of Waterloo
Waterloo, Canada
parth.mittal@uwaterloo.ca

ABSTRACT

We study the communication complexity of $(\Delta + 1)$ vertex coloring, where the edges of an n -vertex graph of maximum degree Δ are partitioned between two players. We provide a randomized protocol which uses $O(n)$ bits of communication and ends with both players knowing the coloring. Combining this with a folklore $\Omega(n)$ lower bound, this settles the randomized communication complexity of $(\Delta + 1)$ -coloring up to constant factors.

CCS CONCEPTS

• Theory of computation $\rightarrow$ Communication complexity.

KEYWORDS

Coloring, Communication Complexity, Sparsification

ACM Reference Format:

Maxime Flin and Parth Mittal. 2024. $(\Delta + 1)$ Vertex Coloring in $O(n)$ Communication. In *ACM Symposium on Principles of Distributed Computing (PODC '24)*, June 17–21, 2024, Nantes, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3662158.3662796>

1 INTRODUCTION

Graph coloring is a fundamental problem in computer science. Given an undirected graph, we are asked to assign each vertex a color such that no two adjacent vertices have the same color. Minimizing the number of colors used, even approximately [22, 26], is known to be NP-hard. On the other hand, for a graph with maximum degree Δ , there is a simple greedy algorithm that finds a $(\Delta + 1)$ -coloring in linear time. Due to the sequential nature of this algorithm, the problem has received a great deal of attention in the sublinear and distributed models of computation, with recent sublinear time [2], semi-streaming space [2], distributed [6, 12, 15, 24] massively parallel computation [11, 13], and dynamic [7, 8, 17] algorithms.

We look at this problem in the two-party communication model of [31]. The edges of the input graph G are partitioned between two players, Alice and Bob, who wish to compute some function (or relation) on G while minimizing the number of bits they send to each other. We refer readers to [23, 29] for an extensive overview of the field. Communication lower bounds have been used to obtain an astonishing breadth of impossibility results in distributed

computing [5, 30], streaming [10, 19], data-structures [27], and circuit complexity [21].

Several graph problems that have fast classical algorithms have been studied in the two-party communication model, often leading to creative algorithms that get very close to somewhat trivial lower bounds. For example, deciding whether a graph is connected has an easy $O(n \log n)$ communication protocol (either player just sends a spanning forest of their subgraph); finding the minimum cut has an $O(n \log n)$ communication protocol, and an $\Omega(n \log \log n)$ randomized communication lower bound [3]; finding a maximum matching in bipartite graphs has an $O(n \log^2 n)$ communication protocol [9]. All of these problems have an $\Omega(n \log n)$ deterministic communication lower bound [16], and an $\Omega(n)$ randomized communication lower bound via a reduction from set-disjointness. Closing the gap between this linear lower bound and nearly-linear upper bounds in the randomized case has been a longstanding open problem in communication complexity.

1.1 Our Contributions

In this paper, we settle the randomized communication complexity of $(\Delta + 1)$ -coloring. For this problem, the current state of the art is similar to the aforementioned graph problems. The best upper bound (to our knowledge) is $O(n \log^2 \Delta)$, alluded to by [1] (see Remark 4.5 for a description of this protocol). It would be natural to expect $\Theta(n \log \Delta)$ to be the “right answer” for this problem — after all this is the number of bits one would use to write down a $(\Delta + 1)$ -coloring. However, we show that if Alice has all the edges of the graph, there is a deterministic $O(n)$ -bit protocol with which she can communicate a $(\Delta + 1)$ -coloring to Bob. This is formalized as a non-deterministic communication upper bound in Theorem 2. The upshot is that we can hope to beat $n \log \Delta$ in the general case, and indeed, our main result is a randomized protocol that uses $O(n)$ bits of communication to find a $(\Delta + 1)$ -coloring:

THEOREM 1. *There exists a zero-error randomized protocol that given an n -vertex graph G and its maximum degree Δ , finds a $(\Delta + 1)$ -coloring of G using $O(n)$ bits of communication in expectation.*

Using Markov’s inequality, the protocol of Theorem 1 can be adapted to provide worst-case guarantees if we allow the algorithm to fail with constant probability. We also assume players have access to a source of shared randomness. This assumption can be removed by a classical result of Newman.

Corollary 1.1. *There exists a private-randomness protocol that given an n -vertex graph G and its maximum degree Δ , outputs a proper $(\Delta + 1)$ -coloring of G with probability $2/3$ and use $O(n)$ bits of communication in the worst-case.*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

PODC '24, June 17–21, 2024, Nantes, France

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0668-4/24/06

<https://doi.org/10.1145/3662158.3662796>

To complement this result, we show a (folklore) communication lower bound ruling out constant-error $o(n)$ communication protocols, via a reduction from the identity function (i.e. on an n -bit string x , the output is x). We also investigate better bounds on the probability of linear communication. We show that when Δ is small compared to n , the protocol runs in $O(n)$ communication with high probability. When Δ is large, we show a slightly weaker $O(n \log^* \Delta)$ bound on communication holds with high probability.

1.2 Technical Overview

The Non-Deterministic Upper Bound. We rely on a result of Csikvári [32] about the number of proper $(\Delta + 1)$ -colorings of an n -vertex graph which tells us that a uniformly random coloring is proper with probability e^{-n} . With a straightforward application of the probabilistic method, this gives us a set of $2^{O(n)}$ colorings which contains a proper coloring for *any* n -vertex graph, which immediately implies the non-deterministic communication upper bound in Theorem 2 (the prover can just point out a valid coloring).

The Randomized Upper Bound. The first insight is that if we run the greedy algorithm on a random (instead of arbitrary) order of the vertices, each vertex has (on average) $\Delta/2$ colors available (instead of just 1) when we try to color it. (This is also the core of the proof of Csikvári.) We can exploit this glut of available colors by sampling random colors from $[\Delta + 1]$, and paying 2 bits of communication per sample to decide if it is valid. One can show this yields an $O(n \log \Delta)$ communication protocol (see Section 4.1). However, this is not enough to go all the way; e.g. on a $(\Delta + 1)$ -clique, this algorithm uses $\Theta(\Delta \log \Delta)$ bits of communication.

On the other hand, there is a well known [1] deterministic algorithm that can find an available color for a single vertex in $O(\log^2 \Delta)$ communication, which we briefly sketch here (see Lemma 4.4 for more details). Because the edges of the graph are partitioned, when Alice and Bob try to color the vertex v , the number of colors blocked in $N_A(v)$ and $N_B(v)$ sum to less than $\Delta + 1$ (as opposed to just their union being less than $\Delta + 1$). Hence, at least one part of any partition of $[\Delta + 1]$ retains this “slack” of 1, which allows a binary search to find an available color in $O(\log^2 \Delta)$ communication.

To get to $O(n)$ communication, we combine the two ideas above. Consider some vertex v and let k be its number of available colors. If Alice had all the edges incident to v , she could describe a sampled color that works using $O(\log(\Delta/k))$ bits. To approach this ideal complexity, we implement a form of palette sparsification: Alice and Bob restrict themselves to a random color space of size roughly $(\Delta/k)^2$. When k is large (i.e., very often for a random permutation), this color space is much smaller than $\Delta + 1$, but it still retains the critical “slack” property required by the binary search protocol (see Lemma 4.6). This yields an $O(\log^2(\Delta/k))$ communication protocol for coloring a single vertex, which in turn is enough to get our main theorem (see Section 4.3).

2 PRELIMINARIES

Notation. Throughout this paper, we will work with input graphs $G = (V, E)$ on the vertex set $V = [n] := \{1, 2, \dots, n\}$, with maximum degree Δ . Let $N(v)$ denote the neighborhood of v in G , and $d_v = |N(v)|$ its degree. For any integer $q \geq 1$, a q -coloring of G is a vector

$C \in [q]^n$. We say the coloring is proper if for all edges $\{u, v\} \in E$ we have $C(u) \neq C(v)$.

Model. In the classic two-party model, Alice and Bob are given n and Δ . The edges of a graph G are adversarially partitioned between the two parties. We will use E_A and E_B to denote the edges given to Alice and Bob respectively, and $N_A(v)$ and $N_B(v)$ for the neighborhood of v in E_A and E_B . At the end of the protocol both players should agree on the color of each vertex. We assume both players have access to public randomness. This can be relaxed to needing only private randomness while adding only $O(\log n + \log(1/\delta))$ bits to the communication cost [28], where δ is the success probability.

In the non-deterministic version of the model, the randomness is replaced by a prover who has access to all the edges of the graph. The prover sends a single message to Alice and Bob, who do not communicate with each other, and must agree on a proper coloring. In particular, if the prover’s message encodes an improper coloring, at least one of them must reject it. The cost of the protocol is the length of the prover’s message.

Proposition 2.1. *Let $a_1 \leq \dots \leq a_m$ be a sequence of real numbers, with $m \geq 2$. Then for any $1 \leq j < m$,*

$$\frac{1}{j} \sum_{k=1}^j a_k \leq \frac{1}{m} \sum_{k=1}^m a_k \leq \frac{1}{m-j} \sum_{k=j+1}^m a_k.$$

(In words, adding bigger numbers to an average cannot decrease it, and adding smaller numbers cannot increase it.)

PROOF. Note that each term on the left side is at most a_j , and each term on the right side is at least a_{j+1} , which immediately gives us that the left side is smaller than the right. To complete the proof, we can rewrite the average of the whole sequence as a convex combination of the two:

$$\frac{1}{m} \sum_{k=1}^m a_k = \frac{j}{m} \cdot \left(\frac{1}{j} \sum_{k=1}^j a_k \right) + \frac{m-j}{m} \cdot \left(\frac{1}{m-j} \sum_{k=j+1}^m a_k \right). \quad \blacksquare$$

2.1 Concentration Inequalities

Throughout the paper we use a handful of classic concentration bounds.

Proposition 2.2 (Chernoff Bound [14]). *Let $X_1, X_2, \dots, X_n$ be independent random variables in $[0, 1]$ and $X = \sum_{i=1}^n X_i$. Then, for any $\varepsilon > 0$,*

$$\Pr[X < (1 - \varepsilon) \mathbb{E}[X]], \Pr[X > (1 + \varepsilon) \mathbb{E}[X]] \leq \exp\left(-\frac{\varepsilon^2 \mathbb{E}[X]}{2 + \varepsilon}\right).$$

Proposition 2.3 (Chernoff Bound for Geometric Variables [14, Problem 2.5]). *Let $X_1, \dots, X_n$ be independent geometric variables of parameter $1/2$ and $X = \sum_{i=1}^n X_i$. Then, for any $r \geq 3$,*

$$\Pr[X > (2 + r)n] \leq \exp\left(-\frac{rn}{4}\right).$$

Proposition 2.4 (Hoeffding Bound [18]). *Let $X_1, \dots, X_n$ be independent random variables and $a_1, \dots, a_n, b_1, \dots, b_n$ reals such that for each $i \in [n]$, with probability one $a_i \leq X_i \leq b_i$. Then, if*

$X = \sum_{i=1}^n X_i$, for any $t > 0$,

$$\Pr[X > \mathbb{E}[X] + t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

Proposition 2.5 (Bounded Differences [14, Corollary 5.2]). *Let $X_1, \dots, X_n$ be independent random variables and $f(x_1, \dots, x_n)$ a function such that whenever x and x' differ in just the i -th coordinate, then $|f(x) - f(x')| \leq d_i$. Then, for all $t > 0$,*

$$\Pr[X < \mathbb{E}[f(X)] - t], \Pr[X > \mathbb{E}[f(X)] + t] \leq \exp\left(-\frac{2t^2}{d}\right),$$

where $d = \sum_{i=1}^n d_i^2$.

3 NON-DETERMINISTIC UPPER BOUND

In this section, we will prove our non-deterministic communication upper bound.

THEOREM 2. *There exists a non-deterministic protocol that given an n -vertex graph G with maximum degree Δ , finds a $(\Delta + 1)$ -coloring of G using $O(n)$ bits of communication.*

We will show that, for each $n \geq 1$, there is a set of colorings $C_{n,\Delta}$ in $[\Delta + 1]^n$ of size $2^{O(n)}$ such that any n -vertex graph with maximum degree Δ has a proper coloring in $C_{n,\Delta}$. Note that this immediately gives an $O(n)$ -bit non-deterministic communication protocol, since the prover can just point out a coloring in $C_{n,\Delta}$ that works for the given graph, with Alice and Bob accepting if and only if there are no monochromatic edges in their respective subgraphs.

Our first observation is that any graph G has many $(\Delta + 1)$ colorings. More formally:

Lemma 3.1. *The number of $(\Delta + 1)$ -colorings of a graph G on n vertices with maximum degree Δ is $\geq (\frac{\Delta+1}{e})^n$.*

The proof is an easy modification of a result of Csikvári [32, Thm 8.3] which lower bounds the number of q -colorings of d -regular graphs for all $q \geq d + 1$.

PROOF. Let π be a random permutation of $[n]$. We color the vertices in the order of π . For each $v \in V$, let d_v^π be the number of neighbors of v coming before v in π . When we color v , there are at least $\Delta + 1 - d_v^\pi$ choices of colors. Hence, the total number of colorings we can obtain by coloring in this order is

$$\# \text{ colorings from } \pi \geq \prod_{v \in V} (\Delta + 1 - d_v^\pi).$$

After taking the logarithm, the righthand side becomes $\sum_v \log(\Delta + 1 - d_v^\pi)$. Now, observe that d_v^π is uniformly distributed in

$\{0, 1, \dots, d_v\}$. Hence by linearity of expectation,

$$\begin{aligned} \mathbb{E}_\pi[\log \# \text{ colorings for } \pi] &\geq \sum_{v \in V} \frac{1}{d_v + 1} \sum_{i=0}^{d_v} \log(\Delta + 1 - i) \\ &\geq \sum_{v \in V} \frac{1}{\Delta + 1} \sum_{i=0}^{\Delta} \log(\Delta + 1 - i) \quad (\text{by Proposition 2.1}) \\ &= \sum_{v \in V} \frac{1}{\Delta + 1} \log((\Delta + 1)!) \\ &\geq n \log\left(\frac{\Delta + 1}{e}\right). \quad (\text{by Stirling's approximation}) \end{aligned}$$

To get the lemma, we apply Jensen's inequality on $\log x$ to obtain that $\log \mathbb{E}_\pi[\# \text{ colorings for } \pi] \geq \mathbb{E}_\pi[\log \# \text{ colorings for } \pi]$. ■

Fix a graph G on n vertices, with maximum degree Δ . Lemma 3.1 implies that a uniformly random coloring from $[\Delta + 1]^n$ is proper for G with probability $1/e^n$. This means that if we construct $C_{n,\Delta}$ by sampling $t = e^n \cdot n^2$ colorings at random in $[\Delta + 1]^n$, the probability that *none* of them is proper for G is at most:

$$(1 - 1/e^n)^t \leq \exp(-t/e^n) = \exp(-n^2).$$

Since there are at most $2^{\binom{n}{2}}$ graphs with n vertices (and maximum degree Δ), a union bound over all such graphs yields the probability that $C_{n,\Delta}$ does not contain a proper coloring for one of them is at most $2^{\binom{n}{2}} \cdot e^{-n^2} < 1$. Hence, there exists a set $C_{n,\Delta}$ of $e^n \cdot n^2 \leq 2^{2n}$ colorings in $[\Delta + 1]^n$ such that any n -vertex graph with maximum degree Δ has a proper coloring in $C_{n,\Delta}$.

is

4 TWO RANDOMIZED UPPER BOUNDS

In this section, we will prove our main theorem, which we restate below.

THEOREM 1. *There exists a zero-error randomized protocol that given an n -vertex graph G and its maximum degree Δ , finds a $(\Delta + 1)$ -coloring of G using $O(n)$ bits of communication in expectation.*

4.1 The First Attempt

We will start by presenting a simple $O(n \log \Delta)$ communication protocol for the problem, and then show that making one of its subroutines (even slightly) more efficient results in an $O(n)$ communication protocol.

Algorithm 1. A $O(n \log \Delta)$ communication protocol for $(\Delta + 1)$ -coloring.

- (1) Alice and Bob choose a random permutation π of $[n]$ with public randomness.
- (2) They iterate over the vertices ordered by π , and on vertex v :
 - (a) With public randomness, they sample a color c uniformly at random from $[\Delta + 1]$, and Alice (respectively Bob) sends a bit indicating whether c

is already assigned to a vertex in $N_A(v)$ (respectively $N_B(v)$).

- (b) If they both send 0 (i.e. c is not used in $N_A(v)$ or $N_B(v)$) they assign c to v , and move on to the next vertex. Otherwise, they go to Step 2a.

Lemma 4.1. *Algorithm 1 uses $O(n \log \Delta)$ bits of communication in expectation.*

PROOF. Before we analyze the algorithm, we need to set up some notation. For a vertex v and a permutation π , let d_v^π denote the number of neighbors of v in $G = (V, E_A \cup E_B)$ that appear before v in π , and $k_v^\pi := \Delta + 1 - d_v^\pi$ (think of k_v^π as a lower bound on the number of colors that will be available when we try to color v). First, we will use the randomness of Step 2a: Observe that on each sample in Step 2a, the probability of sampling a valid color for v is at least $k_v^\pi / (\Delta + 1)$. The number of trials required to find a valid color is a geometric random variable, and hence its mean is upper bounded by $(\Delta + 1) / k_v^\pi$. Since Alice and Bob send 1 bit each per trial, the communication cost of coloring v is $\leq 2(\Delta + 1) / k_v^\pi$ in expectation.

Next, we will use the randomness of π : Note that k_v^π is uniformly random from the set $\{\Delta + 1 - d_v, \Delta + 2 - d_v, \dots, \Delta + 1\}$. Because the analysis of Step 2a holds for any permutation π , the expected cost of coloring v is upper bounded by:

$$\frac{1}{d_v + 1} \sum_{i=\Delta+1-d_v}^{\Delta+1} \frac{2(\Delta+1)}{i} \leq \frac{1}{\Delta+1} \sum_{i=1}^{\Delta+1} \frac{2(\Delta+1)}{i} = 2H_{\Delta+1} = O(\log \Delta),$$

where the first inequality is by Proposition 2.1, and H_m is the m -th harmonic number. Summing up over all vertices, by linearity of expectation, the expected communication cost of the algorithm is $O(n \log \Delta)$. ■

Remark 4.2. *Note that this algorithm also works if E_A and E_B are not disjoint, unlike the algorithm in our main result.*

4.2 A Haystack With Many Needles

The analysis of Algorithm 1 is tight when G is a collection of disjoint $(\Delta + 1)$ -cliques. This is because we are spending 2 bits of communication per sampled color, and to even sample all $\Delta + 1$ colors, we need $\Omega(\Delta \log \Delta)$ samples per clique by Coupon Collector. Now suppose we have an “ideal” vertex v , where Alice has all the edge in G incident on v . Then Alice could just point out the first sampled color that works for v in $O(\log(\Delta/k_v^\pi))$ bits of communication, instead of the $O(\Delta/k_v^\pi)$ we spent. To get (close to) this ideal, we abstract coloring a single vertex into the following set-intersection type problem:

Problem 1 (k -Slack-Int). *Alice and Bob are given sets X and $Y \subseteq [m]$ respectively, such that $|X| + |Y| \leq m - k$ for some $k \geq 1$. Both of them are also given the integer m , but neither of them knows k . They wish to find an element in the intersection of $\bar{X} := [m] \setminus X$ and $\bar{Y} := [m] \setminus Y$.*

To see why it is relevant, we recast coloring a single vertex as an instance of Problem 1. As before, fix some vertex v and permutation π in Algorithm 1. Let A and B denote the sets of colors used so far in $N_A(v)$ and $N_B(v)$ respectively; then coloring v is equivalent to solving Problem 1 with $m = \Delta + 1$, $X = A$, $Y = B$ and $k = k_v^\pi$. We emphasize that the promise $|X| + |Y| \leq m - k$ holds because sets $N_A(v)$ and $N_B(v)$ are disjoint (the edges of G are *partitioned* between Alice and Bob). The main lemma of this section gives an efficient algorithm for this problem.

Lemma 4.3. *There exists a randomized protocol that solves k -Slack-Int in $O(\log^2(m/k))$ bits of communication in expectation.*

Note that the promise on $|X| + |Y|$ is crucial. For example, if we only knew $|X \cup Y| < m$, then this is just the set-intersection problem, which is known to require $\Theta(m)$ bits of communication [20].

In this language, the coloring step in Algorithm 1 was taking random samples in $[m]$, succeeding in each step with probability at least k/m , for an expected communication cost of $2m/k$. To get a better algorithm, we will first focus on the hardest case for this random sampling algorithm (i.e. $k = 1$). Using the promise of k -Slack-Int, we can binary search for an element.

Lemma 4.4. *There exists a deterministic protocol that solves k -Slack-Int in $O(\log^2(m))$ bits of communication, for any $k \geq 1$.*

PROOF. The key observation is that we can binary search for the target element. Let $L = \lceil m/2 \rceil$ and $R = [m] \setminus L$, and consider $|X \cap L| + |Y \cap L|$ and $|X \cap R| + |Y \cap R|$. The two numbers add up to $|X| + |Y|$, which is smaller than m . Hence $|X \cap L| + |Y \cap L|$ must be smaller than $m/2$, or $|X \cap R| + |Y \cap R|$ smaller than $m - m/2$, giving a smaller instance of the same problem. Alice and Bob can send $|X \cap L|$ and $|Y \cap L|$ respectively, in $\leq 2 \log m$ bits, and then recurse on the correct half. This gives a protocol for finding an element in $\bar{X} \cap \bar{Y}$ in $O(\log^2 m)$ communication. ■

Remark 4.5. *Note that Lemma 4.4 gives an $O(n \log^2 \Delta)$ deterministic communication protocol for $(\Delta + 1)$ -coloring, since coloring each vertex (ordered by an arbitrary permutation) is a k -Slack-Int instance (but with no guarantee on k other than $k \geq 1$).*

The upshot of the binary search protocol is that the condition $|X| + |Y| < m$ can be preserved while recursing on smaller subproblems. However, this costs too much communication, and we are not exploiting the fact that we have a lot of slack ($|X| + |Y| \leq m - k$, not just $m - 1$).

We know that if we sample m/k elements from $[m]$, we expect to see at least one element from $\bar{X} \cap \bar{Y}$. And indeed, if we let S denote the set of sampled elements, we are trying to solve Problem 1 on $S \cap X$ and $S \cap Y$, with the caveat that $|S \cap X| + |S \cap Y|$ may not be $\leq |S| - k$ (or for that matter even $< |S|$). It turns out that with roughly m^2/k^2 samples in S , we get $|S \cap X| + |S \cap Y| < |S|$ with $\Omega(1)$ probability. This means that we can run the binary search protocol of Lemma 4.4 on this smaller instance instead, and pay only $O(\log^2(m^2/k^2)) = O(\log^2(m/k))$ bits of communication.

A downside of the sketch above is that we need to know k (to compute p). This turns out to not be an issue, since any guess $\tilde{k} \leq k$ suffices (as we will see in the following lemma), and we can arrive at such a $\tilde{k}$ in $O(\log(m/k))$ guesses starting from m (as we will see after the lemma).

Lemma 4.6. *Let $X, Y \subseteq [m]$ be a k -Slack-Int instance, and let $p = \min\{150m/\tilde{k}^2, 1\}$, for some $\tilde{k} \leq k \leq m$. Define the random set S by sampling each element of $[m]$ independently with probability p . Then*

$$\Pr_S[|S \cap X| + |S \cap Y| \geq |S|] \leq 1/2.$$

PROOF. For brevity, let $s := |S|$, $s_x := |S \cap X|$ and $s_y := |S \cap Y|$. By linearity of expectation $\mathbb{E}[s]$ is $m \cdot p$, and $\mathbb{E}[s_x + s_y] \leq (m - k) \cdot p$. This means that $s_x + s_y < s$ as long as:

- $s_x \leq \mathbb{E}[s_x] + k \cdot p/5$,
- $s_y \leq \mathbb{E}[s_y] + k \cdot p/5$,
- $s \geq \mathbb{E}[s] - k \cdot p/5$.

Using a Chernoff bound from Proposition 2.2 on s_x when $\varepsilon \leq 1$, we have:

$$\Pr[s_x \geq (1 + \varepsilon) \cdot \mathbb{E}[s_x]] \leq \exp\left(-\frac{\varepsilon^2 \mathbb{E}[s_x]}{2 + \varepsilon}\right) \leq \exp\left(-\frac{\varepsilon^2 \mathbb{E}[s_x]}{3}\right)$$

Plugging in $\varepsilon = k \cdot p/(5 \mathbb{E}[s_x])$,

$$\begin{aligned} \Pr[s_x \geq \mathbb{E}[s_x] + k \cdot p/5] &\leq \exp\left(-\frac{(k \cdot p)^2 \mathbb{E}[s_x]}{75 \mathbb{E}[s_x]^2}\right) \\ &\leq \exp\left(-\frac{\tilde{k}^2 p^2}{75 \mathbb{E}[s_x]}\right) \quad (k \geq \tilde{k}) \end{aligned}$$

We can afford to be lazy and upper bound $\mathbb{E}[s_x]$ by $m \cdot p$, and get:

$$\begin{aligned} &\leq \exp\left(-\frac{\tilde{k}^2 \cdot p^2}{75m \cdot p}\right) = \exp\left(-\frac{\tilde{k}^2 \cdot p}{75m}\right) \\ &= \exp\left(-\frac{150\tilde{k}^2 \cdot m}{75m \cdot \tilde{k}^2}\right) = e^{-2}. \end{aligned}$$

If $\varepsilon = k \cdot p/(5 \mathbb{E}[s_x]) > 1$, Proposition 2.2 gives

$$\begin{aligned} \Pr[s_x \geq (1 + \varepsilon) \cdot \mathbb{E}[s_x]] &\leq \exp\left(-\frac{\varepsilon \mathbb{E}[s_x]}{3}\right) \leq \exp\left(-\frac{\tilde{k}p}{75}\right) \\ &\leq \exp\left(-\frac{150m}{75\tilde{k}}\right) \leq e^{-2} \quad (\tilde{k} \leq k \leq m) \end{aligned}$$

We can repeat the same argument for s_y and s to obtain the lemma. ■

Remark 4.7. *The choice of the sampling probability p in Lemma 4.6 may seem strange at first, and we attempt to explain it here. Since we are paying $O(\log^2(|S|))$ bits of communication for the binary search, we want the size of S to be $\text{poly}(m/k)$. This restricts us to $p \in \{1/k, m/k^2, \dots\}$. The most natural choice in this sequence, $p = 1/k$, makes it so that the gap between $\mathbb{E}[s_x + s_y]$ and $\mathbb{E}[s]$ is only a constant, whereas the expectations themselves are $\Omega(m/k)$, which means we cannot hope for $s_x + s_y < s$ with significant probability. The second most natural choice makes the expectations $O(m^2/k^2)$ while making the gap $\Omega(m/k)$, which is exactly the area where a binomial random variable is concentrated.*

Using the lemma, we have the following algorithm for k -Slack-Int:

Algorithm 2. A $O(\log^2(m/k))$ communication protocol for k -Slack-Int.

Input: Alice gets a set $X \subseteq [m]$, Bob gets a set $Y \subseteq [m]$ such that $|X| + |Y| \leq m - k$.

Output: Any element from $\bar{X} \cap \bar{Y}$.

- (1) For $\tilde{k} = m, m/2, \dots, 1$ (a sequence of exponentially decreasing guesses for k):
 - (a) Alice and Bob choose S by sampling each element of $[m]$ independently with probability $p = \min\{1, 150m/\tilde{k}^2\}$, using public randomness.
 - (b) They test if $|S \cap X| + |S \cap Y| < |S|$. If not, they continue to the next value of $\tilde{k}$, otherwise:
 - (c) They run the binary search protocol to find an element in $S \setminus (X \cup Y)$, and return it as the answer.

Note that the algorithm always terminates with a correct answer, since for $\tilde{k}$ small enough, $p = 1$, and we just run the binary search on all of $[m]$. To finish the proof of Lemma 4.3, we bound the communication cost of Algorithm 2:

PROOF OF LEMMA 4.3. Observe that we spend $O(\log(m/k))$ iterations before $\tilde{k} \leq k$ and we can apply Lemma 4.6. On each of these iterations, the test in Step 1b costs only $O(\log(m/k))$ bits of communication, since $\tilde{k} > k$ for all of them. If we get lucky, and the test passes, the binary search takes $O(\log^2(m/k))$ bits, and we are done.

Otherwise, once $\tilde{k} \leq k$, the cost of each iteration increases beyond $\log(m/k)$, but the probability of reaching these values of $\tilde{k}$ drops exponentially. By applying Lemma 4.6 inductively, we obtain that the probability that $\tilde{k} \leq k/2^i$ when the algorithm terminates is at most $1/2^i$. This means the expected communication cost of Algorithm 2 after $\tilde{k} \leq k$ is (up to a constant) upper bounded by:

$$\sum_{i \geq 0} \log^2(2^i \cdot m/k) \cdot 2^{-i} = \sum_{i \geq 0} (i^2 + \log^2(m/k) + 2i \log(m/k)) \cdot 2^{-i},$$

which is $O(\log^2(m/k))$. ■

4.3 Stitching Things Together

To get our $O(n)$ -bit communication protocol for $(\Delta + 1)$ -coloring, we simply plug Algorithm 2 into the coloring steps of Algorithm 1.

Algorithm 3. An $O(n)$ communication protocol for $(\Delta + 1)$ -coloring.

- (1) Alice and Bob choose a random permutation π with public randomness.
- (2) They iterate over the vertices ordered by π , and on vertex v , they run Algorithm 2 with $m = \Delta + 1$, X (respectively Y) equal to the set of colors used in $N_A(v)$ (respectively $N_B(v)$).

PROOF OF THEOREM 1. As before, for a vertex v and permutation π , let d_v^π denote the degree of v among its predecessors in π , and k_v^π denote $\Delta + 1 - d_v^\pi$. By Lemma 4.3, Algorithm 2 colors v in $O(\log^2((\Delta + 1)/k_v^\pi))$.

$1/k_v^\pi$) bits of communication. To finish proving Theorem 1, we will use the randomness of π to show that this quantity is a constant in expectation.

Since π is a uniformly random permutation, k_v^π is uniformly random over $\{\Delta + 1 - d_v, \dots, \Delta + 1\}$, and the expected cost of coloring v is (for some constant c):

$$\begin{aligned} \frac{1}{d_v + 1} \sum_{i=\Delta+1-d_v}^{\Delta+1} c \log^2\left(\frac{\Delta+1}{i}\right) &\leq \frac{1}{\Delta+1} \sum_{i=1}^{\Delta+1} c \log^2\left(\frac{\Delta+1}{i}\right) \\ &\leq \frac{1}{\Delta+1} \sum_{i=1}^{\Delta+1} 3c \cdot \sqrt{\frac{\Delta+1}{i}}, \end{aligned}$$

where the first inequality is by Proposition 2.1, and the second inequality holds because $\log^2(x) \leq 3\sqrt{x}$ for all $x \geq 1$. Then by using the standard integral upper bound, this is at most:

$$\begin{aligned} \frac{1}{\Delta+1} \int_1^{\Delta+1} 3c \cdot \sqrt{\frac{\Delta+1}{x}} \cdot dx &\leq \frac{3c}{\sqrt{\Delta+1}} \int_1^{\Delta+1} \frac{1}{\sqrt{x}} \cdot dx \\ &\leq \frac{3c}{\sqrt{\Delta+1}} \cdot 2\sqrt{\Delta+1} = O(1). \quad \blacksquare \end{aligned}$$

Remark 4.8. The analysis above goes through as long as we solve a k -Slack-Int instance in $(m/k)^{1-\varepsilon}$ bits of communication for any $\varepsilon > 0$. This may be useful for other problems where the “coloring a single vertex” task does not admit a very efficient algorithm.

4.4 High Probability Analysis

To complete our analysis, we investigate the probabilistic guarantees given by Algorithm 3. We first show by standard concentration techniques that when Δ is small compared to n , our algorithm uses $O(n)$ bits with high probability. Formally,

Lemma 4.9. Algorithm 3 communicates $O(n)$ bits with probability $1 - 2e^{-\frac{n}{2\Delta^2 \log^2 \Delta}}$.

PROOF. Sample independent random variables $X = (X_1, \dots, X_n)$ uniformly in $(0, 1)$. With probability one, all X_i 's are different and induce the permutation $\pi(v) = |\{u \in [n] : X_u < X_v\}| + 1$ (the plus one is to have the smallest index at one). By coupling, it suffices to analyze permutations produced by the X_i 's. For each vertex v , define $K_v := \Delta + 1 - |\{u \in N(v) : X_u < X_v\}|$. We first show $S = \sum_v \log^2(\frac{\Delta}{K_v})$ is concentrated around $O(n)$. We will then argue that with high probability, Algorithm 3 uses $O(S)$ communication.

We use the method of bounded differences. Let $i \in [n]$ and $x, x' \in (0, 1)^n$ be two vectors which differ only in the i -th coordinate. Let $S(x)$ and $S(x')$ be the value of the sum when $X = x$ and $X = x'$. Compared to x , changing x_i only affects the $\leq \Delta$ neighbors of i by increasing or decreasing K_v by exactly one. Hence, we bound $|\log^2(\frac{\Delta}{k}) - \log^2(\frac{\Delta}{k'})|$ when $|k - k'| \leq 1$. The function $f(y) = \log^2(\frac{\Delta}{y})$ is differentiable when $y \geq 1$, and $|f'(y)| = 2\log(\frac{\Delta}{y})/y \leq 2\log(\Delta)$. Using the mean value theorem, this implies $|f(k) - f(k')| \leq 2\log(\Delta)$. Hence, Proposition 2.5 implies

$$\Pr[S > \mathbb{E}[S] + n] \leq \exp\left(-\frac{2n^2}{n \cdot (\Delta \cdot 2\log \Delta)^2}\right) \leq \exp\left(-\frac{n}{2\Delta^2 \log^2 \Delta}\right).$$

As shown in the proof of Theorem 1, we have $\mathbb{E}[S] = O(n)$, thus we condition on the high probability event that $S = O(n)$. We now bound the communication of the algorithm. In Algorithm 2, the number of trials before $\tilde{k} \leq k$ is bounded deterministically. Once $\tilde{k} \leq k$, each sampling succeeds with (at least) $1/2$ probability, and after $\log \Delta$ unsuccessful trials, it succeeds with probability one. Hence, the number of trials for vertex v is stochastically dominated by T_v , a geometric variable of parameter $1/2$ capped at $\log \Delta$. It is easy to verify $\mathbb{E}[T_v] \leq 2$. The total cost of the algorithm is bounded (up to constant factors) by $\sum_v T_v \log^2(\frac{\Delta}{K_v})$, which is $\leq 2S = O(n)$ in expectation. As $\{T_v\}$ are independent, the Hoeffding bound (Proposition 2.4) on variables $T_v \log^2(\frac{\Delta}{K_v}) \leq \log^3 \Delta$ (recall K_v are now fixed) implies

$$\Pr\left[\sum_v T_v \log^2\left(\frac{\Delta}{K_v}\right) > 2S + n\right] \leq \exp\left(-\frac{2n}{\log^6 \Delta}\right). \quad \blacksquare$$

The reason we fail to provide high probability guarantees when Δ is large is that, when sampling a uniform permutation, changing the position of a single element may affect the cost of many vertices. By revealing the permutation “in batches”, we provide better probabilistic guarantees for large Δ at the cost of a slightly weaker bound on communication.

Lemma 4.10. When $\Delta \geq \frac{\sqrt{n}}{c \log^3 n}$, Algorithm 3 communicates $O(n \log^* \Delta)$ bits w.p. $1 - e^{-\Theta(\frac{\sqrt{n}}{\log^3 n})}$.

PROOF. Let $r = \log^* \Delta - 3$. For each $i \in [r]$, define V_i inductively as the last $n_i = \frac{n}{(\log^{(i)} \Delta)^2}$ vertices in $V \setminus (V_1 \cup \dots \cup V_{i-1})$ according to π . For brevity, write $V_{\leq i} = V_1 \cup \dots \cup V_i$ (and similarly for $< i$, $\geq i$, $> i$). We also define $V_{r+1} = V \setminus V_{\leq r}$. The algorithm colors all vertices in V_{r+1} first, then those in V_r and so on and so forth until it finishes by coloring V_1 . The crux of the proof is to show the following property – which we refer to as $\mathcal{P}$ – about sets V_i holds with high probability:

$$\text{For each } i \in [r] \text{ and } v \in V \setminus V_{\leq i} \text{ with } d_v \geq \Delta/2 \text{ we have } |N(v) \cap V_{\leq i}| \geq k_i = \frac{\Delta}{8(\log^{(i)} \Delta)^2}.$$

Before proving $\mathcal{P}$, we assume it holds and show the lemma. Let $k_0 = 1$. We focus on vertices with $d_v \geq \Delta/2$ as low degree vertices always have $\Delta/2$ available colors (hence, Algorithm 2 ends in $O(1)$ trials). We assume henceforth $d_v \geq \Delta/2$. Observe that each $v \in V_i$ has $k_v^\pi \geq k_{i-1}$ for all $i \in [r+1]$. By Lemma 4.6, each trial in Algorithm 3 succeeds with probability at least $1/2$. Therefore, the Chernoff Bound on geometric variables (Proposition 2.3) implies the total number of trials for coloring V_i is $O(n_i)$ w.p. $1 - \exp(-n_i)$. In particular, coloring V_i takes $O(n_i \log(\Delta/k_{i-1})) = O(n)$ communication. By union bound, communication exceeds $O(nr)$ with probability $\leq \sum_i \exp(-n_i) \leq \exp(-\Omega(n_1)) = \exp(-\Omega(\frac{n}{\log^2 n}))$. Hence, coloring all vertices takes $O(nr) = O(n \log^* \Delta)$ bits.

We now prove $\mathcal{P}$ to wrap up. Fix some $i \in [r]$ and $v \notin V_{\leq i}$ with $d_v \geq \Delta/2$. We show it has at least k_i neighbors in $V_{\leq i}$ with high probability, and the result follows by union bound. Fix $V_{< i}$ such that v has fewer than k_i neighbors in $V_{< i}$ (otherwise v already satisfies $\mathcal{P}$). The set V_i is constructed by sampling n_i vertices without

replacement from the set $V \setminus V_{<i}$ of size $N = n - \sum_{j < i} n_j \geq n - 2n_j \geq n/2$. Note that v has $|N(v) \setminus V_{<i}|$ neighbors which can be sampled in V_i , which is $M := d_v - |N(v) \cap V_{<i}| \geq \Delta/2 - k_i \geq \Delta/4$ vertices. Let K count the number of neighbors of v sampled in V_i . We expect $\mathbb{E}[K] = \frac{M}{N} n_i \geq 2k_i$. The samples are not independent, since V_i is sampled without replacement, but as Hoeffding showed in [18, Theorem 4], the Chernoff Bound for variables sampled with replacement can be transferred to variables without. Hence, using Proposition 2.2,

$$\begin{aligned} \Pr[K < k_i] &= \Pr\left[K < \mathbb{E}[K]/2\right] \\ &\leq \exp\left(-\mathbb{E}[K]/12\right) \leq \exp\left(-\frac{\sqrt{n}}{12c \log^5 n}\right), \end{aligned}$$

where the last inequality holds from the assumption that Δ is large. ■

5 LOWER BOUND

In this section, we provide a simple lower bound matching our upper bounds.

THEOREM 3. *Any constant-error randomized protocol for computing a $\Delta + 1$ -coloring on n -vertices graphs requires $\Omega(n)$ communication in the worst-case.*


The proof is by a simple reduction from the problem of sending an n -bit string. Alice constructs a degree-2 graph G with $4n$ vertices such that if Bob knows a proper 3-coloring of G , he can recover x . We remark that in this construction, Alice has all the edges. Hence, it also lower bounds the non-deterministic complexity of $\Delta + 1$ -coloring.

Proposition 5.1. *Suppose Alice is given a uniformly random string $x \in \{0, 1\}^n$ and Bob wants to learn x . Any deterministic protocol where Bob can recover x with probability $1/2$ must communicate $\Omega(n)$ bits.*

The proof of Proposition 5.1 is easy, for example by using Fano's inequality to show that the mutual information between the transcript of any constant error deterministic protocol and the random string x must be $\Omega(n)$. By Yao's minimax principle, a distributional lower bound against deterministic protocols implies a worst-case lower bound against randomized protocols.

We describe the reduction to complete the proof of Theorem 3.

Figure 1: The gadget encoding a single bit. The dashed red edges are present when the bit is 0, and the dotted blue edges are present when the bit is 1. The solid black edges are always present.



PROOF OF THEOREM 3. Suppose Alice is given a uniformly random $x \in \{0, 1\}^n$. Let $v_1, \dots, v_{4n}$ be the vertices of G . Gadget $i \in [n]$ uses vertices $v_{4(i-1)+1}, v_{4(i-1)+2}, v_{4(i-1)+3}, v_{4(i-1)+4}$. We

describe two graphs H_0 and H_1 on four vertices $V_H = \{1, 2, 3, 4\}$. For each $i \in [n]$, Alice constructs a copy of H_{x_i} by mapping each vertex $j \in [4]$ to $v_{4(i-1)+j}$. Clearly, if Bob can deduce whether $H = H_0$ or $H = H_1$ from a proper coloring of H , he can recover x from a proper coloring of G . Proposition 5.1 implies Alice must send $\Omega(n)$ bits to Bob for the protocol to succeed with constant probability. We now describe the gadget graphs H_0, H_1 (see also Figure 1).

Let $x \in \{0, 1\}$. We describe edges of H_x on vertex set $[4]$. For any value of x , add edges $\{1, 3\}$ and $\{2, 4\}$. If $x = 0$, put edges $\{1, 2\}$ and $\{3, 4\}$. Otherwise, if $x = 1$, put edges $\{1, 4\}$ and $\{2, 3\}$. Let C be a proper 3-coloring of H_0 . Either $C(1) = C(4)$ or $C(2) = C(3)$. Indeed, if $C(1) \neq C(4)$, then only one out of three colors remain available to 2 and 3. On the other hand, no proper coloring of H_1 can have $C(1) = C(4)$ or $C(2) = C(3)$. Hence by checking these two equalities, Bob can deduce x from a proper coloring of H_x , which concludes the proof. ■

6 OPEN PROBLEMS

The most immediate open problem from our work is:

Problem 2. *What is the communication complexity of finding the maximum degree Δ of a graph?*

If we had an $O(n)$ -bit communication protocol for this problem, we could remove the requirement of knowing Δ in Theorem 1. However, we know that the communication complexity of this problem is $\omega(n)$. Consider the following problem: Alice (respectively Bob) is given as input the sequence of integers $a_1, \dots, a_n \in [n]^n$ (resp. $b_1, \dots, b_n$), and they wish to compute $\max_i \{a_i + b_i\}$. This problem can be reduced to finding the maximum degree of a graph on $n + 2n$ vertices as follows: For each $i \in [n]$, Alice (resp. Bob) adds a_i edges from i to vertices in $\{n + 1, \dots, 2n\}$ (resp. $\{2n + 1, \dots, 3n\}$), starting one vertex ahead of where the last edge was added, cycling at $2n$ (resp. $3n$). Note that verifying a guess x for the answer to this problem is equivalent to deciding if there is an i such that $a_i \geq x - b_i$, which has a well known lower bound of $\Omega(n \log \log n)$ (see [3, Proposition 3.1] for some details).

Next, can we remove the assumption that the edges are partitioned between Alice and Bob?

Problem 3. *What is the communication complexity of $(\Delta + 1)$ -coloring when the edge sets E_A and E_B are not disjoint?*

Typically, when we drop this assumption, even trivial decision problems (is G a clique, or a clique minus an edge) balloon to $\Omega(n^2)$ communication complexity. However, we observed in Remark 4.2 that Algorithm 1 still works when the edge sets are not disjoint (albeit with foreknowledge of Δ), so the answer is between n and $n \log \Delta$.

Problem 4. *What is the deterministic communication complexity of $(\Delta + 1)$ -coloring?*

We saw in Remark 4.5 that there is a deterministic protocol with $O(n \log^2 \Delta)$ communication.

Finally, we look at two problems where our algorithmic techniques completely break down.

Problem 5. *What is the communication complexity of Δ -coloring?*

The crucial difference between Δ and $(\Delta + 1)$ -coloring is that a proper partial Δ -coloring may not necessarily extend to a complete coloring. This means that the greedy algorithm that is the backbone of both Lemma 3.1 and our protocols no longer works on an arbitrary permutation of the vertices. It is possible to adapt Lovász's proof [25] of Brooks' Theorem to a $O(n \log n + n \log^2 \Delta)$ deterministic communication protocol, using the protocol of Remark 4.5 as a sub-routine.

Problem 6. Let κ denote the degeneracy of the input graph G . What is the communication complexity of $(\kappa + 1)$ -coloring?

Recall that the degeneracy of a graph is the minimum over all permutations of the maximum left-degree of a vertex. The classical algorithm to find a $(\kappa + 1)$ -coloring is to greedily color the vertices ordered by a permutation achieving minimum left degree κ .

A recent manuscript [4] shows that this permutation can be found deterministically in $O(n \log^3 n)$ bits of communication, and hence there is an $O(n \log^3 n + n \log^2 \Delta)$ communication protocol for $(\kappa + 1)$ -coloring.

ACKNOWLEDGMENTS

We thank Sepehr Assadi for several useful discussions and comments on an earlier version of this paper. We also thank the anonymous reviews of PODC '24 for their helpful comments.

The first author is supported by the Icelandic Research Fund (grant 2310015-051). The second author is supported in part by a David R. Cheriton Graduate Scholarship and Sepehr Assadi's Sloan Research Fellowship, an NSERC Discovery grant, and a startup grant from University of Waterloo.

REFERENCES

- [1] Sepehr Assadi, Amit Chakrabarti, Prantar Ghosh, and Manuel Stoeckl. 2023. Coloring in Graph Streams via Deterministic and Adversarially Robust Algorithms. In *Proceedings of the 42nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2023, Seattle, WA, USA, June 18-23, 2023*, Floris Geerts, Hung Q. Ngo, and Stavros Sintos (Eds.). ACM, 141–153. <https://doi.org/10.1145/3584372.3588681> arXiv:2212.10641
- [2] Sepehr Assadi, Yu Chen, and Sanjeev Khanna. 2019. Sublinear algorithms for $(\Delta + 1)$ vertex coloring. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, Philadelphia, PA, 767–786. <https://doi.org/10.1137/1.9781611975482.48> arXiv:1807.08886
- [3] Sepehr Assadi and Aditi Dudeja. 2021. A Simple Semi-Streaming Algorithm for Global Minimum Cuts. In *4th Symposium on Simplicity in Algorithms, SOSA 2021, Virtual Conference, January 11-12, 2021*, Hung Viet Le and Valerie King (Eds.). SIAM, 172–180. <https://doi.org/10.1137/1.9781611976496.19>
- [4] Sepehr Assadi, Prantar Ghosh, Bruno Loff, Parth Mittal, and Sagnik Mukhopadhyay. [n. d.]. Polynomial Pass Semi-Streaming Lower Bounds for K-Cores and Degeneracy. ([n. d.]). Under submission.
- [5] Nir Bachrach, Keren Censor-Hillel, Michal Dory, Yuval Efron, Dean Leitersdorf, and Ami Paz. 2019. Hardness of Distributed Optimization. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 238–247. <https://doi.org/10.1145/3293611.3331597> arXiv:1905.10284
- [6] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. 2016. The locality of distributed symmetry breaking. *J. ACM* 63, 3 (2016), Art. 20, 45. <https://doi.org/10.1145/2903137> arXiv:1202.1983
- [7] Sayan Bhattacharya, Deeparnab Chakrabarty, Monika Henzinger, and Danupon Nanongkai. 2018. Dynamic Algorithms for Graph Coloring. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, Artur Czumaj (Ed.). SIAM, 1–20. <https://doi.org/10.1137/1.9781611975031.1> arXiv:1711.04355
- [8] Sayan Bhattacharya, Fabrizio Grandoni, Janardhan Kulkarni, Quanquan C. Liu, and Shay Solomon. 2022. Fully dynamic $(\Delta + 1)$ -coloring in $O(1)$ update time. *ACM Trans. Algorithms* 18, 2 (2022), Art. 10, 25. <https://doi.org/10.1145/3494539> arXiv:1910.02063
- [9] Joakim Blikstad, Jan van den Brand, Yuval Efron, Sagnik Mukhopadhyay, and Danupon Nanongkai. [2022] ©2022. Nearly optimal communication and query complexity of bipartite matching. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science—FOCS 2022*. IEEE Computer Soc., Los Alamitos, CA, 1174–1185. arXiv:2208.02526
- [10] Amit Chakrabarti and Oded Regev. 2012. An optimal lower bound on the communication complexity of gap-Hamming-distance. *SIAM J. Comput.* 41, 5 (2012), 1299–1317. <https://doi.org/10.1137/120861072> arXiv:1009.3460
- [11] Yi-Jun Chang, Manuela Fischer, Mohsen Ghaffari, Jara Uitto, and Yufan Zheng. 2019. The Complexity of $(\Delta + 1)$ Coloring in Congested Clique, Massively Parallel Computation, and Centralized Local Computation. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*. ACM, 471–480. <https://doi.org/10.1145/3293611.3331607> arXiv:1808.08419
- [12] Yi-Jun Chang, Wenzheng Li, and Seth Pettie. 2020. Distributed $(\Delta + 1)$ -coloring via ultrafast graph shattering. *SIAM J. Comput.* 49, 3 (2020), 497–539. <https://doi.org/10.1137/19M1249527>
- [13] Artur Czumaj, Peter Davies, and Merav Parter. 2021. Simple, deterministic, constant-round coloring in congested clique and MPC. *SIAM J. Comput.* 50, 5 (2021), 1603–1626. <https://doi.org/10.1137/20M1366502> arXiv:2009.06043
- [14] Devdatt P. Dubhashi and Alessandro Panconesi. 2009. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, Cambridge. xvi+196 pages. <https://doi.org/10.1017/CBO9780511581274>
- [15] Mohsen Ghaffari and Fabian Kuhn. [2022] ©2022. Deterministic distributed vertex coloring: simpler, faster, and without network decomposition. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021*. IEEE Computer Soc., Los Alamitos, CA, 1009–1020. <https://doi.org/10.1109/FOCS52979.2021.00101> arXiv:2011.04511
- [16] András Hajnal, Wolfgang Maass, and György Turán. 1988. On the Communication Complexity of Graph Properties. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, Janos Simon (Ed.). ACM, 186–191. <https://doi.org/10.1145/62212.62228>
- [17] Monika Henzinger and Pan Peng. 2022. Constant-time dynamic $(\Delta + 1)$ -coloring. *ACM Trans. Algorithms* 18, 2 (2022), Art. 16, 21. <https://doi.org/10.1145/3501403> arXiv:1907.04745
- [18] Wassily Hoeffding. 1994. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding* (1994), 409–426.
- [19] Piotr Indyk and David P. Woodruff. 2003. Tight Lower Bounds for the Distinct Elements Problem. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*. IEEE Computer Society, 283–288. <https://doi.org/10.1109/SFCS.2003.1238202>
- [20] Bala Kalyanasundaram and Georg Schnitger. 1992. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math.* 5, 4 (1992), 545–557. <https://doi.org/10.1137/0405044>
- [21] Mauricio Karchmer and Avi Wigderson. 1990. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discrete Math.* 3, 2 (1990), 255–265. <https://doi.org/10.1137/0403021>
- [22] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. 2000. On the hardness of approximating the chromatic number. *Combinatorica* 20, 3 (2000), 393–415. <https://doi.org/10.1007/s004930070013>
- [23] Eyal Kushilevitz and Noam Nisan. 1997. *Communication complexity*. Cambridge University Press, Cambridge. xiv+189 pages.
- [24] Nathan Linial. 1992. Locality in distributed graph algorithms. *SIAM J. Comput.* 21, 1 (1992), 193–201. <https://doi.org/10.1137/0221015>
- [25] L. Lovász. 1975. Three short proofs in graph theory. *J. Combinatorial Theory Ser. B* 19, 3 (1975), 269–271. [https://doi.org/10.1016/0095-8956\(75\)90089-1](https://doi.org/10.1016/0095-8956(75)90089-1)
- [26] Carsten Lund and Mihalis Yannakakis. 1994. On the hardness of approximating minimization problems. *J. Assoc. Comput. Mach.* 41, 5 (1994), 960–981. <https://doi.org/10.1145/185675.306789>
- [27] Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. 1998. On data structures and asymmetric communication complexity. Vol. 57. 37–49. <https://doi.org/10.1006/jcss.1998.1577> 27th Annual ACM Symposium on the Theory of Computing (STOC'95) (Las Vegas, NV).
- [28] Ilan Newman. 1991. Private vs. common random bits in communication complexity. *Inform. Process. Lett.* 39, 2 (1991), 67–71. [https://doi.org/10.1016/0020-0190\(91\)90157-D](https://doi.org/10.1016/0020-0190(91)90157-D)
- [29] Anup Rao and Amir Yehudayoff. 2020. *Communication complexity and applications*. Cambridge University Press, Cambridge. xvii+251 pages.
- [30] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. 2011. Distributed verification and hardness of distributed approximation. In *STOC'11—Proceedings of the 43rd ACM Symposium on Theory of Computing*. ACM, New York, 363–372. <https://doi.org/10.1145/1993636.1993686> arXiv:1011.3049
- [31] Andrew Chi-Chih Yao. 1979. Some Complexity Questions Related to Distributive Computing (Preliminary Report). In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, Michael J. Fischer, Richard A. DeMillo, Nancy A. Lynch, Walter A. Burkhard, and Alfred V. Aho (Eds.). ACM, 209–213. <https://doi.org/10.1145/800135.804414>

- [32] Yufei Zhao. 2017. Extremal regular graphs: independent sets and graph homomorphisms. *Amer. Math. Monthly* 124, 9 (2017), 827–843. <https://doi.org/10.4169/amer.math.monthly.124.9.827> arXiv:1610.09210

[org/10.4169/amer.math.monthly.124.9.827](https://doi.org/10.4169/amer.math.monthly.124.9.827) arXiv:1610.09210