

Adversarially Robust Coloring for Graph Streams

Amit Chakrabarti  

Department of Computer Science, Dartmouth College, Hanover, NH, USA

Prantar Ghosh 

Department of Computer Science, Dartmouth College, Hanover, NH, USA

Manuel Stoeckl  

Department of Computer Science, Dartmouth College, Hanover, NH, USA

Abstract

A streaming algorithm is considered to be adversarially robust if it provides correct outputs with high probability even when the stream updates are chosen by an adversary who may observe and react to the past outputs of the algorithm. We grow the burgeoning body of work on such algorithms in a new direction by studying robust algorithms for the problem of maintaining a valid vertex coloring of an n -vertex graph given as a stream of edges. Following standard practice, we focus on graphs with maximum degree at most Δ and aim for colorings using a small number $f(\Delta)$ of colors.

A recent breakthrough (Assadi, Chen, and Khanna; SODA 2019) shows that in the standard, non-robust, streaming setting, $(\Delta + 1)$ -colorings can be obtained while using only $\tilde{O}(n)$ space. Here, we prove that an adversarially robust algorithm running under a similar space bound must spend almost $\Omega(\Delta^2)$ colors and that robust $O(\Delta)$ -coloring requires a *linear* amount of space, namely $\Omega(n\Delta)$. We in fact obtain a more general lower bound, trading off the space usage against the number of colors used. From a complexity-theoretic standpoint, these lower bounds provide (i) the first significant separation between adversarially robust algorithms and ordinary randomized algorithms for a *natural* problem on insertion-only streams and (ii) the first significant separation between randomized and deterministic coloring algorithms for graph streams, since deterministic streaming algorithms are automatically robust.

We complement our lower bounds with a suite of positive results, giving adversarially robust coloring algorithms using sublinear space. In particular, we can maintain an $O(\Delta^2)$ -coloring using $\tilde{O}(n\sqrt{\Delta})$ space and an $O(\Delta^3)$ -coloring using $\tilde{O}(n)$ space.

2012 ACM Subject Classification Theory of computation \rightarrow Streaming, sublinear and near linear time algorithms

Keywords and phrases Data streaming, graph algorithms, graph coloring, lower bounds, online algorithms

Digital Object Identifier 10.4230/LIPIcs.ITCS.2022.37

Related Version *Full Version:* <https://arxiv.org/abs/2109.11130>

Funding This work was supported in part by NSF under awards 1907738 and 2006589.

Acknowledgements Prantar Ghosh would like to thank Sayan Bhattacharya for a helpful conversation regarding this work.

1 Introduction

A data streaming algorithm processes a huge input, supplied as a long sequence of elements, while using working memory (i.e., space) much smaller than the input size. The main algorithmic goal is to compute or estimate some function of the input σ while using space *sublinear* in the size of σ . For most – though not all – problems of interest, a streaming algorithm *needs* to be randomized in order to achieve sublinear space. For a randomized algorithm, the standard correctness requirement is that for each possible input stream it return a valid answer with high probability. A burgeoning body of work – much of it very



© Amit Chakrabarti, Prantar Ghosh, and Manuel Stoeckl;
licensed under Creative Commons License CC-BY 4.0

13th Innovations in Theoretical Computer Science Conference (ITCS 2022).

Editor: Mark Braverman; Article No. 37; pp. 37:1–37:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

recent [10, 11, 21, 24, 17, 29, 8, 9] but traceable back to [20] – addresses streaming algorithms that seek an even stronger correctness guarantee, namely that they produce valid answers with high probability even when working with an input generated by an active adversary. There is compelling motivation from practical applications for seeking this stronger guarantee: for instance, consider a user continuously interacting with a database and choosing future queries based on past answers received; or think of an online streaming or marketing service looking at a customer’s transaction history and recommending them products based on it.

We may view the operation of streaming algorithm \mathcal{A} as a game between a *solver*, who executes \mathcal{A} , and an *adversary*, who generates a “hard” input stream σ . The standard notion of \mathcal{A} having error probability δ is that for every fixed σ that the adversary may choose, the probability over \mathcal{A} ’s random choices that it errs on σ is at most δ . Since the adversary has to make their choice before the solver does any work, they are *oblivious* to the actual actions of the solver. In contrast to this, an *adaptive adversary* is not required to fix all of σ in advance, but can generate the elements (tokens) of σ incrementally, based on outputs generated by the solver as it executes \mathcal{A} . Clearly, such an adversary is much more powerful and can attempt to learn something about the solver’s internal state in order to generate input tokens that are bad for the particular random choices made by \mathcal{A} . Indeed, such adversarial attacks are known to break many well known algorithms in the streaming literature [20, 10]. Motivated by this, one defines a δ -error *adversarially robust streaming algorithm* to be one where the probability that an adaptive adversary can cause the solver to produce an incorrect output at some point of time is at most δ . Notice that a *deterministic* streaming algorithm (which, by definition, must always produce correct answers) is automatically adversarially robust.

Past work on such adversarially robust streaming algorithms has focused on statistical estimation problems and on sampling problems but, with the exception of [17], there has not been much study of graph theoretic problems. This work focuses on graph coloring, a fundamental algorithmic problem on graphs. Recall that the goal is to efficiently process an input graph given as a stream of edges and assign colors to its vertices from a small palette so that no two adjacent vertices receive the same color. The main messages of this work are that (i) while there exist surprisingly efficient sublinear-space algorithms for coloring under standard streaming, it is provably harder to obtain adversarially robust solutions; but nevertheless, (ii) there *do* exist nontrivial sublinear-space robust algorithms for coloring.

To be slightly more detailed, suppose we must color an n -vertex input graph G that has maximum degree Δ . Producing a coloring using only $\chi(G)$ colors, where $\chi(G)$ is the chromatic number, is NP-hard while producing a $(\Delta + 1)$ -coloring admits a straightforward greedy algorithm, given offline access to G . Producing a good coloring given only streaming access to G and sublinear (i.e., $o(n\Delta)$ bits of) space is a nontrivial problem and the subject of much recent research [13, 7, 12, 4, 14], including the breakthrough result of Assadi, Chen, and Khanna [7] that gives a $(\Delta + 1)$ -coloring algorithm using only semi-streaming (i.e., $\tilde{O}(n)$ bits of) space.¹ However, all of these algorithms were designed with only the standard, oblivious adversary setting in mind; an adaptive adversary can make all of them fail. This is the starting point for our exploration in this work.

1.1 Our Results and Contributions

We ask whether the graph coloring problem is inherently harder under an adversarial robustness requirement than it is for standard streaming. We answer this question affirmatively with the first major theorem in this work, which is the following (we restate the theorem with more detail and formality as Theorem 8).

¹ The notation $\tilde{O}(\cdot)$ hides factors polylogarithmic in n .

► **Theorem 1.** *A constant-error adversarially robust algorithm that processes a stream of edge insertions into an n -vertex graph and, as long as the maximum degree of the graph remains at most Δ , maintains a valid K -coloring (with $\Delta + 1 \leq K \leq n/2$) must use at least $\Omega(n\Delta^2/K)$ bits of space.*

We spell out some immediate corollaries of this result because of their importance as conceptual messages.

- **Robust coloring using $O(\Delta)$ colors.** In the setting of Theorem 1, if the algorithm is to use only $O(\Delta)$ colors, then it must use $\Omega(n\Delta)$ space. In other words, a sublinear-space solution is ruled out.
- **Robust coloring using semi-streaming space.** In the setting of Theorem 1, if the algorithm is to run in only $\tilde{O}(n)$ space, then it must use $\tilde{\Omega}(\Delta^2)$ colors.
- **Separating robust from standard streaming with a natural problem.** Contrast the above two lower bounds with the guarantees of the [7] algorithm, which handles the non-robust case. This shows that “maintaining an $O(\Delta)$ -coloring of a graph” is a natural (and well-studied) algorithmic problem where, even for insertion-only streams, the space complexities of the robust and standard streaming versions of the problem are well separated: in fact, the separation is roughly quadratic, by taking $\Delta = \Theta(n)$. This answers an open question of [24], as we explain in greater detail in Section 1.2.
- **Deterministic versus randomized coloring.** Since every deterministic streaming algorithm is automatically adversarially robust, the lower bound in Theorem 1 applies to such algorithms. In particular, this settles the deterministic complexity of $O(\Delta)$ -coloring. Also, turning to semi-streaming algorithms, whereas a combinatorially optimal² $(\Delta + 1)$ -coloring is possible using randomization [7], a deterministic solution must spend at least $\tilde{\Omega}(\Delta^2)$ colors. These results address a broadly-stated open question of Assadi [5]; see Section 1.2 for details.

We prove the lower bound in Theorem 1 using a reduction from a novel two-player communication game that we call SUBSET-AVOIDANCE. In this game, Alice is given an a -sized subset of the universe $[t]$;³ she must communicate a possibly random message to Bob that causes him to output a b -sized subset of $[t]$ that, with high probability, avoids Alice’s set completely. We give a fairly tight analysis of the communication complexity of this game, showing an $\Omega(ab/t)$ lower bound, which is matched by an $\tilde{O}(ab/t)$ deterministic upper bound. The SUBSET-AVOIDANCE problem is a natural one. We consider the definition of this game and its analysis – which is not complicated – to be additional conceptual contributions of this work; these might be of independent interest for future applications.

We complement our lower bound with some good news: we give a suite of upper bound results by designing adversarially robust coloring algorithms that handle several interesting parameter regimes. Our focus is on maintaining a valid coloring of the graph using $\text{poly}(\Delta)$ colors, where Δ is the current maximum degree, as an adversary inserts edges. In fact, some of these results hold even in a *turnstile model*, where the adversary might both add and delete edges. In this context, it is worth noting that the [7] algorithm also works in a turnstile setting.

² If one must use at most $f(\Delta)$ colors for some function f , the best possible function that always works is $f(\Delta) = \Delta + 1$.

³ The notation $[t]$ denotes the set $\{1, 2, \dots, t\}$.

► **Theorem 2.** *There exist adversarially robust algorithms for coloring an n -vertex graph achieving the following tradeoffs (shown in Table 1) between the space used for processing the stream and the number of colors spent, where Δ denotes the evolving maximum degree of the graph and, in the turnstile setting, m denotes a known upper bound on the stream length.*

■ **Table 1** A summary of our adversarially robust coloring algorithms. A “strict graph turnstile” model requires the input to describe a simple graph at all times; see Section 3.

Model	Colors	Space	Notes	Reference
Insertion-only	$O(\Delta^3)$	$\tilde{O}(n)$	$\tilde{O}(n\Delta)$ external random bits	Theorem 10
Insertion-only	$O(\Delta^k)$	$\tilde{O}(n\Delta^{1/k})$	any $k \in \mathbb{N}$	Corollary 15
Strict Graph Turnstile	$O(\Delta^k)$	$\tilde{O}(n^{1-1/k}m^{1/k})$	constant $k \in \mathbb{N}$	Theorem 14

In each of these algorithms, for each stream update or query made by the adversary, the probability that the algorithm fails either by returning an invalid coloring or aborting is at most $1/\text{poly}(n)$.

We give a more detailed discussion of these results, including an explanation of the technical caveat noted in Table 1 for the $O(\Delta^3)$ -coloring algorithm, in Section 2.2.

1.2 Motivation, Context, and Related Work

Graph streaming has become widely popular [25], especially since the advent of large and evolving networks including social media, web graphs, and transaction networks. These large graphs are regularly mined for knowledge and such knowledge often informs their future evolution. Therefore, it is important to have adversarially robust algorithms for working with these graphs. Yet, the recent explosion of interest in robust algorithms has not focused much on graph problems. We now quickly recap some history.

Two influential works [26, 20] identified the challenge posed by adaptive adversaries to sketching and streaming algorithms. In particular, Hardt and Woodruff [20] showed that many statistical problems, including the ubiquitous one of ℓ_2 -norm estimation, do not admit adversarially robust linear sketches of sublinear size. Recent works have given a number of positive results. Ben-Eliezer, Jayaram, Woodruff, and Yogev [10] considered such fundamental problems as distinct elements, frequency moments, and heavy hitters (these date back to the beginnings of the literature on streaming algorithms); for $(1 \pm \varepsilon)$ -approximating a function value, they gave two generic frameworks that can “robustify” a standard streaming algorithm, blowing up the space cost by roughly the *flip number* $\lambda_{\varepsilon, m}$, defined as the maximum number of times the function value can change by a factor of $1 \pm \varepsilon$ over the course of an m -length stream. For *insertion-only streams* and monotone functions, $\lambda_{\varepsilon, m}$ is roughly $O(\varepsilon^{-1} \log m)$, so this overhead is very small. Subsequent works [21, 29, 8] have improved this overhead with the current best-known one being $O(\sqrt{\varepsilon \lambda_{\varepsilon, m}})$ [8].

For insertion-only graph streams, a number of well-studied problems such as triangle counting, maximum matching size, and maximum subgraph density can be handled by the above framework because the underlying functions are monotone. For some problems such as counting connected components, there are simple deterministic algorithms that achieve an asymptotically optimal space bound, so there is nothing new to say in the robust setting. For graph sparsification, [17] showed that the Ahn–Guha sketch [2] can be made adversarially robust with a slight loss in the quality of the sparsifier. Thanks to efficient adversarially robust sampling [11, 17], many sampling-based graph algorithms should yield corresponding

robust solutions without much overhead. For problems calling for Boolean answers, such as testing connectivity or bipartiteness, achieving low error against an oblivious adversary automatically does so against an adaptive adversary as well, since a sequence of correct outputs from the algorithm gives away no information to the adversary. This is a particular case of a more general phenomenon captured by the notion of pseudo-determinism, discussed at the end of this section.

Might it be that for all interesting data streaming problems, efficient standard streaming algorithms imply efficient robust ones? The above framework does not automatically give good results for *turnstile* streams, where each token specifies either an insertion or a deletion of an item, or for estimating non-monotone functions. In either of these situations, the flip number can be very large. As noted above, linear sketching, which is the preeminent technique behind turnstile streaming algorithms (including ones for graph problems), is vulnerable to adversarial attacks [20]. This does not quite provide a separation between standard and robust space complexities, since it does not preclude efficient non-linear solutions. The very recent work [24] gives such a separation: it exhibits a function estimation problem for which the ratio between the adversarial and standard streaming complexities is as large as $\tilde{\Omega}(\sqrt{\lambda_{\varepsilon,m}})$, which is exponential upon setting parameters appropriately. However, their function is highly artificial, raising the important question: *Can a significant gap be shown for a natural streaming problem?*⁴

It is easy to demonstrate such a gap in graph streaming. Consider the problem of finding a spanning forest in a graph undergoing edge insertions and deletions. The celebrated Ahn–Guha–McGregor sketch [3] solves this in $\tilde{O}(n)$ space, but this sketch is not adversarially robust. Moreover, suppose that \mathcal{A} is an adversarially robust algorithm for this problem. Then we can argue that the memory state of \mathcal{A} upon processing an unknown graph G must contain enough information to recover G entirely: an adversary can repeatedly ask \mathcal{A} for a spanning forest, delete all returned edges, and recurse until the evolving graph becomes empty. Thus, for basic information theoretic reasons, \mathcal{A} must use $\Omega(n^2)$ bits of space, resulting in a quadratic gap between robust and standard streaming space complexities. Arguably, this separation is not very satisfactory, since the hardness arises from the turnstile nature of the stream, allowing the adversary to delete edges. Meanwhile, the [24] separation does hold for insert-only streams, but as we (and they) note, their problem is rather artificial.

Hardness for Natural Problems. We now make a simple, yet crucial, observation. Let MISSING-ITEM-FINDING (MIF) denote the problem where, given an evolving set $S \subseteq [n]$, we must be prepared to return an element in $[n] \setminus S$ or report that none exists. When the elements of S are given as an input stream, MIF admits the following $O(\log^2 n)$ -space solution against an oblivious adversary: maintain an ℓ_0 -sampling sketch [23] for the characteristic vector of $[n] \setminus S$ and use it to randomly sample a valid answer. In fact, this solution extends to turnstile streams. Now suppose that we have an adversarially robust algorithm \mathcal{A} for MIF, handling insert-only streams. Then, given the memory state of \mathcal{A} after processing an unknown set T with $|T| = n/2$, an adaptive adversary can repeatedly query \mathcal{A} for a missing item x , record x , insert x as the next stream token, and continue until \mathcal{A} fails to find an item. At that point, the adversary will have recorded (w.h.p.) the set $[n] \setminus T$, so he can reconstruct T . As before, by basic information theory, this reconstructability implies that \mathcal{A} uses $\Omega(n)$ space.

⁴ This open question was explicitly raised in the STOC 2021 workshop *Robust Streaming, Sketching, and Sampling* [28].

This exponential gap between standard and robust streaming, based on well-known results, seems to have been overlooked – perhaps because MIF does not conform to the type of problems, namely estimation of real-valued functions, that much of the robust streaming literature has focused on. That said, though MIF is a natural problem and the hardness holds for insert-only streams, there is one important box that MIF does not tick: it is not important enough on its own and so does not command a serious literature. This leads us to refine the open question of [24] thus: *Can a significant gap be shown for a natural and well-studied problem with the hardness holding even for insertion-only streams?*

With this in mind, we return to graph problems, searching for such a gap. In view of the generic framework of [10] and follow-up works, we should look beyond estimating some monotone function of the graph with scalar output. What about problems where the output is a big *vector*, such as approximate maximum matching (not just its size) or approximate densest subgraph (not just the density)? It turns out that the *sketch switching* technique of [10] can still be applied: since we need to change the output only when the estimates of the associated numerical values (matching size and density, respectively) change enough, we can proceed as in that work, switching to a new sketch with fresh randomness that remains unrevealed to the adversary. This gives us a robust algorithm incurring only logarithmic overhead.

But graph coloring is different. As our Theorem 1 shows, it does exhibit a quadratic gap for the right setting of parameters and it is, without doubt, a heavily-studied problem, even in the data streaming setting.

The above hardness of MIF provides a key insight into why graph coloring is hard; see Section 2.1.

Connections with Other Work on Streaming Graph Coloring. Graph coloring is, of course, a heavily-studied problem in theoretical computer science. For this discussion, we stick to streaming algorithms for this problem, which already has a significant literature [13, 1, 7, 12, 4, 14].

Although it is not possible to $\chi(G)$ -color an input graph in sublinear space [1], as [7] shows, there is a semi-streaming algorithm that produces a $(\Delta + 1)$ -coloring. This follows from their elegant *palette sparsification* theorem, which states that if each vertex samples roughly $O(\log n)$ colors from a palette of size $\Delta + 1$, then there exists a proper coloring of the graph where each vertex uses a color only from its sampled list. Hence, we only need to store edges between vertices whose lists intersect. If the edges of G are independent of the algorithm’s randomness, then the expected number of such “conflict” edges is $O(n \log^2 n)$, leading to a semi-streaming algorithm. But note that an adaptive adversary can attack this algorithm by using a reported coloring to learn which future edges would definitely be conflict edges and inserting such edges to blow up the algorithm’s storage.

There are some other semi-streaming algorithms (in the standard setting) that aim for $\Delta(1 + \varepsilon)$ -colorings. One is palette-sparsification based [4] and so, suffers from the above vulnerability against an adaptive adversary. Others [13, 12] are based on randomly partitioning the vertices into clusters and storing only intra-cluster edges, using pairwise disjoint palettes for the clusters. Here, the semi-streaming space bound hinges on the random partition being likely to assign each edge’s endpoints to different clusters. This can be broken by an adaptive adversary, who can use a reported coloring to learn many vertex pairs that are intra-cluster and then insert new edges at such pairs.

Finally, we highlight an important theoretical question about sublinear algorithms for graph coloring: *Can they be made deterministic?* This was explicitly raised by Assadi [5] and, prior to this work, it was open whether, for $(\Delta + 1)$ -coloring, *any* sublinear space bound could

be obtained deterministically. Our Theorem 1 settles the deterministic space complexity of this problem, showing that even the weaker requirement of $O(\Delta)$ -coloring forces $\Omega(n\Delta)$ space, which is linear in the input size.

Parameterizing Theorem 1 differently, we see that a robust (in particular, a deterministic) algorithm that is limited to semi-streaming space must spend $\tilde{\Omega}(\Delta^2)$ colors. A major remaining open question is whether this can be matched, perhaps by a deterministic semi-streaming $O(\Delta^2)$ -coloring algorithm. In fact, it is not known how to get even a $\text{poly}(\Delta)$ -coloring deterministically.⁵ Our algorithmic results, summarized in Theorem 2, make partial progress on this question. Though we do not obtain deterministic algorithms, we obtain adversarially robust ones, and we do obtain $\text{poly}(\Delta)$ -colorings, though not all the way down to $O(\Delta^2)$ in semi-streaming space.

Other Related Work. Pseudo-deterministic streaming algorithms[19] fall between adversarially robust and deterministic ones. Such an algorithm is allowed randomness, but for each particular input stream it must produce one fixed output (or output sequence) with high probability. Adversarial robustness is automatic, because when such an algorithm succeeds, it does not reveal any of its random bits through the outputs it gives. Thus, there is nothing for an adversary to base adaptive decisions on.

The well-trodden subject of dynamic graph algorithms deals with a model closely related to the adaptive adversary model: one receives a stream of edge insertions/deletions and seeks to maintain a solution after each update. There have been a few works on the Δ -based graph coloring problem in this setting [15, 16, 22]. However, the focus of the dynamic setting is on optimizing the update *time* without any restriction on the space usage; this is somewhat orthogonal to the streaming setting where the primary goal is space efficiency, and update time, while practically important, is not factored into the complexity.

2 Overview of Techniques

2.1 Lower Bound Techniques

As might be expected, our lower bounds are best formalized through communication complexity. Recall that a typical communication-to-streaming reduction for proving a one-pass streaming space lower bound works as follows. We set up a communication game for Alice and Bob to solve, using one message from Alice to Bob. Suppose that Alice and Bob have inputs x and y in this game. The players simulate a purported efficient streaming algorithm \mathcal{A} (for P , the problem of interest) by having Alice feed some tokens into \mathcal{A} based on x , communicating the resulting memory state of \mathcal{A} to Bob, having Bob continue feeding tokens into \mathcal{A} based on y , and finally querying \mathcal{A} for an answer to P , based on which Bob can give a good output in the communication game. When this works, it follows that the space used by \mathcal{A} must be at least the one-way (and perhaps randomized) communication complexity of the game. Note, however, that this style of argument where it is possible to solve the game by querying the algorithm only once, is also applicable to an oblivious adversary setting. Therefore, it cannot prove a lower bound any higher than the standard streaming complexity of P .

⁵ Shortly after the announcement of the full version of our paper [18], Assadi, Chen, and Sun [6] announced an independent work proving that this cannot be done.

The way to obtain stronger lower bounds by using the purported adversarial robustness of \mathcal{A} is to design communication protocols where Bob, after receiving Alice's message, proceeds to query \mathcal{A} repeatedly, feeding tokens into \mathcal{A} based on answers to such queries. In fact, in the communication games we shall use for our reductions, Bob will not have any input at all and the goal of the game will be for Bob to recover information about Alice's input, perhaps indirectly. It should be clear that the lower bound for the MIF problem, outlined in Section 1.2, can be formalized in this manner. For our main lower bound (Theorem 1), we use a communication game that can be seen as a souped-up version of MIF.

The Subset-Avoidance Problem. Recall the SUBSET-AVOIDANCE problem described in Section 1.1 and denote it $\text{AVOID}(t, a, b)$. To restate: Alice is given a set $A \subseteq [t]$ of size a and must induce Bob to output a set $B \subseteq [t]$ of size b such that $A \cap B = \emptyset$. The one-way communication complexity of this game can be lower bounded from first principles. Since each output of Bob is compatible with only $\binom{t-b}{a}$ possible input sets of Alice, she cannot send the same message on more than that many inputs. Therefore, she must be able to send roughly $\binom{t}{a} / \binom{t-b}{a}$ distinct messages for a protocol to succeed with high probability. The number of bits she must communicate in the worst case is roughly the logarithm of this ratio, which we show is $\Omega(ab/t)$. Interestingly, we observe that this lower bound is tight and can in fact be matched by a deterministic protocol, as noted in Lemma 7.

In the sequel, we shall need to consider a direct sum version of this problem that we call $\text{AVOID}^k(t, a, b)$, where Alice has a list of k subsets and Bob must produce his own list of subsets, with his i th avoiding the i th subset of Alice. We extend our lower bound argument to show that the one-way complexity of $\text{AVOID}^k(t, a, b)$ is $\Omega(kab/t)$.

Using Graph Coloring to Solve Subset-Avoidance. To explain how we reduce the AVOID^k problem to graph coloring, we focus on a special case of Theorem 1 first. Suppose we have an adversarially robust $(\Delta + 1)$ -coloring streaming algorithm \mathcal{A} . We describe a protocol for solving $\text{AVOID}(t, a, b)$. Let us set $t = \binom{n}{2}$ to have the universe correspond to all possible edges of an n -vertex graph. Suppose Alice's set A has size $a \approx n^2/8$. We show that, given a set of n vertices, Alice can use public randomness to randomly map her elements to the set of vertex-pairs so that the corresponding edges induce a graph G that, w.h.p., has max-degree $\Delta \approx n/4$. Alice proceeds to feed the edges of G into \mathcal{A} and then sends Bob the state of \mathcal{A} .

Bob now queries \mathcal{A} to obtain a $(\Delta + 1)$ -coloring of G . Then, he pairs up like-colored vertices to obtain a maximal pairing. Observe that he can pair up all but at most one vertex from each color class. Thus, he obtains at least $(n - \Delta - 1)/2$ such pairs. Since each pair is monochromatic, they don't share an edge, and hence, Bob has retrieved $(n - \Delta - 1)/2$ missing edges that correspond to elements absent in Alice's set. Since Alice used public randomness for the mapping, Bob knows exactly which elements these are. He now forms a matching with these pairs and inserts the edges to \mathcal{A} . Once again, he queries \mathcal{A} to find a coloring of the modified graph. Observe that the matching can increase the max-degree of the original graph by at most 1. Therefore, this new coloring uses at most $\Delta + 2$ colors. Thus, Bob would retrieve at least $(n - \Delta - 2)/2$ new missing edges. He again adds to the graph the matching formed by those edges and queries \mathcal{A} . It is crucial to note here that he can repeatedly do this and expect \mathcal{A} to output a correct coloring because of its adversarial robustness. Bob stops once the max-degree reaches $n - 1$, since now the algorithm can color each vertex with a distinct color, preventing him from finding a missing edge.

Summing up the sizes of all the matchings added by Bob, we see that he has found $\Theta((n - \Delta)^2)$ elements missing from Alice's set. Since $\Delta \approx n/4$, this is $\Theta(n^2)$. Thus, Alice and Bob have solved the $\text{AVOID}(t, a, b)$ problem where $t = \binom{n}{2}$ and $a, b = \Theta(n^2)$. As outlined above, this requires $\Omega(ab/t) = \Omega(n^2)$ communication. Hence, \mathcal{A} must use at least $\Omega(n^2) = \Omega(n\Delta)$ space.

With some further work, we can generalize the above argument to work for any value of Δ with $1 \leq \Delta \leq n/2$. For this generalization, we use the communication complexity of $\text{AVOID}^k(t, a, b)$ for suitable parameter settings. With more rigorous analysis, we can further generalize the result to apply not only to $(\Delta + 1)$ -coloring algorithms but to any $f(\Delta)$ -coloring algorithm. That is, we can prove Theorem 8.

2.2 Upper Bound Techniques

It is useful to outline our algorithms in an order different from the presentation in Section 5.

A Sketch-Switching-Based $O(\Delta^2)$ -Coloring. The main challenge in designing an adversarially robust coloring algorithm is that the adversary can compel the algorithm to change its output at every point in the stream: he queries the algorithm, examines the returned coloring, and inserts an edge between two vertices of the same color. Indeed, the sketch switching framework of [10] shows that for *function estimation*, one can get around this power of the adversary as follows. Start with a basic (i.e., oblivious-adversary) sketch for the problem at hand. Then, to deal with an adaptive adversary, run multiple independent basic sketches in parallel, changing outputs only when forced to because the underlying function has changed significantly. More precisely, maintain λ independent parallel sketches where λ is the *flip number*, defined as the maximum number of times the function value can change by the desired approximation factor over the course of the stream. Keep track of which sketch is currently being used to report outputs to the adversary. Upon being queried, re-use the most recently given output unless forced to change, in which case discard the current sketch and switch to the next in the list of λ sketches. Notice that this keeps the adversary oblivious to the randomness being used to compute future outputs: as soon as our output reveals any information about the current sketch, we discard it and never use it again to process a stream element.

This way of switching to a new sketch only when forced to ensures that λ sketches suffice, which is great for function estimation. However, since a graph coloring output can be forced to change at every point in a stream of length m , naively implementing this idea would require m parallel sketches, incurring a factor of m in space. We have to be more sophisticated. We combine the above idea with a chunking technique so as to reduce the number of times we need to switch sketches.

Suppose we split the m -length stream into k chunks, each of size m/k . We initialize k parallel sketches of a standard streaming $(\Delta + 1)$ -coloring algorithm \mathcal{C} to be used one at a time as each chunk ends. We store (buffer) an entire chunk explicitly and when we reach its end, we say we have reached a “checkpoint,” use a fresh copy of \mathcal{C} to compute a $(\Delta + 1)$ -coloring of the entire graph at that point, delete the chunk from our memory, and move on to store the next chunk. When a query arrives, we deterministically compute a $(\Delta + 1)$ -coloring of the partial chunk in our buffer and “combine” it with the coloring we computed at the last checkpoint. The combination uses at most $(\Delta + 1)^2 = O(\Delta^2)$ colors. Since a single copy of \mathcal{C} takes $\tilde{O}(n)$ space, the total space used by the sketches is $\tilde{O}(nk)$. Buffering a chunk uses an additional $\tilde{O}(m/k)$ space. Setting k to be $\sqrt{m/n}$, we get the total space usage to be $\tilde{O}(\sqrt{mn}) = \tilde{O}(n\sqrt{\Delta})$, since $m = O(n\Delta)$.

37:10 Adversarially Robust Coloring for Graph Streams

Handling edge deletions is more delicate. This is because we can no longer express the current graph as a union of G_1 (the graph up to the most recent checkpoint) and G_2 (the buffered subgraph) as above. A chunk may now contain an update that deletes an edge which was inserted before the checkpoint, and hence, is not in store. Observe, however, that deleting an edge doesn't violate the validity of a coloring. Hence, if we ignore these edge deletions, the only worry is that they might substantially reduce the maximum degree Δ causing us to use many more colors than desired. Now, note that if we have a $(\Delta_1 + 1)$ -coloring at the checkpoint, then as long as the current maximum degree Δ remains above $\Delta_1/2$, we have a 2Δ -coloring in store. Hence, combining that with a $(\Delta + 1)$ -coloring of the current chunk gives an $O(\Delta^2)$ -coloring. Furthermore, we can keep track of the maximum degree of the graph using only $\tilde{O}(n)$ space and detect the points where it falls below half of what it was at the last checkpoint. We declare each such point as a new "ad hoc checkpoint," i.e., use a fresh sketch to compute a $(\Delta + 1)$ -coloring there. Since the max-degree can decrease by a factor of 2 at most $\log n$ times, we show that it suffices to have only $\log n$ times more parallel sketches initialized at the beginning of the stream. This incurs only an $O(\log n)$ -factor overhead in space. We discuss the algorithm and its analysis in detail in Algorithm 3 and Lemma 13 respectively.

To generalize the above to an $O(\Delta^k)$ -coloring in $\tilde{O}(n\Delta^{1/k})$ space, we use recursion in a manner reminiscent of streaming coreset construction algorithms. Split the stream into $\Delta^{1/k}$ chunks, each of size $n\Delta^{1-1/k}$. Now, instead of storing a chunk entirely and coloring it deterministically, we can recursively color it with Δ^{k-1} colors in $O(n\Delta^{1/k})$ space and combine the coloring with the $(\Delta + 1)$ -coloring at the last checkpoint. The recursion makes the analysis of this algorithm even more delicate, and careful work is needed to argue the space usage and to properly handle deletions in the turnstile setting. The details appear in Theorem 14.

A Palette-Sparsification-Based $O(\Delta^3)$ -Coloring. This algorithm uses a different approach to the problem of the adversary forcing color changes. It ensures that, every time an edge is added, one of its endpoints is randomly recolored, where the color is drawn uniformly from a set $C \setminus K$ of colors, where C is determined by the degree of the endpoint, and K is the set of colors currently held by neighboring vertices. Let R_v denote the random string that drives this color-choosing process at vertex v . When the adversary inserts an edge $\{u, v\}$, the algorithm uses R_u and R_v to determine whether this edge could with significant probability end up with the same vertex color on both ends in the future. If so, the algorithm stores the edge; if not, it can be ignored entirely. It will turn out that when the number of colors is set to establish an $O(\Delta^3)$ -coloring, only an $\tilde{O}(1/\Delta)$ fraction of edges need to be stored, so the algorithm only needs to store $\tilde{O}(n)$ bits of data related to the input. The proof of this storage bound has to contend with an adaptive adversary. We do so by first arguing that despite this adaptivity, the adversary cannot cause the algorithm to use more storage than the worst oblivious adversary could have. We can then complete the proof along traditional lines, using concentration bounds. The details appear in Algorithm 2 and Theorem 10.

There is a technical caveat here. The random string R_v used at each vertex v is about $\tilde{O}(\Delta)$ bits long. Thus, the algorithm can only be called semi-streaming if we agree that these $\tilde{O}(n\Delta)$ random bits do not count towards the storage cost. In the standard streaming setting, this "randomness cost" is not a concern, for we can use the standard technique of invoking Nisan's space-bounded pseudorandom generator [27] to argue that the necessary bits can be generated on the fly and never stored. Unfortunately, it is not clear that this transformation preserves adversarial robustness. Despite this caveat, the algorithmic result

is interesting as a contrast to our lower bounds, because the lower bounds do apply even in a model where random bits are free, and only actually computed input-dependent bits count towards the space complexity.

3 Preliminaries

Defining Adversarial Robustness. For the purposes of this paper, a “streaming algorithm” is always one-pass and we always think of it as working against an adversary. In the *standard streaming* setting, this adversary is oblivious to the algorithm’s actual run. This can be thought of as a special case of the setup we now introduce in order to define *adversarially robust streaming* algorithms.

Let \mathcal{U} be a universe whose elements are called tokens. A data stream is a sequence in \mathcal{U}^* . A data streaming problem is specified by a relation $f \subseteq \mathcal{U}^* \times \mathcal{Z}$ where \mathcal{Z} is some output domain: for each input stream $\sigma \in \mathcal{U}^*$, a valid solution is any $z \in \mathcal{Z}$ such that $(\sigma, z) \in f$. A randomized streaming algorithm \mathcal{A} for f running in s bits of space and using r random bits is formalized as a triple consisting of (i) a function $\text{INIT}: \{0, 1\}^r \rightarrow \{0, 1\}^s$, (ii) a function $\text{PROCESS}: \{0, 1\}^s \times \mathcal{U} \times \{0, 1\}^r \rightarrow \{0, 1\}^s$, and (iii) a function $\text{QUERY}: \{0, 1\}^s \times \{0, 1\}^r \rightarrow \mathcal{Z}$. Given an input stream $\sigma = (x_1, \dots, x_m)$ and a random string $R \in_R \{0, 1\}^r$, the algorithm starts in state $w_0 = \text{INIT}(R)$, goes through a sequence of states w_1, \dots, w_m , where $w_i = \text{PROCESS}(w_{i-1}, x_i, R)$, and provides an output $z = \text{QUERY}(w_m, R)$. The algorithm is δ -error in the standard sense if $\Pr_R[(\sigma, z) \in f] \geq 1 - \delta$.

To define adversarially robust streaming, we set up a game between two players: Solver, who runs an algorithm as above, and Adversary, who adaptively generates a stream $\sigma = (x_1, \dots, x_m)$ using a next-token function $\text{NEXT}: \mathcal{Z}^* \rightarrow \mathcal{U}$ as follows. With w_0, \dots, w_m as above, put $z_i = \text{QUERY}(w_i, R)$ and $x_i = \text{NEXT}(z_0, \dots, z_{i-1})$. In words, Adversary is able to query the algorithm at each point of time and can compute an arbitrary *deterministic* function of the history of outputs provided by the algorithm to generate his next token. Fix (an upper bound on) the stream length m . Algorithm \mathcal{A} is δ -error *adversarially robust* if

$$\forall \text{ function NEXT} : \Pr_R[\forall i \in [m] : ((x_1, \dots, x_i), z_i) \in f] \geq 1 - \delta.$$

In this work, we prove lower bounds for algorithms that are only required to be $O(1)$ -error adversarially robust. On the other hand, the algorithms we design will achieve vanishingly small error of the form $1/\text{poly}(m)$ and moreover, they will be able to detect when they are about to err and can abort at that point.

Graph Streams and the Coloring Problem. Throughout this paper, an *insert-only graph stream* describes an undirected graph on the vertex set $[n]$, for some fixed n that is known in advance, by listing its edges in some order: each token is an edge. A *strict graph turnstile stream* describes an evolving graph G by using two types of tokens – $\text{INS-EDGE}(\{u, v\})$, which causes $\{u, v\}$ to be added to G , and $\text{DEL-EDGE}(\{u, v\})$, which causes $\{u, v\}$ to be removed – and satisfies the promises that each insertion is of an edge that was not already in G and that each deletion is of an edge that was in G . When we use the term “graph stream” without qualification, it should be understood to mean an insert-only graph stream, unless the context suggests that either flavor is acceptable.

In this context, a semi-streaming algorithm is one that runs in $\tilde{O}(n) := O(n \text{ polylog } n)$ bits of space.

In the K -coloring problem, the input is a graph stream and a valid answer to a query is a vector in $[K]^n$ specifying a color for each vertex such that no two adjacent vertices receive the same color. The quantity K may be given as a function of some graph parameter, such

37:12 Adversarially Robust Coloring for Graph Streams

as the maximum degree Δ . In reading the results in this paper, it will be helpful to think of Δ as a growing but sublinear function of n , such as n^α for $0 < \alpha < 1$. Since an output of the K -coloring problem is a $\Theta(n \log K)$ -sized object, we think of a semi-streaming coloring algorithm running in $\tilde{O}(n)$ space as having “essentially optimal” space usage.

One-Way Communication Complexity. In this work, we shall only consider a special kind of two-player communication game: one where all input belongs to the speaking player Alice and her goal is to induce Bob to produce a suitable output. Such a game, g , is given by a relation $g \in \mathcal{X} \times \mathcal{Z}$, where \mathcal{X} is the input domain and \mathcal{Z} is the output domain. In a protocol Π for g , Alice and Bob share a random string R . Alice is given $x \in \mathcal{X}$ and sends Bob a message $\text{msg}(x, R)$. Bob uses this to compute an output $z = \text{out}(\text{msg}(x, R))$. We say that Π solves g to error δ if $\forall x \in \mathcal{X} : \Pr_R[(x, z) \in g] \geq 1 - \delta$. The communication cost of Π is $\text{cost}(\Pi) := \max_{x, R} \text{length}(\text{msg}(x, R))$. The (one-way, randomized, public-coin) δ -error communication complexity of g is $R_\delta^\rightarrow(g) := \min\{\text{cost}(\Pi) : \Pi \text{ solves } g \text{ to error } \delta\}$.

If Π never uses R , it is deterministic. Minimizing over zero-error deterministic protocols gives us the one-way deterministic communication complexity of g , denoted $D^\rightarrow(g)$.

A Result on Random Graphs. During the proof of our main lower bound (in Section 4.2), we shall need the following basic lemma on the maximum degree of a random graph. A proof of this lemma is given in the full version of this paper.

► **Lemma 3.** *Let G be a graph with M edges and n vertices, drawn uniformly at random. Define Δ_G to be its maximum degree. Then for $0 \leq \varepsilon \leq 1$:*

$$\Pr \left[\Delta_G \geq \frac{2M}{n}(1 + \varepsilon) \right] \leq 2n \exp \left(-\frac{\varepsilon^2}{3} \cdot \frac{2M}{n} \right).$$

Algorithmic Results From Prior Work. Our adversarially robust graph coloring algorithms in Section 5.2 will use, as subroutines, some previously known standard streaming algorithms for coloring. We summarize the key properties of these existing algorithms.

► **Fact 4** (Restatement of [7], Result 2). *There is a randomized turnstile streaming algorithm for $(\Delta + 1)$ -coloring a graph with max-degree Δ in the oblivious adversary setting that uses $\tilde{O}(n)$ bits of space and $\tilde{O}(n)$ random bits. The failure probability can be made at most $1/n^p$ for any large constant p .*

In the adversarial model described above, we need to answer a query after each stream update. The algorithm mentioned in Fact 4 or other known algorithms using “about” Δ colors (e.g., [12]) use at least $\tilde{\Theta}(n)$ post-processing time in the worst case to answer a query. Hence, using such algorithms in the adaptive adversary setting might be inefficient. We observe, however, that at least for insert-only streams, there exists an algorithm that is efficient in terms of both space and time. This is obtained by combining the algorithms of [12] and [22] (see the discussion towards the end of Section 5.2 for details).

► **Fact 5.** *In the oblivious adversary setting, there is a randomized streaming algorithm that receives a stream of edge insertions of a graph with max-degree Δ and degeneracy κ and maintains a proper coloring of the graph using $\kappa(1 + \varepsilon) \leq \Delta(1 + \varepsilon)$ colors, $\tilde{O}(\varepsilon^{-2}n)$ space, and $O(1)$ amortized update time. The failure probability can be made at most $1/n^p$ for any large constant p .*

4 Hardness of Adversarially Robust Graph Coloring

In this section, we prove our first major result, showing that graph coloring is significantly harder when working against an adaptive adversary than it is in the standard setting of an oblivious adversary. We carry out the proof plan outlined in Section 2.1, first describing and analyzing our novel communication game of SUBSET-AVOIDANCE (henceforth, AVOID) and then reducing the AVOID problem to robust coloring.

4.1 The Subset Avoidance Problem

Let $\text{AVOID}(t, a, b)$ denote the following one-way communication game.

- Alice is given $S \subseteq [t]$ with $|S| = a$;
- Bob must produce $T \subseteq [t]$ with $|T| = b$ for which T is disjoint from S .

Let $\text{AVOID}^k(t, a, b)$ be the problem of simultaneously solving k instances of $\text{AVOID}(t, a, b)$.

► **Lemma 6.** *The public-coin δ -error communication complexity of $\text{AVOID}^k(t, a, b)$ is bounded thus:*

$$R_{\delta}^{\rightarrow}(\text{AVOID}^k(t, a, b)) \geq \log(1 - \delta) + k \log \left(\binom{t}{a} / \binom{t-b}{a} \right) \quad (1)$$

$$\geq \log(1 - \delta) + kab/(t \ln 2). \quad (2)$$

Proof. Let Π be a δ -error protocol for $\text{AVOID}^k(t, a, b)$ and let $d = \text{cost}(\Pi)$, as defined in Section 3. Since, for each input $(S_1, \dots, S_k) \in \binom{[t]}{a}^k$, the error probability of Π on that input is at most δ , there must exist a fixing of the random coins of Π so that the resulting deterministic protocol Π' is correct on all inputs in a set

$$\mathcal{C} \subseteq \binom{[t]}{a}^k, \quad \text{with } |\mathcal{C}| \geq (1 - \delta) \binom{t}{a}^k.$$

The protocol Π' is equivalent to a function $\phi: \mathcal{C} \rightarrow \binom{[t]}{b}^k$ where

- the range size $|\text{Im}(\phi)| \leq 2^d$, because $\text{cost}(\Pi) \leq d$, and
- for each $(S_1, \dots, S_k) \in \mathcal{C}$, the tuple $(T_1, \dots, T_k) := \phi((S_1, \dots, S_k))$ is a correct output for Bob, i.e., $S_i \cap T_i = \emptyset$ for each i .

For any fixed $(T_1, \dots, T_k) \in \binom{[t]}{b}^k$, the set of all $(S_1, \dots, S_k) \in \binom{[t]}{a}^k$ for which each coordinate S_i is disjoint from the corresponding T_i is precisely the set $\binom{[t] \setminus T_1}{S_1} \times \dots \times \binom{[t] \setminus T_k}{S_k}$. The cardinality of this set is exactly $\binom{t-b}{a}^k$. Thus, for any subset \mathcal{D} of $\binom{[t]}{b}^k$, it holds that $|\mathcal{C} \cap \phi^{-1}(\mathcal{D})| \leq \binom{t-b}{a}^k |\mathcal{D}|$. Consequently,

$$(1 - \delta) \binom{t}{a}^k \leq |\mathcal{C}| = |\phi^{-1}(\text{Im}(\phi))| \leq \binom{t-b}{a}^k |\text{Im}(\phi)| \leq \binom{t-b}{a}^k 2^d,$$

which, on rearrangement, gives Equation (1).

To obtain Equation (2), we note that

$$\begin{aligned} \binom{t}{a} / \binom{t-b}{a} &= \frac{t!a!(t-a-b)!}{(t-a)!a!(t-b)!} = \frac{t \cdot (t-1) \cdots (t-a+1)}{(t-b) \cdot (t-b-1) \cdots (t-a-b+1)} \\ &\geq \left(\frac{t}{t-b} \right)^a = \left(\frac{1}{1-b/t} \right)^a > e^{ab/t}, \end{aligned} \quad (3)$$

which implies

$$\log(1 - \delta) + k \log \left(\binom{t}{a} / \binom{t-b}{a} \right) \geq \log(1 - \delta) + kab/(t \ln 2). \quad \blacktriangleleft$$

Since our data streaming lower bounds are based on the AVOID^k problem, it is important to verify that we are not analyzing its communication complexity too loosely. In the full version of this paper [18], we prove the following result, which says that the lower bound in Lemma 6 is close to being tight. In fact, a nearly matching upper bound can be obtained deterministically.

► **Lemma 7.** *For any $t \in \mathbb{N}$, $0 < a + b \leq t$, the deterministic complexity of $\text{AVOID}(t, a, b)$ is bounded thus:*

$$D^\rightarrow(\text{AVOID}(t, a, b)) \leq \log \left(\binom{t}{a} / \binom{t-b}{a} \right) + \log \left(\ln \binom{t}{a} \right) + 2.$$

4.2 Reducing Multiple Subset Avoidance to Graph Coloring

Having introduced and analyzed the AVOID communication game, we are now ready to prove our main lower bound result, on the hardness of adversarially robust graph coloring.

► **Theorem 8** (Main lower bound). *Let L, n, K be integers with $2K \leq n$, and $L + 1 \leq K$, and $L \geq 12 \ln(4n)$.*

Assume there is an adversarially robust coloring algorithm \mathcal{A} for insert-only streams of n -vertex graphs which works as long as the input graph has maximum degree $\leq L$, and maintains a coloring with $\leq K$ colors so that all colorings are correct with probability $\geq 1/4$. Then \mathcal{A} requires at least C bits of space, where

$$C \geq \frac{1}{40 \ln 2} \cdot \frac{nL^2}{K} - 3.$$

Proof. Given an algorithm \mathcal{A} as specified, we can construct a public-coin protocol to solve the communication problem $\text{AVOID}^{\lfloor n/(2K) \rfloor}(\binom{2K}{2}, \lfloor LK/4 \rfloor, \lfloor L/2 \rfloor \lceil K/2 \rceil)$ using exactly as much communication as \mathcal{A} requires storage space. The protocol for the more basic problem $\text{AVOID}(\binom{2K}{2}, \lfloor LK/4 \rfloor, \lfloor L/2 \rfloor \lceil K/2 \rceil)$ is described in Algorithm 1.

To use \mathcal{A} to solve $s := \lfloor n/2K \rfloor$ instances of AVOID , we pick s disjoint subsets V_1, \dots, V_s of the vertex set $[n]$, each of size $2K$. A streaming coloring algorithm on the vertex set $[2K]$ with degree limit L and using at most K colors can be implemented by relabeling the vertices in $[2K]$ to the vertices in some set V_i and using \mathcal{A} . This can be done s times in parallel, as the sets $(V_i)_{i=1}^s$ are disjoint. Note that a coloring of the entire graph on vertex set $[n]$ using $\leq K$ colors is also a K -coloring of the s subgraphs supported on V_1, \dots, V_s . To minimize the number of color queries made, Algorithm 1 can be implemented by alternating between adding elements from the matching M in each instance (for Algorithm 1), and making single color queries to the n -vertex graph (for Algorithm 1).

The guarantee that \mathcal{A} uses fewer than K colors depends on the input graph stream having maximum degree at most L . In Bob's part of the protocol, adding a matching to the graph only increases the maximum degree of the graph represented by Z by at most one; since he does this $\lfloor L/2 \rfloor$ times, in order for the maximum degree of the graph represented by Z to remain at most L , we would like the random graph Alice inserts into the algorithm to have maximum degree $\leq L/2 \leq L - \lfloor L/2 \rfloor$. By Lemma 3, the probability that, given some i , this random graph on V_i has maximum degree $\Delta_i \geq L/2$ is

$$\Pr \left[\Delta_i \geq \frac{L}{4}(1+1) \right] \leq 4K e^{-L/12}.$$

Taking a union bound over all s graphs, we find that

$$\Pr \left[\max_{i \in [s]} \Delta_i \geq L/2 \right] \leq 4K \left\lfloor \frac{n}{2K} \right\rfloor e^{-L/12} \leq 2n e^{-L/12}.$$

We can ensure that this happens with probability at most $1/2$ by requiring $L \geq 12 \ln(4n)$.

■ **Algorithm 1** Protocol for $\text{AVOID}(\binom{2K}{2}, \lfloor LK/4 \rfloor, \lfloor L/2 \rfloor \lceil K/2 \rceil)$.

Require: Algorithm \mathcal{A} that colors graphs up to maximum degree L , always using $\leq K$ colors

- 1: $R \leftarrow$ publicly random bits to be used by \mathcal{A}
- 2: $\pi \leftarrow$ publicly random permutation of $\{1, \dots, \binom{2K}{2}\}$, drawn uniformly
- 3: $e_1, \dots, e_{\binom{2K}{2}} \leftarrow$ an enumeration of the edges of the complete graph on $2K$ vertices

- 4: **function** ALICE(S):
- 5: $Z \leftarrow \mathcal{A}::\text{INIT}(R)$, the initial state of \mathcal{A}
- 6: **for** i from 1 to $\binom{2K}{2}$ **do**
- 7: **if** $\pi_i \in S$ **then**
- 8: $Z \leftarrow \mathcal{A}::\text{INSERT}(Z, R, e_i)$
- 9: **return** Z

- 10: **function** BOB(Z):
- 11: $J \leftarrow$ empty list
- 12: **for** i from 1 to $\lfloor L/2 \rfloor$ **do**
- 13: $\text{CLR} \leftarrow \mathcal{A}::\text{QUERY}(Z, R)$
- 14: $M \leftarrow$ maximal pairing of like-colored vertices, according to CLR
- 15: **for** each pair $\{u, v\} \in M$ **do**
- 16: $Z \leftarrow \mathcal{A}::\text{INSERT}(Z, R, \{u, v\})$ $\triangleright M$ is turned into a matching and inserted
- 17: $J \leftarrow J \cup M$
- 18: **if** $\text{length}(J) \leq \lfloor L/2 \rfloor \lceil K/2 \rceil$ **then**
- 19: **return** fail
- 20: **else**
- 21: $T \leftarrow \{\pi_i : e_i \in \text{first } \lfloor L/2 \rfloor \lceil K/2 \rceil \text{ edges of } J\}$
- 22: **return** T

If all the random graphs produced by Alice have maximum degree $\leq L/2$, and the $\lfloor L/2 \rfloor$ colorings requested by the protocol are all correct, then we will show that Bob's part of the protocol recovers at least $\lfloor L/2 \rfloor \lceil K/2 \rceil$ edges for each instance. Since the algorithm \mathcal{A} 's random bits R and permutation random bits π are independent, the probability that the maximum degree is low and the algorithm gives correct colorings on graphs of maximum degree at most L is $\geq (1/2) \cdot (1/4) = 1/8$.

The list of edges that Bob inserts (Algorithm 1) are fixed functions of the query output of \mathcal{A} on its state Z and random bits R . None of the edges can already have been inserted by Alice or Bob, since each edge connects two vertices which have the same color. Because these edges only depend on the query output of \mathcal{A} , conditioned on this query output they are independent of Z and R . This ensures that \mathcal{A} 's correctness guarantee against an adversary applies here, and thus the colorings reported on Algorithm 1 are correct.

Assuming all queries succeed, and the initial graph that Alice added has maximum degree $\leq L/2$, for each $i \in [\lfloor L/2 \rfloor]$, the coloring produced will have at most K colors. Let B be the set of vertices covered by the matching M , so that $[2K] \setminus B$ are the unmatched vertices. Since no pair of unmatched vertices can have the same color, $|[2K] \setminus B| \leq K$. This implies $|B| \geq K$, and since $|M| = |B|/2$ is an integer, we have $|M| \geq \lceil K/2 \rceil$. Thus each **for** loop iteration will add at least $\lceil K/2 \rceil$ new edges to J . The final value of the list J will contain at least $\lfloor L/2 \rfloor \lceil K/2 \rceil$ edges that were not added by Alice; Algorithm 1 converts the first $\lfloor L/2 \rfloor \lceil K/2 \rceil$ of these to elements of $\{1, \dots, \binom{2K}{2}\}$ not in the set S given to Alice.

37:16 Adversarially Robust Coloring for Graph Streams

Finally, by applying Lemma 6, we find that the communication C needed to solve s independent copies of $\text{AVOID}(\binom{2K}{2}, \lfloor LK/4 \rfloor, \lfloor L/2 \rfloor \lceil K/2 \rceil)$ with failure probability $\leq 7/8$ satisfies

$$\begin{aligned} C &\geq \log \left(1 - \frac{7}{8} \right) + \left\lfloor \frac{n}{2K} \right\rfloor \frac{\lfloor LK/4 \rfloor \cdot \lfloor L/2 \rfloor \lceil K/2 \rceil}{\binom{2K}{2} \ln 2} \\ &\geq \frac{n}{4K} \frac{L^2 K^2 / 20}{\frac{1}{2} (2K)^2 \ln 2} - 3 \geq \frac{nL^2}{40K \ln 2} - 3, \end{aligned}$$

where we used $K > L \geq 12 \ln(4n) \geq 12 \ln 4$ to conclude $\lfloor LK/4 \rfloor \lfloor L/2 \rfloor \lceil K/2 \rceil \geq (LK)^2/20$. \blacktriangleleft

Applying the above Theorem 8 with “ $K = f(L)$,” we immediately obtain the following corollary, which highlights certain parameter settings that are particularly instructive.

► **Corollary 9.** *Let f be a monotonically increasing function, and L an integer for which $L = \Omega(\log n)$ and $f(L) \leq n/2$. Let \mathcal{A} be a coloring algorithm which works for graphs of maximum degree up to L ; which at any point in time uses $\leq f(\Delta)$ colors, where Δ is the current graph’s maximum degree; and which has total failure probability $\leq 3/4$ against an adaptive adversary. Then the number of bits S of space used by \mathcal{A} is lower-bounded as $S = \Omega(nL^2/f(L))$. In particular:*

- If $f(\Delta) = \Delta + 1$ – or, more generally, $f(\Delta) = O(\Delta)$ – then $S = \Omega(nL)$ space is needed.
- To ensure $S = \tilde{O}(n)$ space, $f(\Delta) = \tilde{\Omega}(\Delta^2)$ is needed.
- If $f(L) = \Theta(n)$, then $S = \Omega(L^2)$.

5 Upper Bounds: Adversarially Robust Coloring Algorithms

We now turn to positive results. We show how to maintain a $\text{poly}(\Delta)$ -coloring of a graph in an adversarially robust fashion. We design two broad classes of algorithms. The first, described in Section 5.1, is based on palette sparsification as in [7, 4], with suitable enhancements to ensure robustness. The resulting algorithm maintains an $O(\Delta^3)$ -coloring and uses $\tilde{O}(n)$ bits of working memory. As noted in Section 2.2, the algorithm comes with the caveat that it requires a large pool of random bits: up to $\tilde{O}(n\Delta)$ of them. As also noted there, it makes sense to treat this randomness cost as separate from the space cost.

The second class of algorithms, described in Section 5.2, is built on top of the sketch switching technique of [10], suitably modified to handle non-real-valued outputs. This time, the amount of randomness used is small enough that we can afford to store all random bits in working memory. These algorithms can be enhanced to handle strict graph turnstile streams as described in Section 3. For any such turnstile stream of length at most m , we maintain an $O(\Delta^2)$ -coloring using $\tilde{O}(\sqrt{nm})$ space. More generally, we maintain an $O(\Delta^k)$ -coloring in $O(n^{1-1/k}m^{1/k})$ space for any $k \in \mathbb{N}$. In particular, for insert-only streams, this implies an $O(\Delta^k)$ -coloring in $O(n\Delta^{1/k})$ space.

5.1 An Algorithm Based on Palette Sparsification

We proceed to describe our palette-sparsification-based algorithm. It maintains a $3\Delta^3$ -coloring of the input graph G , where Δ is the evolving maximum degree of the input graph G . With high probability, it will store only $O(n(\log n)^4) = \tilde{O}(n)$ bits of information about G ; an easy modification ensures that this bound is always maintained by having the algorithm abort if it is about to overshoot the bound.

■ **Algorithm 2** Adversarially robust $3\Delta^3$ -coloring algorithm, assuming $0 < \Delta \leq L$.

Input: Stream of edges of a graph $G = (V, E)$, with maximum degree always $\leq L$

Random bits:

- 1: **for** each vertex $x \in [n]$ **do**
- 2: **for** each $i \in [L]$ **do**
- 3: $P_x^i \leftarrow$ list of $4 \log n$ colors sampled u.a.r. with replacement from $[2i^2]$

Initialize:

- 4: **for** each vertex $x \in [n]$ **do**
- 5: $\text{DEG}(x) \leftarrow 0$ ▷ tracks degree of x
- 6: $\text{CLR}(x) \leftarrow (0, 0)$ ▷ maintains color of x ; in general $\in \bigcup_{i=1}^L \{i\} \times [2i^2]$
- 7: $A \leftarrow$ empty list of edges

Process(edge $\{u, v\}$):

- 8: $\text{DEG}(u), \text{DEG}(v) \leftarrow \text{DEG}(u) + 1, \text{DEG}(v) + 1$ ▷ maintain vertex degrees
- 9: $k \leftarrow \max\{\text{DEG}(u), \text{DEG}(v)\}$
- 10: **for** i from k to L **do** ▷ store edges that might be needed in the future
- 11: **if** P_u^i and P_v^i overlap **then**
- 12: $A \leftarrow A \cup \{u, v\}$
- 13: $\text{USED} \leftarrow \{\text{CLR}(w) : \{u, w\} \in A\}$ ▷ prepare to recolor vertex u : collect colors of neighbors
- 14: **for** j from 1 to $4 \log n$ **do**
- 15: $c \leftarrow (\text{DEG}(u), P_u^{\text{DEG}(u)}[j])$ ▷ try the next color in the random list
- 16: **if** $c \notin \text{USED}$ **then**
- 17: $\text{CLR}(u) \leftarrow c$; **return**
- 18: **abort** ▷ failed to find a color

Query(\cdot):

- 19: **return** the vector CLR

The algorithm does need a large number of random bits – up to $O(nL(\log n)^2)$ of them – where L is the maximum degree of the graph at the end of the stream or an upper bound on the same. Due to the way the algorithm looks ahead at future random bits, L must be known in advance.

The algorithm uses these available random bits to pick, for each vertex, L lists of random color palettes, one at each of L “levels.” The level- i list at vertex x is called P_x^i and consists of $4 \log n$ colors picked uniformly at random with replacement from the set $[2i^2]$. The algorithm tracks each vertex’s degree. Whenever a vertex x is recolored, its new color is always of the form (d, p) , where $d = \deg(x)$ and $p \in P_x^d$. Thus, when the maximum degree in G is Δ , the only colors that have been used are the initial default $(0, 0)$ and colors from $\bigcup_{i=1}^{\Delta} \{i\} \times [2i^2]$. The total number of colors is therefore at most $1 + \sum_{i=1}^{\Delta} 2i^2 \leq 3\Delta^3$.

In the full version of this paper [18], we prove the following result.

► **Theorem 10.** *Algorithm 2 is an adversarially robust $O(\Delta^3)$ -coloring algorithm for insertion streams which stores $O(n(\log n)^4)$ bits related to the graph, requires access to $\tilde{O}(nL)$ random bits, and even against an adaptive adversary succeeds with probability $\geq 1 - O(1/n^2)$.*

5.2 Sketch-Switching Based Algorithms for Turnstile Streams

We present a class of sketch switching based algorithms for $\text{poly}(\Delta)$ -coloring. First, we give an outline of a simple algorithm for insert-only streams that colors the graph using $O(\Delta^2)$ colors and $\tilde{O}(n\sqrt{\Delta})$ space, where Δ is the max-degree of the graph at the time of query. Next, we show how to modify it to handle deletions. This is captured in Algorithm 3, whose correctness is given by Lemma 13. Then we describe how it can be generalized to get an $O(\Delta^k)$ -coloring in $\tilde{O}(n\Delta^{1/k})$ space for insert-only streams for any constant $k \in \mathbb{N}$. Finally, we prove the fully general result giving an $O(\Delta^k)$ -coloring in $\tilde{O}(n^{1-1/k}m^{1/k})$ space for turnstile streams, which is given by Theorem 14. Finally, we discuss how we can get rid of some assumptions, albeit reasonable, that we make for our algorithms.

Throughout this section, we make the standard assumption that the stream length m for turnstile streams is bounded by $\text{poly}(n)$. When we say that a statement holds with high probability (w.h.p.), we mean that it holds with probability at least $1 - 1/\text{poly}(n)$. In our algorithms, we often take the *product* of colorings of multiple subgraphs of a graph G . We define this notion below and record its key property.

► **Definition 11** (Product of Colorings). *Let $G_1 = (V, E_1), \dots, G_k = (V, E_k)$ be graphs on a common vertex set V . Given a coloring χ_i of G_i , for each $i \in [k]$, the product of these colorings is defined to be a coloring where each vertex $v \in V$ is assigned the color $\langle \chi_1(v), \chi_2(v), \dots, \chi_k(v) \rangle$.*

► **Lemma 12.** *Given a proper c_i -coloring χ_i of a graph $G_i = (V, E_i)$ for each $i \in [k]$, the product of the colorings χ_i is a proper $(\prod_{i=1}^k c_i)$ -coloring of $\cup_{i=1}^k G_i := (V, \cup_{i=1}^k E_i)$.*

Proof. An edge in $\cup_{i=1}^k G_i$ comes from G_{i^*} for some $i^* \in [k]$, and hence the colors of its endpoints in the product coloring differ in the i^* th coordinate. For $i \in [k]$, the i th coordinate can take c_i different values and hence the color bound holds. ◀

Insert-Only Streams and $O(\Delta^2)$ -Coloring. Split the $O(n\Delta)$ -length stream into $\sqrt{\Delta}$ chunks of size $O(n\sqrt{\Delta})$ each. Let \mathcal{A} be a standard (i.e., oblivious-adversary) semi-streaming algorithm for $O(\Delta)$ -coloring a graph (by Fact 4 and Fact 5, such algorithms exist). At the start of the stream, initialize $\sqrt{\Delta}$ parallel copies of \mathcal{A} , called $A_1, \dots, A_{\sqrt{\Delta}}$; these will be our “parallel sketches.” At any point of time, only a suffix of this list of parallel sketches will be active.

We use the sketch switching idea of [10] as follows. With each edge insertion, we update each of the active parallel sketches. Whenever we arrive at the end of a chunk, we say we have reached a “checkpoint” and query the least-numbered active sketch (this is guaranteed to be “fresh” in the sense that it has not been queried before) to produce a coloring of the entire graph until that point. By design, the randomness of the queried sketch is independent of the edges it has processed. Therefore, it returns a correct $O(\Delta)$ -coloring of the graph until that point, w.h.p. Henceforth, we mark the just-queried sketch as inactive and never update it, but continue to update all higher-numbered sketches. Thus, each copy of \mathcal{A} actually processes a stream independent of its randomness and hence, works correctly while using $\tilde{O}(n)$ space. By a union bound over all sketches, w.h.p., all of them generate correct colorings at the respective checkpoints and simultaneously use $\tilde{O}(n)$ space each, i.e., $\tilde{O}(n\sqrt{\Delta})$ space in total.

Conditioned on the above good event, we can always return an $O(\Delta^2)$ -coloring as follows. We store (buffer) the most recent partial chunk explicitly, using our available $\tilde{O}(n\sqrt{\Delta})$ space. Now, when a query arrives, we can express the current graph G as $G_1 \cup G'$, where G_1 is the subgraph of G until the last checkpoint and G' is the subgraph in our buffer. Observe

that we computed an $O(\Delta(G_1))$ -coloring of G_1 at the last checkpoint. Further, we can deterministically compute a $(\Delta(G') + 1)$ -coloring of G' since we explicitly store it. We output the product of the colorings (Definition 11) of G_1 and G' , which must be a proper $O(\Delta(G_1) \cdot \Delta(G')) = O(\Delta(G)^2)$ -coloring of the graph G (Lemma 12).

Extension to Handle Deletions. The algorithm above doesn't immediately work for turnstile streams. The chunk currently being processed by the algorithm may contain an update that deletes an edge which was inserted before the start of the chunk, and hence, is not in store. Thus, we can no longer express the current graph as a union of the graphs G_1 and G' as above. Overcoming this difficulty complicates the algorithm enough that it is useful to lay it out more formally as pseudocode (see Algorithm 3). This new algorithm maintains an $O(\Delta^2)$ -coloring, works even on turnstile streams, and uses $\tilde{O}(\sqrt{mn})$ space. Note that while the black-box algorithm \mathcal{A} used in Algorithm 3 might be any generic $O(\Delta)$ -coloring semi-streaming algorithm with error $1/m$, it can be, for instance, chosen to be the one given by Fact 4 or, for insert-only streams, the one in Fact 5. The former gives a tight $(\Delta + 1)$ -coloring but possibly large query time, while the latter answers queries fast using possibly a few more colors, up to $\Delta(1 + \epsilon)$.

Before proceeding to the analysis, let us set up some terminology. Recall from Section 3 that we work with strict graph turnstile streams, so each deletion of an edge e can be matched to a unique previous token that most recently inserted e . An edge deletion, where the corresponding insertion did not occur inside the same chunk, is called a *negative edge*. Call a point in the stream a *checkpoint* if we use a *fresh* parallel copy of \mathcal{A} , i.e., a copy A_i that hasn't been queried before, to generate an $O(\Delta)$ -coloring of the graph at that point. We define two types of checkpoints, namely *fixed* and *ad hoc*. We have a fixed checkpoint at the end of each chunk; this means that whenever the last update of a chunk arrives, we compute a coloring of the graph seen so far using a fresh copy of \mathcal{A} . The ad hoc checkpoints are made on the fly inside a current chunk, precisely when a query appears and we see that the max-degree of the current graph is less than half of what it was at the *last* checkpoint (which might be fixed or ad hoc).

A detailed analysis of Algorithm 3, given in the full version of this paper [18], proves the following lemma.

► **Lemma 13.** *For any strict graph turnstile stream of length at most m for a graph G given by an adaptive adversary, the following hold simultaneously, w.h.p.:*

- *Algorithm 3 outputs an $O(\Delta^2)$ -coloring after each query, where Δ is the maximum degree of the graph at the time a query is made.*
- *Algorithm 3 uses $\tilde{O}(\sqrt{mn})$ bits of space.*

Generalization to $O(\Delta^k)$ -Coloring in $\tilde{O}(n\Delta^{1/k})$ Space for Insert-Only Streams. The recursive application of the sketch-switching techniques explained in the previous section is described in detail in the full version of this paper [18]. There, we prove the following result, and its immediate corollary.

► **Theorem 14.** *For any strict graph turnstile stream of length at most m , and for any constant $k \in \mathbb{N}$, there exists an adversarially robust algorithm \mathcal{A} such that the following hold simultaneously w.h.p.:*

- *After each query, \mathcal{A} outputs an $O(\Delta^k)$ -coloring, where Δ is the max-degree of the current graph.*
- *\mathcal{A} uses $\tilde{O}(n^{1-1/k}m^{1/k})$ bits of space.*

37:20 Adversarially Robust Coloring for Graph Streams

■ **Algorithm 3** Adversarially robust $O(\Delta^2)$ -coloring in $\tilde{O}(\sqrt{nm})$ space for turnstile streams.

Input: Stream of edge insertions/deletions of an n -vertex graph $G = (V, E)$; parameter m

Require: Semi-streaming algorithm \mathcal{A} that works on turnstile graph streams and provides an $O(\Delta)$ -coloring with error $\leq 1/m$ against an oblivious adversary

Initialize:

- 1: $s \leftarrow C \cdot \sqrt{m/n} \log n$ for some sufficiently large constant C
- 2: $A_1, \dots, A_s \leftarrow$ independent parallel initializations of \mathcal{A}
- 3: $c \leftarrow 0$ ▷ index into list (A_1, \dots, A_s)
- 4: $\text{CLR} \leftarrow n$ -vector of vertex colors, initialized to all-1s ▷ valid $O(\Delta)$ -coloring until last checkpoint
- 5: $\text{DEG} \leftarrow n$ -vector of vertex degrees, initialized to all-0s
- 6: $G' \leftarrow (V, \emptyset)$ ▷ buffer to store current chunk
- 7: $\text{CHUNKSIZE} \leftarrow 0$ ▷ current buffer size
- 8: $\text{CHECKPTMAXDEG} \leftarrow 0$ ▷ max-degree at last checkpoint

Process(operation OP, edge $\{u, v\}$): ▷ OP says whether to insert or delete

- 9: **for** i from $c + 1$ to s **do**
- 10: A_i . **Process**(OP, $\{u, v\}$) ▷ if this aborts, report FAIL
- 11: **if** OP = “insert” **then**
- 12: increment $\text{DEG}(u), \text{DEG}(v)$
- 13: add $\{u, v\}$ to G'
- 14: **else if** OP = “delete” **then**
- 15: decrement $\text{DEG}(u), \text{DEG}(v)$
- 16: **if** $\{u, v\} \in G'$ **then:** ▷ else, negative edge; not stored
- 17: delete $\{u, v\}$ from G'
- 18: $\text{CHUNKSIZE} \leftarrow \text{CHUNKSIZE} + 1$
- 19: $\Delta \leftarrow \max_{v \in [n]} \text{DEG}(v)$
- 20: **if** $\text{CHUNKSIZE} = \sqrt{nm}$ **then:**
- 21: **NewCheckpoint**() ▷ fixed checkpoint encountered
- 22: $\text{CHUNKSIZE} \leftarrow 0$
- 23: **if** $\Delta < \text{CHECKPTMAXDEG}/2$ **then:**
- 24: **NewCheckpoint**() ▷ ad hoc checkpoint created

Query():

- 25: $\text{CLR}' \leftarrow (\Delta_{G'} + 1)$ -coloring of G'
- 26: **return** $\langle (\text{CLR}(v), \text{CLR}'(v)) : v \in [n] \rangle$ ▷ take the product of the two colorings

NewCheckpoint():

- 27: $c \leftarrow c + 1$ ▷ switch to next fresh sketch
- 28: $\text{CLR} \leftarrow A_c$. **Query**() ▷ if A_c fails, report FAIL
- 29: $G' \leftarrow (V, \emptyset)$
- 30: $\text{CHECKPTMAXDEG} \leftarrow \max_{v \in [n]} \text{DEG}(v)$

► **Corollary 15.** *For any stream of edge insertions describing a graph G , and for any $k \in \mathbb{N}$, there exists an adversarially robust algorithm \mathcal{A} such that the following hold simultaneously w.h.p.:*

- *After each query, \mathcal{A} outputs an $O(\Delta^k)$ -coloring, where Δ is the max-degree of the current graph.*
- *\mathcal{A} uses $\tilde{O}(n\Delta^{1/k})$ bits of space.*

Removing the Assumption of Prior Knowledge of m . In Algorithm 3 as well as in the algorithm described in Theorem 14, we assume that a value m , an upper bound on the number of edges, is given to us in advance. Without it, we do not know how many sketches to initialize at the start of the stream. In this context, we note the following. First, knowledge of an upper bound on the number of edges is a reasonable assumption, especially for turnstile streams, since an algorithm typically knows how large of an input stream it can handle. Second, for insert-only streams, we can set $m = n\Delta/2$ if an upper bound Δ on the max-degree of the final graph is known; a knowledge of such a bound is reasonable since $f(\Delta)$ -coloring is usually studied with a focus on bounded-degree graphs. Third, we can remove the assumption of knowing either m or Δ for insert-only streams at the cost of a factor of Δ in the number of colors and an additive $\tilde{O}(n)$ factor in space, which we outline next.

At the beginning of the stream, we initialize $\lceil \log n \rceil$ copies of the oblivious $O(\Delta)$ -coloring semi-streaming algorithm A for the checkpoints where Δ first attains values of the form 2^i for some $i \in [\lceil \log n \rceil]$. For each i , the substream between the checkpoints with $\Delta = 2^i$ and $\Delta = 2^{i+1}$ can be handled using our algorithm as a black box since we know that the stream length is at most $2^i n$. This way, we need not initialize $O(D^{1/k})$ sketches for $D \gg \Delta_{\max}$ at the very beginning of the stream, where Δ_{\max} is the final max-degree of the graph, and incur such a huge factor in space; we can initialize $d^{1/k}$ sketches for the substream with $d \leq \Delta \leq 2d$ only when (if at all) Δ reaches the value d . Thus, the maximum space used is $O(n\Delta_{\max}^{1/k})$, which we can afford. When queried in a substream between checkpoints at $\Delta = 2^i$ and $\Delta = 2^{i+1}$, we use our $O(\Delta^k)$ -coloring algorithm to get a coloring of the substream, and take a product with the $O(\Delta)$ -coloring at the checkpoint at $\Delta = 2^i$. Thus, we get an $O(\Delta^{k+1})$ -coloring of the current graph. The additional space usage is $\tilde{O}(n)$ due to the initial $\lceil \log n \rceil$ sketches taking up $\tilde{O}(n)$ space each; hence, the total space usage is still $O(n\Delta_{\max}^{1/k})$.

References

- 1 Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Ami Paz. Smaller cuts, higher lower bounds. *CoRR*, abs/1901.01630, 2019.
- 2 Kook Jin Ahn and Sudipto Guha. Graph sparsification in the semi-streaming model. In *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part II*, volume 5556 of *Lecture Notes in Computer Science*, pages 328–338. Springer, 2009. doi:10.1007/978-3-642-02930-1_27.
- 3 Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. In *Proc. 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 459–467, 2012. doi:10.1137/1.9781611973099.40.
- 4 Noga Alon and Sepehr Assadi. Palette sparsification beyond $(\Delta+1)$ vertex coloring. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2020, August 17-19, 2020, Virtual Conference*, volume 176 of *LIPICs*, pages 6:1–6:22, 2020. doi:10.4230/LIPICs.APPROX/RANDOM.2020.6.
- 5 Sepehr Assadi. Sublinear algorithms for $(\Delta + 1)$ vertex coloring. Lecture at Sublinear Algorithms and Nearest-Neighbor Search Workshop, Simons Institute; available online at https://www.youtube.com/watch?v=VU7Y_8ZcNu0&t=2206, 2018.

- 6 Sepehr Assadi, Andrew Chen, and Glenn Sun. Deterministic graph coloring in the streaming model. *arXiv preprint arXiv:2109.14891*, 2021.
- 7 Sepehr Assadi, Yu Chen, and Sanjeev Khanna. Sublinear algorithms for $(\Delta + 1)$ vertex coloring. In *Proc. 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 767–786, 2019. doi:10.1137/1.9781611975482.48.
- 8 Idan Attias, Edith Cohen, Moshe Shechner, and Uri Stemmer. A framework for adversarial streaming via differential privacy and difference estimators. *CoRR*, abs/2107.14527, 2021. arXiv:2107.14527.
- 9 Omri Ben-Eliezer, Talya Eden, and Krzysztof Onak. Adversarially robust streaming via dense–sparse trade-offs. *CoRR*, abs/2109.03785, 2021.
- 10 Omri Ben-Eliezer, Rajesh Jayaram, David P. Woodruff, and Eylon Yogev. A framework for adversarially robust streaming algorithms. In *Proc. 39th ACM Symposium on Principles of Database Systems*, pages 63–80, 2020. doi:10.1145/3375395.3387658.
- 11 Omri Ben-Eliezer and Eylon Yogev. The adversarial robustness of sampling. In *Proc. 39th ACM Symposium on Principles of Database Systems*, pages 49–62. ACM, 2020. doi:10.1145/3375395.3387643.
- 12 Suman K. Bera, Amit Chakrabarti, and Prantar Ghosh. Graph coloring via degeneracy in streaming and other space-conscious models. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPICs*, pages 11:1–11:21, 2020. doi:10.4230/LIPICs.ICALP.2020.11.
- 13 Suman Kalyan Bera and Prantar Ghosh. Coloring in graph streams. *CoRR*, abs/1807.07640, 2018. arXiv:1807.07640.
- 14 Anup Bhattacharya, Arijit Bishnu, Gopinath Mishra, and Anannya Upasana. Even the easiest(?) graph coloring problem is not easy in streaming! In *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPICs*, pages 15:1–15:19, 2021. doi:10.4230/LIPICs.ITCS.2021.15.
- 15 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 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1–20. SIAM, 2018. doi:10.1137/1.9781611975031.1.
- 16 Sayan Bhattacharya, Fabrizio Grandoni, Janardhan Kulkarni, Quanquan C. Liu, and Shay Solomon. Fully dynamic $(\Delta+1)$ -coloring in constant update time. *CoRR*, abs/1910.02063, 2019.
- 17 Vladimir Braverman, Avinatan Hassidim, Yossi Matias, Mariano Schain, Sandeep Silwal, and Samson Zhou. Adversarial robustness of streaming algorithms through importance sampling. *CoRR*, abs/2106.14952, 2021. arXiv:2106.14952.
- 18 Amit Chakrabarti, Prantar Ghosh, and Manuel Stoeckl. Adversarially robust coloring for graph streams. *CoRR*, abs/2109.11130, 2021. arXiv:2106.14952.
- 19 Shafi Goldwasser, Ofer Grossman, Sidhant Mohanty, and David P. Woodruff. Pseudo-Deterministic Streaming. In *Proc. 20th Conference on Innovations in Theoretical Computer Science*, volume 151, pages 79:1–79:25, 2020. doi:10.4230/LIPICs.ITCS.2020.79.
- 20 Moritz Hardt and David P. Woodruff. How robust are linear sketches to adaptive inputs? In *Proc. 45th Annual ACM Symposium on the Theory of Computing*, pages 121–130, 2013. doi:10.1145/2488608.2488624.
- 21 Avinatan Hassidim, Haim Kaplan, Yishay Mansour, Yossi Matias, and Uri Stemmer. Adversarially robust streaming algorithms via differential privacy. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- 22 Monika Henzinger and Pan Peng. Constant-time dynamic $(\Delta+1)$ -coloring. In *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020*,

- Montpellier, France, volume 154 of *LIPICs*, pages 53:1–53:18, 2020. doi:10.4230/LIPICs.STACS.2020.53.
- 23 Hossein Jowhari, Mert Saglam, and Gábor Tardos. Tight bounds for l_p samplers, finding duplicates in streams, and related problems. In *Proc. 30th ACM Symposium on Principles of Database Systems*, pages 49–58, 2011. doi:10.1145/1989284.1989289.
 - 24 Haim Kaplan, Yishay Mansour, Kobbi Nissim, and Uri Stemmer. Separating adaptive streaming from oblivious streaming using the bounded storage model. In *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 94–121. Springer, 2021. doi:10.1007/978-3-030-84252-9_4.
 - 25 Andrew McGregor. Graph stream algorithms: a survey. *ACM SIGMOD Record*, 43(1):9–20, 2014. doi:10.1145/2627692.2627694.
 - 26 Ilya Mironov, Moni Naor, and Gil Segev. Sketching in adversarial environments. *SIAM J. Comput.*, 40(6):1845–1870, 2011. doi:10.1137/080733772.
 - 27 Noam Nisan. Pseudorandom generators for space-bounded computation. In *Proc. 22nd Annual ACM Symposium on the Theory of Computing*, pages 204–212, 1990. doi:10.1007/BF01305237.
 - 28 Uri Stemmer. Separating adaptive streaming from oblivious streaming. Lecture at STOC 2021 Workshop: Robust Streaming, Sketching and Sampling, available online at <https://www.youtube.com/watch?v=svgv-xw9DZc&t=7679s>, 2021. Based on joint work with Haim Kaplan, Yishay Mansour, and Kobbi Nissim.
 - 29 David P. Woodruff and Samson Zhou. Tight bounds for adversarially robust streams and sliding windows via difference estimators. In *Proc. 62nd Annual IEEE Symposium on Foundations of Computer Science*, page to appear, 2021.