

Jan Grebík* and Zoltán Vidnyánszky**

FROM DESCRIPTIVE TO DISTRIBUTED

Abstract. In the past couple of years a rich connection has been found between the fields of *descriptive set theory* and *distributed computing*. Frequently, and less surprisingly, finitary algorithms can be adopted to the infinite setting, resulting in theorems about infinite, definable graphs. In this survey, we take a different perspective and illustrate how results and ideas from descriptive set theory provide new insights and techniques to the theory of distributed computing. We focus on the two classical topics from graph theory, *vertex* and *edge colorings*. After summarizing the up-to-date results from both areas, we discuss the adaptation of Marks' games method to the LOCAL model of distributed computing and the development of the multi-step Vizing's chain technique, which led to the construction of the first non-trivial distributed algorithms for Vizing colorings. We provide a list of related open problems to complement our discussion. Finally, we describe an efficient deterministic distributed algorithm for Brooks coloring on graphs of subexponential growth.

Mathematics Subject Classification (2020): 03E15, 68W15; 28A05, 05C15

Keywords: descriptive, Borel combinatorics, LOCAL model

*Masaryk University, Faculty of Informatics, Brno, Czech Republic

`grebikj@gmail.com`

ORCID: 0000-0002-9980-4660

**Eötvös Loránd University, Institute of Mathematics, Budapest, Hungary

`zoltan.vidnyanszky@ttk.elte.hu`

ORCID: 0000-0001-8168-9353

DOI: <https://doi.org/10.64191/zr24040410108>

CONTENTS

| | |
|---|-----|
| 1. Introduction | 263 |
| 2. Classical results and notation | 265 |
| 2.1. Greedy algorithm | 266 |
| 2.2. Brooks' and Vizing's theorems | 266 |
| 3. LOCAL model | 268 |
| 3.1. Formal definition and basic results | 269 |
| 3.2. Global picture of the LOCAL model | 271 |
| 3.3. Upper bound for Brooks and Vizing colorings | 273 |
| 4. Measurable combinatorics | 274 |
| 4.1. Definitions and basic results | 275 |
| 4.2. Measure, as the analogue of randomness | 278 |
| 5. Summary of coloring results and related topics | 281 |
| 5.1. Decidability and complexity | 281 |
| 5.2. Formal connections | 284 |
| 6. Marks' games | 285 |
| 6.1. Games | 287 |
| 6.2. Homomorphism graphs | 289 |
| 6.3. Transfer principles | 290 |
| 6.4. The descriptive case | 293 |
| 7. Versions of Vizing's theorem | 295 |
| 7.1. Perfect matching on regular trees | 296 |
| 7.2. Multi-step Vizing chains | 298 |
| 7.3. Applications | 301 |
| 8. Brooks' theorem for subexponential growth graphs | 303 |
| 9. Open problems | 305 |
| Edge colorings. | 306 |
| Acknowledgments | 309 |
| References | 310 |

1. Introduction

Graph colorings is a fundamental concept in combinatorics and theoretical computer science. Moreover, a myriad of practical problems can be also reformulated in terms of graphs and their colorability (e.g., scheduling and clustering problems,

register allocation, etc.¹), thus, these concepts have numerous applications beyond their theoretical usefulness. A basic question about graph colorings asks for the smallest number of colors needed to color the vertices (or edges) of a given graph G so that adjacent vertices (or edges) have different colors. In the case of vertex coloring this number is called the *chromatic number* of G , denoted by $\chi(G)$, and in case of edge coloring it is called the *chromatic index* of G , denoted by $\chi'(G)$. Classical theorems of Brooks and Vizing give an upper bound on $\chi(G)$ and $\chi'(G)$, respectively, in terms of the maximum degree $\Delta(G)$ of the graph G . Colorings matching these upper bounds can be found in time polynomial in the size of the graph. In contrast, deciding (in the case of vertex colorings even approximately) the chromatic number or chromatic index of a graph is an NP-complete problem.

In this survey, we discuss algorithmic aspects of these classical problems from the perspectives of the *LOCAL model of distributed computing* and *measurable graph combinatorics*. At first glance, these fields may seem to be at the opposite ends of the spectrum: one focuses on distributed algorithms for large but finite networks, while the other investigates uncountable graphs, and stems from pure measure theory and logic. However, the recently discovered formal connections, originated in the work of Bernshteyn [12] and Elek [50], and the ensuing surge of activity proved otherwise. In fact, researchers in both areas have very often studied *exactly* the same problems, and, now, there are black-box translations of results from one field to the other. As we only briefly comment on this beautiful correspondence in Section 5.2, we refer the reader to the surveys of Bernshteyn [10] and Pikhurko [107] for more comprehensive treatment.

What phenomenon links these fields? At a high level, in both areas the combinatorial core of the questions of interest boils down to the understanding when a solution to local graph coloring problems can be produced *efficiently in a decentralized way*. In the seminal paper [12], Bernshteyn proved that efficient distributed algorithms solving a given coloring problem on finite graphs yield automatically the existence of such a measurable coloring on measurable graphs. As a consequence, many results from the distributed computing literature immediately yield new results in the relatively younger field of measurable graph combinatorics. While this fundamental result follows the typical direction, namely, ideas from theoretical computer science are used to solve major problems in pure mathematics, it has also opened the gates for myriads of follow up questions. The most important one for this survey might be the following:

Can measure theoretic techniques have a direct impact on the evolution of distributed algorithms on finite graphs?

In this paper, we aim to answer this question in the positive. Our goal is to complement the views presented in the survey papers by Bernshteyn [10] and Pikhurko [107] by describing the flow of ideas from the infinite world to the finite. The two main results from measurable graph combinatorics and their application in the theory of distributed computing that we discuss are (a) the *determinacy method*

¹ For more on the historical evolution of these problems, and in particular, the amusing history of the four coloring problem see [122] or https://en.wikipedia.org/wiki/Graph_coloring.

of Marks for showing the *inexistence* of Borel vertex colorings, and (b) *measurable Vizing's theorem*, which proves the *existence* of measurable edge colorings that satisfy the Vizing's bound.

The paper is structured as follows. Sections 2–5 serve as an introduction and a comparison of graph coloring results in the classical, distributed and measurable setting. We discuss the adaptation of Marks' technique to the LOCAL model of distributed computing in Section 6 and the measurable and distributed analogs of Vizing's theorem in Section 7. In Section 8, we describe the efficient distributed algorithms for Brooks' coloring for graphs of subexponential growth and, finally, in Section 9 we finish the survey with a collection of related and fascinating open problems.

2. Classical results and notation

We start by fixing a standard graph theoretic notation that will be used throughout the paper. A *graph* G is a pair (V, E) , where V is a set and $E \subseteq \binom{V}{2}$ is a set of unordered pairs of V , in particular, we do not allow multiple edges or loops. We will also use the notation $V(G)$ for the set of vertices of G and $E(G)$ or just simply G for the set of edges.

To distinguish finite and countable graphs from the usually uncountable graphs coming from the measurable context we use G for the former and \mathcal{G} for the latter. In case G is finite, we set $n := |V|$ and $m := |E|$. The *maximum degree* of G is denoted by $\Delta(G)$. Throughout the paper we assume that $\Delta(G) \in \mathbb{N}$, in fact, most of the time we treat $\Delta(G)$ as a constant that is fixed in advance and independent of n . The line graph of a graph $G = (V, E)$, denoted by $L(G)$ is defined as the pair (E, F) , where $\{e, e'\} \in F$ if $e \cap e' \neq \emptyset$. We write $N_G(v)$ for the set of neighbors of $v \in V$ in G . We set $\deg_G(v) = |N_G(v)|$ and $\deg_G(v, A) = |N_G(v) \cap A|$. We denote as $B_G(v, t)$ the t -neighborhood of v in G , that is, the restriction of G to vertices of graph distance at most t from v , that includes additional labeling that might be defined on G .

Standard conventions are used to denote asymptotics, that is, if $f, g: \mathbb{N} \rightarrow \mathbb{N}$ are functions, we say that $f \in o(g)$, $f \in O(g)$, $f \in \Omega(g)$ and $f \in \omega(g)$ if $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$, $\limsup_{n \rightarrow \infty} f(n)/g(n) < \infty$, $g \in O(f)$ and $g \in o(f)$, respectively. Additionally, we write $f \in \Theta(g)$ if $f \in O(g)$ and $g \in O(f)$.

For $k \in \mathbb{N}$, we write $[k] = \{1, \dots, k\}$. A *vertex coloring* of G (with $k \in \mathbb{N}$ colors) or simply *k -coloring* is a map $c: V \rightarrow [k]$ such that $c(v) \neq c(w)$ whenever $\{v, w\} \in E$. An *edge coloring* of G (with $\ell \in \mathbb{N}$ colors) is a map $d: E \rightarrow [\ell]$ such that $d(e) \neq d(f)$ whenever $e \neq f$ and $e \cap f \neq \emptyset$. The *chromatic number* of G , denoted by $\chi(G)$, is the minimum $k \in \mathbb{N}$ such that there is a vertex coloring of G with k colors and the *chromatic index* of G , denoted by $\chi'(G)$, is the minimum $\ell \in \mathbb{N}$ such that there is an edge coloring of G with ℓ colors. Note that we have $\chi(L(G)) = \chi'(G)$ by the definition.

A *partial vertex coloring* (with $k \in \mathbb{N}$ colors) is a partial map $c: V \rightarrow [k]$ that is a vertex coloring of the graph induced from G on $\text{dom}(c) \subseteq V$. We set $U_c = V \setminus \text{dom}(c)$ to denote the set of uncolored vertices. A *partial edge coloring* $d: E \rightarrow [k]$ and the set of uncolored edges U_d is defined analogously.

2.1. Greedy algorithm. We start with the classical greedy upper bound.

Theorem 2.1 (Greedy coloring). *Let G be a graph such that $\Delta(G) \in \mathbb{N}$. Then $\chi(G) \leq \Delta(G) + 1$ and $\chi'(G) \leq 2\Delta(G) - 1$.*

The proof of this theorem for vertex colorings is very often presented in the form of a sequential algorithm. Namely, after enumerating the vertices of a finite graph G as $(v_1, \dots, v_{|V|})$, we go through the list inductively and color the current vertex with any of the available colors at this step. This is always possible as we are allowed to use strictly more colors than is the maximum degree of G . The argument for $\chi'(G)$ is analogous after, for example, passing to $L(G)$.

Note that the sequential algorithm runs in time $O(m)$. In case that we treat Δ as a constant, which we usually do, the running time is $O(n)$. It is one of the basic results in the LOCAL model, which is discussed in the next section, that there is in fact an efficient distributed algorithm producing a greedy coloring, see Theorem 3.4.

2.2. Brooks' and Vizing's theorems. The greedy upper bound for vertex colorings is tight in general. Namely, if $n \in \mathbb{N}$ and K_n is the complete graph on $[n]$ vertices, then $\chi(K_n) = n$ and $\Delta(K_n) = n - 1$. A classical result of Brooks [28] states that apart from this situation and the case of odd cycles, $\Delta(G)$ colors are enough.

Theorem 2.2 (Brooks coloring). *Let G be a graph such that $\Delta(G) \in \mathbb{N}$ and assume that G does not contain a copy of $K_{\Delta(G)+1}$, or an odd cycle in case $\Delta(G) = 2$. Then $\chi(G) \leq \Delta(G)$.*

Brooks coloring can be found in time $O(m)$ as was proved by Skulrattanakulchai [113], however, and that is one of the main results discussed in this paper, there is no efficient deterministic distributed algorithm even if we assume that the graphs that we consider have high *girth*, i.e., they look locally like trees. Recall that a classical theorem of Erdős [53], proven by the probabilistic method, states that there exist graphs of arbitrarily high girth and chromatic number. Recall that the *girth* of a graph G is the size of the shortest cycle in G . More related to our investigation is the same result of Bollobás [19] for Δ -regular random graphs: random Δ regular graphs have girth tending to ∞ with positive probability as $n \rightarrow \infty$ and the chromatic number of such graphs is almost surely $\Theta(\Delta/\log(\Delta))$ as $\Delta \rightarrow \infty$. This gives a lower bound on the number of colors needed for *any* efficient distributed algorithm on locally tree-like graphs.

In the case of edge colorings the situation is much simpler as was proved by Vizing [120] and independently by Gupta [77].

Theorem 2.3 (Vizing coloring). *Let G be a graph such that $\Delta(G) \in \mathbb{N}$. Then $\chi'(G) \in \{\Delta(G), \Delta(G) + 1\}$.*

The standard proof of Vizing's theorem, which is based on the augmenting chain technique, implies immediately that edge coloring with $\Delta(G)+1$ colors can be found in time polynomial in m . Currently, the fastest deterministic algorithm for Vizing

coloring for general graphs due to Sinnamoni [112] runs in time $O(m\sqrt{n})$, and the fastest randomized one due to Bhattacharya, Carmon, Costa, Solomon, and Zhang [18] runs in time $O(mn^{1/3})$. We refer the reader to [14, 15, 18] for more details on the current development and better bounds under some additional assumptions on the maximum degree. While it is known that there is no efficient deterministic distributed algorithm, the existence of a randomized one is one of the main open problems in the theory of distributed computing, see Section 7 and Problem 9.1.

We formalize the *augmenting chain technique* as it plays a crucial role both in our investigation of edge colorings in Section 7 and in our construction of an efficient deterministic distributed algorithm for Brooks' theorem on graphs of subexponential growth in Section 8.

Definition 2.4 (Augmenting subgraph). Let $G = (V, E)$ be a graph, $d: E \rightarrow [k]$ be a partial edge coloring and $e \in U_d$ be an uncolored edge. A subgraph $H \subseteq G$ is called *augmenting* for e (and d) if it is connected, $E(H) \setminus \{e\} \subseteq \text{dom}(d)$ and there is a partial edge coloring $d': E \rightarrow [k]$ such that $\text{dom}(d') = \text{dom}(d) \cup \{e\}$ and $d(f) = d'(f)$ for every $f \in \text{dom}(d) \setminus E(H)$.

An augmenting subgraph for partial vertex colorings and an uncolored vertex is defined analogously.

Vizing's theorem for finite graphs then follows from the fact that there is an augmenting subgraph for every uncolored edge $e \in E$ and any partial edge coloring of G with $\Delta(G) + 1$ colors.

In order to produce Vizing coloring by a distributed algorithm, or measurably, we need to control the size and structure of the augmenting subgraphs. Already the standard proof of Vizing's theorem provides some information in this direction, the augmenting subgraph is of the form *fan* and *alternating path*, see Figure 4 in Section 7. A detailed explanation of the recent progress on this problem is done in Section 7.3. We finish this section by stating a purely graph-theoretic consequence of this line of research that was proved by Christiansen [36].

Theorem 2.5 (Small augmenting subgraphs for edge colorings, Christiansen [36]). *Let G be a graph, $d: E \rightarrow [\Delta(G) + 1]$ be a partial edge coloring and $e \in U_d$ be an uncolored edge. Then there exists an augmenting subgraph H for e such that $|E(H)| \in O(\Delta(G)^7 \log(n))$.*

We remark that up to Δ factors the bound in the theorem matches the lower bound of Chang, He, Li, Pettie, and Uitto [33] who found an instance of partial edge coloring and an uncolored edge such that every augmenting subgraph has size $\Omega(\Delta(G) \log(n/\Delta(G)))$.

An analogous result for Brooks' coloring was proved by Panconesi and Srinivasan [105, 106]. Namely, every uncolored vertex in a partially vertex colored graph G with $\Delta(G) \geq 3$ colors is contained in an augmenting path of length bounded by $O(\log_{\Delta(G)}(n))$. They also demonstrated that this bound is tight. We state this result formally in Section 8, where we use it in a black-box manner in our construction of an efficient deterministic distributed algorithm for Brooks coloring on graphs of subexponential growth.

3. LOCAL model

The definition of the *LOCAL model* of distributed computing by Linial [91] was motivated by the desire to understand *distributed algorithms* in huge networks, see the book of Barenboim and Elkin [6] for a historical account. As an example, consider a huge network of wifi routers. Let us think of two routers as connected by an edge if they are close enough to exchange messages. It is desirable that such close-by routers communicate with user devices on different channels to avoid interference. In graph-theoretic language, we want to properly color the underlying network. Observe that producing the greedy coloring, which was easily solved by a sequential algorithm in the previous section, remains a highly interesting and non-trivial problem in the distributed setting, as, ideally, each vertex decides on its output color after only a few rounds of communication with its neighbors.

The LOCAL model of distributed computing formalizes this setup: we have a large network, where each vertex knows the size of the network n , and perhaps some other parameters like the maximum degree Δ . In the case of randomized algorithms, each vertex has access to a private random bit string, while in the case of deterministic algorithms, each vertex is equipped with a unique identifier from a range polynomial in the size n of the network. In one round, each vertex can exchange any message with its neighbors and can perform an arbitrary computation. The goal is to find a solution to a given problem, in our case a vertex or edge coloring of the network, in as few communication rounds as possible. As the allowed message size is unbounded, a t -round distributed algorithm can be equivalently described as a function that maps t -neighborhoods to outputs—the output of a vertex is then simply the output of this function applied to the t -neighborhood of this vertex. An algorithm is correct if and only if the collection of outputs at all vertices constitutes a correct solution to the problem.

The theory of distributed algorithms is extremely rich and the landscape of possible complexities of local graph coloring problems is very well understood. It should not come as a surprise that vertex and edge colorings play a prominent role in this endeavor. However, it might not be immediately clear why *any* non-trivial problem can be solved by a distributed algorithm.

Before we formalize the definitions, we describe a simple, yet profound, procedure that is due to Cole and Vishkin [38], which modifies a given vertex coloring of a graph G with k colors to a vertex coloring with roughly $\log(k)^{\Delta(G)}$ colors in one communication round.

Example 3.1 (Color reduction algorithm, [6, Section 3.5]). Let $G = (V, E)$ be a graph of maximum degree $\Delta \in \mathbb{N}$ and $c: V \rightarrow [k]$ be a vertex coloring. We describe how to produce a new vertex coloring with $(2\lceil \log(k) \rceil)^\Delta$ colors in one round of communication having c as an input at every vertex of G .

Let us view the color $c(v)$ as a $\{0, 1\}$ -string of length $\lceil \log(k) \rceil$. Given $v \in V$, we list for each $w \in N_G(v)$ the pair that encodes the position of the minimal bit where $c(v)$ and $c(w)$ differ together with the value of $c(v)$ at that bit. That is, each v is assigned a string $c'(v)$ of at most $\Delta(G)$ many pairs from the set $\lceil \log(k) \rceil \times \{0, 1\}$. We leave as an exercise to verify that c' is indeed a proper vertex coloring of G with $(2\lceil \log(k) \rceil)^\Delta$ colors.

A more complicated combinatorial argument, which is due to Linial [91], decreases the number of colors to $4(\Delta + 1)^2 \log^2(k)$ in one round of communication, see Hirvonen and Suomela [79, Section 4.8] or Rozhoň [109, Section 2.1.1] for a detailed explanation.

3.1. Formal definition and basic results. Our main references are the books of Barenboim and Elkin [6] and of Hirvonen and Suomela [79], and the recent survey of Rozhoň [109].

When talking about distributed algorithms, we always assume that there is a fixed class of finite graphs on which they operate. Prominent examples of classes studied in the literature include all finite graphs, graphs of degree bounded by $\Delta \in \mathbb{N}$, acyclic graphs, graphs of bounded neighborhood growth, or graphs that look locally like a fixed graph, e.g., grid graphs, high girth graphs etc.

In this survey, we work, unless stated otherwise, with the class of *graphs of degree bounded by $\Delta \in \mathbb{N}$* . In particular, the dependence on Δ is hidden in the asymptotic notation.

In the high-level overview, we mentioned that a t -round distributed algorithm assigns a color to a vertex based on its t -neighborhood and every vertex must run the same algorithm. In particular, if there are vertices with isomorphic t -neighborhoods, they should receive the same color. In order to make the model non-trivial, we need to introduce a way how to *break symmetries* in graphs that are highly symmetric. In the case of *deterministic distributed algorithms* this is done by assigning a *unique identifier* to each vertex from a range that is polynomial in the size of V . In other words, when we talk about deterministic distributed algorithms, we always assume that the graph G comes equipped with an adversarial injective map $id: V \rightarrow n^c$, where $c \geq 1$ is a fixed constant. In the case of *randomized distributed algorithms*, each vertex generates independently uniformly an *infinite sequence of random bits*. Formally, the input graph comes equipped with an adversarial map $i: V \rightarrow \{0, 1\}^{\mathbb{N}}$.

Definition 3.2 (Distributed algorithm). A *deterministic distributed algorithm* $\mathcal{A} = (\mathcal{A}_n)_n$ of complexity $(t_n)_n \in \mathbb{R}^{\mathbb{N}}$ is a function such that for every graph G of size n and every assignment of unique identifiers $id: V \rightarrow n^c$ outputs for every $v \in V$ a color

$$v \mapsto \mathcal{A}_n(B_G(v, t_n))$$

that only depends on the t_n -neighborhood of v , that is, on $B_G(v, t_n)$ together with $id \upharpoonright B_G(v, t_n)$.

A *randomized distributed algorithm* is defined analogously by replacing an assignment of unique identifiers by an assignment of infinite sequences of random bits.

Note that we implicitly assume that n , the size of the vertex set of G , is part of the input of \mathcal{A} . In particular, the functions in the sequence $\mathcal{A} = (\mathcal{A}_n)_n$ can be completely independent. We also remark that under a mild assumption on the class of graphs under consideration, the constant $c \geq 1$ can be freely thought of as being equal to 1. For instance, it is trivial to check that if there is an algorithm for $c \geq 1$, then there is one for any $c \geq c' \geq 1$, and if the class of graphs is closed under enlarging the vertex set the same holds for any $c' \geq c$.

As it is always clear from the context whether we talk about deterministic or randomized algorithms, we suppress the mentioning of unique identifiers or random bits and simply say that we apply \mathcal{A} on G . We emphasize again that when constructing distributed algorithms one should think of these assignments as being given to us by an adversary and our algorithm should perform the given task successfully for every such assignment in the case of deterministic algorithms, and with high-probability in case of randomized algorithms.

As our main goal is to discuss graph coloring problems and to keep the notation simple, we formalize the definition of *LOCAL complexity* for vertex colorings only. The definition of LOCAL complexity for edge colorings, or more generally for the so-called *locally checkable labeling (LCL) problems* that are discussed in Section 3.2, can be done completely analogously. We also remind the reader that we implicitly work with the class of graphs of degree bounded by $\Delta \in \mathbb{N}$, but the definition can be generalized to any class of finite graphs.

Definition 3.3 (LOCAL complexity of graph colorings). The *vertex coloring problem with $k \in \mathbb{N}$ colors* has a *deterministic LOCAL complexity* $(t_n)_n$ if there is a deterministic distributed algorithm $\mathcal{A} = (\mathcal{A}_n)_n$ of complexity $(t_n)_n$ such that for every graph G the output coloring

$$v \mapsto \mathcal{A}_n(B_G(v, t_n))$$

produced by $\mathcal{A} = (\mathcal{A}_n)_n$ is a vertex coloring of G with k colors,

The *randomized LOCAL complexity* is defined analogously, with the main difference being that we demand that the output coloring is a vertex coloring of G with k colors with probability at least $1 - 1/n^d$, where $d > 0$ is a fixed constant.

Note that the deterministic complexity is an upper bound on the randomized complexity as interpreting the first $\lfloor \log(n^c) \rfloor$ bits as an element of the space of unique identifiers n^c gives an injective map with probability at least $1 - \frac{1}{n}$ provided that $c \geq 3$.

In Example 3.1, we discussed how to decrease the number of colors of a vertex coloring in one round of communication. Once we are familiar with this one step procedure, nothing can stop us from iterating it. For example, if G is a 2-regular graph and we start with a vertex coloring with $k = 2^{2^{2^2}} \approx 2 \cdot 10^{19729}$ many colors, then this procedure decreases the number of colors in 4 rounds of communication as follows

$$2 \cdot 10^{19729} \mapsto 17179869184 \mapsto 4624 \mapsto 676 \mapsto 400.$$

It is not coincidental that it took us 4 steps to reduce a 4-times iterated exponential to a reasonably small number. A general argument (using the better bound discussed after Example 3.1 and treating Δ as a constant) shows that starting from vertex coloring with n^c colors, where $c > 0$ is a constant, we arrive to a coloring with $O(1)$ colors after at most $O(\log^*(n))$ many communication rounds.² A particular instance of a starting coloring might be the assignment of unique

² Recall that $\log^*(n)$ is defined as the minimum number of times we need to apply log to decrease n to a value that is at most 2.

identifiers that is given to us in the setting of deterministic algorithms. Combining these ideas allows to prove the upper bound in the following fundamental result. Note that the last step, that is, getting from constant colors to $\Delta(G) + 1$, can be done by a simple distributed greedy algorithm by iterating over the constantly many color classes and simulating the sequential greedy algorithm locally, see again [79, Section 4.8] and [109, Section 2.1.1] for more details. We also refer the reader to [6, Section 3.6-3.10] for a more detailed treatment of the dependency on Δ .

Theorem 3.4 (Distributed Greedy coloring³, Cole–Vishkin [38], Goldberg–Plotkin–Shannon [65], Linial [91]). *Let $\Delta \in \mathbb{N}$.*

- *There is a deterministic distributed algorithm of complexity $O(\log^*(n))$ that produces a vertex coloring with $\Delta + 1$ colors on the class of graphs of degree bounded by Δ .*
- *There is a deterministic distributed algorithm of complexity $O(\log^*(n))$ that produces an edge coloring with $2\Delta - 1$ colors on the class of graphs of degree bounded by Δ .*

Moreover, in both cases $\Omega(\log^(n))$ rounds are necessary.*

For the proof of the lower bound, which is originally due to Linial [91], together with a description of randomized distributed algorithm for these colorings that work *without knowledge of the size of the graph*, we refer the reader to the paper of Holroyd, Schramm and Wilson [80, Sections 1 and 2], where the connection to *random processes*, in particular, to *finitary factors of iid*, is discussed.

In the light of Theorem 3.4, we are compelled to raise the following questions. Can we decrease the number of colors? Are there similarly efficient distributed algorithms for Brooks or Vizing colorings? Does randomness give us any advantage?

In order to answer these questions, it is useful to discuss the global picture of possible LOCAL complexities.

3.2. Global picture of the LOCAL model. The landscape of possible LOCAL complexities of graph coloring problems where the correctness of a solution can be checked locally, so-called *locally checkable labeling (LCL) problems* (see, e.g., [109, Section 2.4.2] for a formal definition),⁴ is very well understood. This understanding has been a major success of several recent lines of work on local algorithms, including [2, 3, 22, 27, 32, 34, 35, 60–62, 64, 76, 103, 110] and many others. For the class of graphs with degree bounded by Δ , a clean picture emerges, see Figure 1.⁵

The two most interesting features of this classification is that (i) there is only one regime when randomness helps, (ii) there are automatic *speed-up* results that

³This is typically referred to as $\Delta+1$ -coloring, rather than “greedy”.

⁴Apart from graph colorings, the class of LCL problems include for instance the perfect matching problem, the sinkless orientation problem and various versions of the unfriendly coloring problem.

⁵The picture is slightly simplified, in particular, $\log(n)$ and $\log(\log(n))$ complexities are only conjectured [35], the best up-to-date general upper bounds in the corresponding regimes are $O(\log^3(n))$ and $O(\log^3(\log(n)))$.

often turn a proposed distributed algorithm into a faster one in a completely black-box manner. In particular, item (ii) automatically answers question of the form, is there an LCL problem that has (deterministic) LOCAL complexity $\Theta(\sqrt{\log(n)})$?

The fundamental result of Bernshteyn [12], mentioned in the introduction, turns colorings produced by efficient distributed algorithms into measurable colorings. Here, *efficient distributed algorithms* refers to the $o(\log(n))$ regime, see Figure 1. In particular, one needs to specify if randomness is allowed as this is the regime, where it gives an advantage.

To understand the LOCAL complexities of our graph coloring problems, we need to locate the corresponding blue dots in Figure 1. We have seen in Theorem 3.4 that the LOCAL complexity of *greedy coloring*, both deterministic and randomized, is $\Theta(\log^*(n))$, the class of LCL problems with this LOCAL complexity is sometimes referred to as the class of *symmetry breaking problems*.

Deterministic lower bounds $\Omega(\log(n))$ for Brooks and strong version of Vizing coloring are fundamental results in the theory of the LOCAL model. The original argument uses the *round elimination method* of Brandt, see [22] for more details.




FIGURE 1. (Courtesy of Jukka Suomela) The blue dots represent classes of LCL problems (size of the dots suggest the importance of the class), that is, every LCL problem, e.g., vertex coloring with Δ colors, belongs to one of the dots. *Deterministic LOCAL complexity* is given by the projection of the dot to the horizontal axis and *randomized LOCAL complexity* is given by the projection to the diagonal. Light salmon colored areas do not contain any complexity class of LCL problems, in particular, they represent the *speed-up* results.

In Figure 1, we see that this automatically implies the randomized lower bound $\Omega(\log(\log(n)))$. Rather surprisingly the lower bounds also follow from the groundbreaking work of Marks [96] in Borel combinatorics by applying Bernshteyn's correspondence [12]. We remark that Marks' argument involves infinite games and one of the deepest results in mathematical logic, the Borel determinacy theorem of Martin [99]. In Section 6, we show how to adapt Marks' game technique directly in the LOCAL model and get an alternative and very simple proof of these fundamental lower bounds based on determinacy of finite games. The following theorem should be compared with Theorem 4.7.

Theorem 3.5 (Distributed lower bounds for colorings, [26, 33]). *Let $\Delta \geq 3$.*

- *Vertex coloring with Δ colors on the class of finite graphs of degree bounded by Δ that do not contain $K_{\Delta+1}$ needs $\Omega(\log(n))$ rounds in the deterministic LOCAL model.*
- *Edge coloring with $2\Delta - 2$ colors on the class of finite graphs of degree bounded by Δ needs $\Omega(\log(n))$ rounds in the deterministic LOCAL model.*

3.3. Upper bound for Brooks and Vizing colorings. Finding optimal distributed algorithms for graph colorings and decreasing the number of colors, ideally matching the Brooks and Vizing bounds, has been a major challenge in the theory of distributed computing. In this section, we give a brief overview of the up-to-date results. Recall that for simplicity we treat Δ as a constant, see the survey of Rozhoň [109, Section 1.6] for references, results and open problems in the regime when the complexity depends on Δ as well.

The local complexity, both deterministic and randomized, of the greedy vertex coloring is $\Theta(\log^*(n))$ by Theorem 3.4. The first non-trivial result concerning distributed Brooks colorings, that remained the state of the art for more than two decades, was achieved by Panconesi and Srinivasan [105, 106]. The combinatorial core of their algorithm is based on the augmenting graph result for Brooks' theorem that is discussed after Theorem 2.5, and yields both a deterministic and a randomized distributed algorithm of local complexity $O(\log^3(n))$.

In the special case of trees of degree bounded by $\Delta \geq 3$, Barenboim and Elkin [5] described a deterministic distributed algorithm of local complexity $O(\log(n))$ and Chang, Kopelowitz and Pettie [34] gave a randomized one of local complexity $O(\log(\log(n)))$, when Δ is sufficiently large. Importantly, these upper bounds exactly match the lower bounds of Brandt, Fischer, Hirvonen, Keller, Lempäälä, Rybicki, Suomela and Uitto [26] that were proved for the class of graphs of girth $\Omega(\log(n))$, in particular, they also apply to the class of finite trees.

The following are the up-to-date results of Ghaffari, Hirvonen, Kuhn and Maus [63], see also [54]. Note that these results are only one logarithmic factor away from the lower bounds.

Theorem 3.6 (Distributed Brooks coloring, [63]). *Let $\Delta \geq 3$.*

- (1) *Vertex coloring with Δ colors on the class of finite graphs of degree bounded by Δ that do not contain $K_{\Delta+1}$ has deterministic local complexity $O(\log^2(n))$.*

- (2) *Vertex coloring with Δ colors on the class of finite graphs of degree bounded by Δ that do not contain $K_{\Delta+1}$ has randomized local complexity $O(\log^2(\log(n)))$.*

For results that concern decreasing the number of colors even further, we refer the reader to the papers of Bamas and Esperet [4] for matching optimal coloring and Chung, Pettie, and Su [37] for colorings of graphs of girth at least 3, 4 or 5.

Applying the vertex coloring results to the line graph of a given graph G implies immediately the same upper bounds for edge colorings of G . Namely, the greedy edge coloring with $2\Delta - 1$ colors have a local complexity $\Theta(\log^*(n))$. Similarly, the same bounds as in Theorem 3.6 apply to edge colorings with $2\Delta - 2$ colors. Chang, He, Li, Pettie, and Uitto [33] showed that edge coloring with $2\Delta - 2$ colors have a deterministic lower bound $\Omega(\log(n))$ even when restricted to the class of finite trees of degree bounded by Δ . In the same paper [33], the authors developed a randomized algorithm of local complexity $O(\text{poly}(\log(\log(n))))$ for edge coloring with $\Delta + \sqrt{\Delta} \text{poly}(\log(n))$ colors.

Getting closer to the Vizing bound, Su and Vu [115] devised a randomized distributed algorithm for edge coloring with $\Delta+2$ colors of local complexity $\text{poly}(\log(n))$. Finally, Bernshteyn [11] matched the Vizing bound by both deterministic and randomized distributed algorithms of local complexity $\text{poly}(\log(n))$. This was subsequently improved by Christiansen [36], and Bernshteyn and Dhawan [14] to the current state of art, see Section 7 for more details.

Theorem 3.7 (Distributed Vizing coloring, [14]). *Let $\Delta \in \mathbb{N}$.*

- (1) *Edge coloring with $\Delta+1$ colors on the class of finite graphs of degree bounded by Δ has deterministic local complexity $O(\text{poly}(\log(\log(n))) \log^5(n))$.*
- (2) *Edge coloring with $\Delta+1$ colors on the class of finite graphs of degree bounded by Δ has randomized local complexity $O(\text{poly}(\log(\log(n))) \log^2(n))$.*

4. Measurable combinatorics

Generalizing concepts from finite graph theory to uncountable graphs in the most straightforward way often results in counter-intuitive behavior: the most spectacular example of this is probably the Banach-Tarski paradox. It states that a three-dimensional ball of unit volume can be decomposed into finitely many pieces that can be moved by isometries (distance preserving transformations such as rotations and translations) to form two three-dimensional balls each of them with unit volume(!). The graph theoretic problem lurking behind this paradox is the following: fix finitely many isometries of \mathbb{R}^3 and then consider a graph where x and y are connected if there is an isometry that sends x to y . Then our task becomes to find a perfect matching in the appropriate subgraph of this graph – namely, the bipartite subgraph where one partition contains points of the first ball and the other contains points of the other two balls. Banach and Tarski have shown that, with a suitably chosen set of isometries, the Axiom of Choice implies the existence of such a matching. In contrast, since isometries preserve the Lebesgue measure, the pieces in the decomposition cannot be Lebesgue measurable. Such observations lead to

the following heuristic: if we want to use (uncountably) infinite graphs to model the behavior of large finite graphs, we must impose measurability constraints on the objects we are investigating.

Recently, results in this direction that lie on the border of combinatorics, logic, group theory, and ergodic theory led to an emergence of a new field often called *descriptive* or *measurable combinatorics*. The field focuses on the connection between the discrete and continuous and is largely concerned with the investigation of graph-theoretic concepts. See [12, 40, 43, 46, 49, 56, 66, 88, 90, 95, 96, 98] for some of the most important results, and [86, 97, 107] for surveys.

The usual setup in descriptive combinatorics is that we have a graph with uncountably many connected components, each being a countable graph of bounded degree. For example, in case of the Banach-Tarski paradox, the vertices of the underlying graph are the points of the three balls, edges correspond to isometries, and the degree of each vertex is bounded by the number of chosen isometries.

Before we formally introduce Borel graphs and their colorings, we illustrate the subject on a classical -and impossible to omit- example from ergodic theory that simulates nicely the type of behavior that one should expect when dealing with measurable sets, and how such behaviors can be dramatically different from the classical case.

Example 4.1 (Irrational rotations). Let $\alpha \in [0, \pi]$ be such that $\frac{\alpha}{\pi}$ is irrational and denote by T_α the rotation of the circle \mathbb{S}^1 by α . Define the graph \mathcal{G}_α with vertex set \mathbb{S}^1 as follows. We say that $x, y \in \mathbb{S}^1$ form an edge in \mathcal{G}_α if $T_\alpha(x) = y$ or $T_\alpha(y) = x$.

Observe that \mathcal{G}_α is 2-regular and, as $\frac{\alpha}{\pi}$ is irrational, acyclic. In particular, there are uncountably many connectivity components in \mathcal{G}_α each isomorphic to the canonical graph on \mathbb{Z} . As each of these connectivity components has chromatic number 2, it follows that $\chi(\mathcal{G}_\alpha) = 2$. In other words, we can write $\mathbb{S}^1 = B_0 \sqcup B_1$ such that $T_\alpha(B_0) = B_1$, $T_\alpha(B_0) \cap B_0 = \emptyset$ and $T_\alpha(B_1) \cap B_1 = \emptyset$.

A typical question in measurable combinatorics asks if the sets B_0 and B_1 can be Lebesgue measurable, that is, if the *measurable chromatic number* of \mathcal{G}_α is equal to 2 as well. We demonstrate that this is *not* possible. Suppose for a contradiction that B_0 and B_1 are Lebesgue measurable. Then by the properties above combined with the fact that T_α preserves measure, we conclude that $\lambda(B_0) = \lambda(B_1) = 1/2$. By the Lebesgue density point theorem, there is a nonempty open interval U with the property that $\lambda(U \setminus B_0) < \frac{1}{3}\lambda(U)$. But then as $2\alpha/\pi$ is also irrational, there is an odd $n \in \mathbb{N}$ such that $\lambda(T_\alpha^n(U) \cap U) > \frac{2}{3}\lambda(U)$, where T_α^n denotes the application of T_α n -times. Note that $T_\alpha^n(B_0) = B_1$ as this holds for every odd $n \in \mathbb{N}$. An easy volume argument in U then implies that $B_1 \cap B_0 = T_\alpha^n(B_0) \cap B_0 \neq \emptyset$, a contradiction.

In a stark contrast, a vertex coloring with 3 colors can be achieved with Lebesgue measurable sets, see Example 4.5.

4.1. Definitions and basic results. The connection between the local and global structure of the graphs that are studied in measurable combinatorics is formalized through classical descriptive set theoretic notions, see [84] for introduction to the subject. Namely, the vertex set of a graph comes equipped with a σ -algebra of

measurable, or definable, subsets. The structure and all the operations on the graphs that we are interested in, for example, graph colorings or partitions, have to be measurable with respect to this σ -algebra or its powers.

To avoid any pathological behavior, we always assume that the σ -algebra turns the vertex set into a *standard Borel space*. Formally, a standard Borel space is a pair (X, \mathcal{B}) , where X is a set and \mathcal{B} is a σ -algebra of subsets of X that coincide with a σ -algebra of Borel sets for some complete separable metric on X . Slightly abusing the notation, we refer to elements of \mathcal{B} as *Borel sets*. By the *Borel isomorphism theorem*, all uncountable standard Borel spaces are Borel isomorphic, so one may safely assume that the vertex set of every Borel graph in this paper is, for example, the unit interval $[0, 1]$ endowed with the σ -algebra of Borel sets. If (X, \mathcal{B}) is a standard Borel space, then the product $X \times X$ and the set of all unordered pairs $\binom{X}{2}$ is naturally a standard Borel space.

Definition 4.2. A *Borel graph* \mathcal{G} is a triple (V, E, \mathcal{B}) , where (V, \mathcal{B}) is a standard Borel space, (V, E) is a graph and $E \subseteq \binom{V}{2}$ is a Borel subset of the standard Borel space $\binom{V}{2}$ of unordered pairs of V .

Borel graphs often appear in the nature. The most relevant examples for our purposes are so called Schreier graphs of group actions. These are frequently investigated in the theory of random processes, measured group theory, ergodic theory and dynamics.

Example 4.3 (The Schreier graph $\text{Sch}(\Gamma, S, X)$). Let Γ be a countable group and $S \subseteq \Gamma$ be a generating set of Γ . Assume that $\Gamma \curvearrowright X$ is an action of Γ on a set X . The *Schreier graph* $\text{Sch}(\Gamma, S, X)$ of the action is a graph on the set X , where $x \neq x'$ form an edge if there is $\gamma \in S$ such that $\gamma \cdot x = x'$ or $\gamma \cdot x' = x$.

If (X, \mathcal{B}) is a standard Borel space and the action $\Gamma \curvearrowright X$ is Borel, meaning that $\gamma \cdot \cdot : X \rightarrow X$ is a Borel isomorphism for every $\gamma \in \Gamma$, then $\text{Sch}(\Gamma, S, X)$ is a Borel graph.

The simplest example of a Schreier graph is the (right) *Cayley graph of Γ with respect to S* , denoted by $\text{Cay}(\Gamma, S)$, that is, the Schreier graph that comes from the right multiplication action of Γ on itself. Observe that the graph in Example 4.1 is of the form $\text{Sch}(\Gamma, S, X)$, where $\Gamma = \mathbb{Z}$, $S = \{1\}$ and $X = \mathbb{S}^1$.

Note that if S is finite, then the degree of $\text{Sch}(\Gamma, S, X)$ is bounded by $2|S|$. In particular, the degree is uniformly bounded. Unless explicitly stated otherwise, we are exclusively concerned with graphs with uniformly bounded degree.

Definition 4.4. Let $\mathcal{G} = (V, E, \mathcal{B})$ be a Borel graph. The *Borel chromatic number of \mathcal{G}* , denoted by $\chi_{\mathcal{B}}(\mathcal{G})$, is the minimal $k \in \mathbb{N}$ such that there is $c: V \rightarrow [k]$ that is a vertex coloring with k colors and $c^{-1}(\{i\}) \in \mathcal{B}$ for every $i \in [k]$. In other words, there is a partition $V = B_1 \sqcup \dots \sqcup B_k$ into Borel \mathcal{G} -independent sets, where $X \subseteq V$ is called *\mathcal{G} -independent* if it does not span any edge in \mathcal{G} .

Analogously, we can define the *Borel chromatic index of \mathcal{G}* and denote it as $\chi'_{\mathcal{B}}(\mathcal{G})$.

Is there an analogue of greedy colorings in the Borel context? As we most of the time implicitly assume that the underlying vertex set (V, \mathcal{B}) is uncountable, it is a

priori not even clear that there is a countable decomposition $V = X_1 \sqcup X_2 \sqcup \dots$ into \mathcal{G} -independent Borel sets. Let us remark that the main result of the seminal paper of Kechris-Solecki-Todorćević [88], which initiated the study of Borel chromatic numbers, was to give a complete characterization of the Borel graphs that admit such a countable decomposition. In particular, one of their results states that in the bounded degree case such a decomposition always exists. Compare the following example with Example 3.1.

Example 4.5 (Countable colorings from finite degrees). Let $\mathcal{G} = (V, E, \mathcal{B})$ be a Borel graph that is locally finite, that is, $\deg_{\mathcal{G}}(v) \in \mathbb{N}$ for every $v \in V$. By the Borel isomorphism theorem, we may assume that $V = 2^{\mathbb{N}}$, that is, every $v \in V$ can be represented by an infinite sequence of 0s and 1s. The crucial observation is that for every $v \in V$ and a finite set $A \subseteq V \setminus \{v\}$, there is $\ell \in \mathbb{N}$ such that $v \upharpoonright [\ell] \neq x \upharpoonright [\ell]$ for every $x \in A$. Hence, for every $v \in V$ we may set $c'(v) = v \upharpoonright [k]$, where $k \in \mathbb{N}$ is the minimal number that have this property with respect to the finite set $N_{\mathcal{G}}(v)$. It is a matter of routine work with Boolean operations to check that the map $c': V \rightarrow 2^{<\mathbb{N}}$, where $2^{<\mathbb{N}}$ is the set of *finite* sequences of 0s and 1s, is a Borel map that has the property that $c'(v) \neq c'(w)$ for every $\{v, w\} \in E$. In particular, composing c' with *any* bijection between the countable sets $2^{<\mathbb{N}}$ and \mathbb{N} defines a Borel map $c: V \rightarrow \mathbb{N}$ that is a vertex coloring with \mathbb{N} colors, that is,

$$V = \bigsqcup_{i \in \mathbb{N}} c^{-1}(\{i\})$$

is a decomposition into Borel \mathcal{G} -independent sets.

Example 4.5 serves as a basic symmetry breaking device in the measurable context that allows to simulate local algorithms. It is an interesting philosophical problem if the computational power of Borel constructions consists merely of transfinite iteration of this process possibly on increasing graph powers of the underlying graph.

Similarly as in the proof of the upper bound from Theorem 3.4, it is possible to modify the countable coloring above into a greedy coloring by simulating locally the sequential greedy algorithm.

Theorem 4.6 (Borel greedy colorings, [88]). *Let $\mathcal{G} = (V, E, \mathcal{B})$ be a Borel graph such that $\Delta(\mathcal{G}) \in \mathbb{N}$. Then we have $\chi_{\mathcal{B}}(\mathcal{G}) \leq \Delta(\mathcal{G}) + 1$ and $\chi'_{\mathcal{B}}(\mathcal{G}) \leq 2\Delta(\mathcal{G}) - 1$.*

In a groundbreaking paper [96], Marks proved that the Borel greedy colorings are optimal even for Borel graphs that are regular and acyclic. The proof uses one of the deepest theorems in descriptive set theory, the determinacy of two player games with a Borel pay-off set, proven by Martin [99].

Theorem 4.7 (Borel lower bounds for colorings, [96]). *Let $\Delta > 2$.*

- (1) *There is an acyclic Δ -regular Borel graph \mathcal{G} that does not admit Borel vertex coloring with Δ colors.*
- (2) *There is an acyclic Δ -regular Borel graph \mathcal{G} that does not admit Borel edge coloring with $2\Delta - 2$ colors.*

One should compare the above statement with Theorem 3.5. We will discuss the proof of this result in some detail in Section 6.

4.2. Measure, as the analogue of randomness. In ergodic theory, measured group theory or the theory of random processes on countable graphs, people study combinatorial problems on Borel graphs that are equipped with an additional structure, *Borel probability measure* that satisfies the *mass transport principle*. This concept is also known under the names *pmp* (*probability measure preserving*) *action*, *graphing* or *unimodular random network*.

Formally, let $\mathcal{G} = (V, E, \mathcal{B})$ be a Borel graph and μ be a Borel probability measure on (V, \mathcal{B}) . We say that μ satisfies the *mass transport principle* if

$$(4.1) \quad \int_A \deg_{\mathcal{G}}(x, B) \, d\mu = \int_B \deg_{\mathcal{G}}(x, A) \, d\mu$$

for every $A, B \in \mathcal{B}$. That is, counting edges from a set A to B and from B to A is the same.

Typical examples of Borel graphs that satisfies the mass transport principle are the Schreier graphs, introduced in Example 4.3, under the additional assumption that the Borel action preserves a given Borel probability measure μ . We complement this family of examples with examples coming from the theory of random processes.

Example 4.8 (Examples of mass transport). (a) *iid graphings*: Let Γ be a countable group and $S \subseteq \Gamma$ be a finite *symmetric* generating set of Γ . The vertex set of the iid graphing \mathcal{G}_{Γ} consists of all maps $i: \Gamma \rightarrow \{0, 1\}^{\mathbb{N}}$, and two such maps $i \neq j$ form an edge if there is $s \in S$, such that

$$i(s^{-1} \cdot \gamma) = j(\gamma)$$

for every $\gamma \in \Gamma$. If we endow the vertex set with the iid power of the coin-flip measure, we get that almost surely every connectivity component of \mathcal{G}_{Γ} is isomorphic to $\text{Cay}(\Gamma, S)$. It is easy to check that \mathcal{G}_{Γ} is indeed a graphing.

(b) *percolation graph on \mathbb{Z}^d* (see [1] for an introduction to the subject): A more complicated example of a graphing can be described using the notion of *site* (or vertex) *percolation*. Consider the iid graphing $\mathcal{G}_{\mathbb{Z}^d}$ and fix $p \in [0, 1]$ with the property that the site percolation on \mathbb{Z}^d contains infinite cluster with positive probability. Define a subset of the vertex set A to consist of those functions $i: \mathbb{Z}^d \rightarrow \{0, 1\}^{\mathbb{N}}$ such the the origin belongs to an infinite cluster. Define $\mathcal{G}_{\mathbb{Z}^d, p}$ to be the graph induced in A . It can be again easily shown that, after re-normalizing the measure, $\mathcal{G}_{\mathbb{Z}^d, p}$ is indeed a graphing. Observe that $\mathcal{G}_{\mathbb{Z}^d, 1} = \mathcal{G}_{\mathbb{Z}^d}$.

The mass transport principle translates some of the intuition about finite graphs into the measurable setup. In the following definition of a *measured graph*, we work in bigger generality and do not require any condition on the Borel probability measure.

Definition 4.9 (Measured graph and measurable chromatic numbers). A *measured graph* is a quadruple $\mathcal{G} = (V, E, \mathcal{B}, \mu)$, where (V, E, \mathcal{B}) is a Borel graph and μ is a

Borel probability measure on (V, \mathcal{B}) . A measured graph \mathcal{G} is called a *graphing* if μ satisfies the mass transport principle.

The μ -measurable chromatic number of \mathcal{G} , denoted by $\chi_\mu(\mathcal{G})$ is the minimal $k \in \mathbb{N}$ such that there is a μ -null set $N \in \mathcal{B}$ and a decomposition $V = N \sqcup B_1 \sqcup \dots \sqcup B_k$ such that B_i is a Borel \mathcal{G} -independent set for every $i \in [k]$. In other words, the Borel chromatic number is equal to k off of a μ -null set.

The μ -measurable chromatic index of \mathcal{G} is defined analogously and is denoted by $\chi'_\mu(\mathcal{G})$.

It follows immediately that $\chi_\mu(\mathcal{G}) \leq \chi_{\mathcal{B}}(\mathcal{G})$ for every measured graph $\mathcal{G} = (V, E, \mathcal{B}, \mu)$, where we abuse the notation and define $\chi_{\mathcal{B}}(\mathcal{G})$ to be the Borel chromatic number of the underlying Borel graph of (V, E, \mathcal{B}) of \mathcal{G} . We remark that it is one of the most intriguing problems in the area to find examples of local graph coloring problem that could be solved on graphings but not on measured graphs in general.

Does ignoring events that happen on null sets help with graph colorings, in particular, can the inequality $\chi_\mu(\mathcal{G}) \leq \chi_{\mathcal{B}}(\mathcal{G})$ be strict? Does the analogy between deterministic and randomized algorithms, and Borel and measurable colorings go deeper? The answer to the first question is clear yes as we demonstrate next. The answer to the second question is more complicated, see Remark 4.13.

Investigating measurable chromatic numbers on bounded degree Borel graphs has a rich history that connects the theory of random graphs, graph limits, random processes and measured group theory, see the survey of Kechris and Marks [86, Section 6] for many results in this direction. An illustrative example of this rich connection is the general lower bound $\Omega(\Delta/\log(\Delta))$ on measurable chromatic numbers of acyclic Δ -regular graphs (where Δ is large) that follows from the seminal result of Bollobás [19] on the size of maximal independent set in random regular graphs combined with the theory of *local-global* convergence developed by Hatami, Lovász and Szegedy [78].

Complementing Theorem 3.6, the measurable analogue of Brooks theorem was proved by Conley, Marks and Tucker-Drob [42].

Theorem 4.10 (Measurable Brooks coloring, [42]). *Let $\Delta \geq 3$ and $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ be a measured graph that does not contain $K_{\Delta+1}$ as a subgraph and such that $\Delta(\mathcal{G}) \leq \Delta$. Then $\chi_\mu(\mathcal{G}) \leq \Delta$.*

Combining Theorem 4.7 and Theorem 4.10 yields for every $\Delta \geq 3$ the existence of an acyclic Δ -regular Borel graph \mathcal{G} such that $\chi_\mu(\mathcal{G}) \leq \Delta < \Delta + 1 = \chi_{\mathcal{B}}(\mathcal{G})$ for every Borel probability measure μ .⁶ Note that this result exactly matches with Theorem 3.5 and Theorem 3.6 from the theory of distributed computing. While all four theorems were proven independently, it follows from the correspondence of Bernshteyn [12] that Theorem 4.7 implies Theorem 3.5, and Theorem 3.6 implies Theorem 4.10 in a black-box manner.

⁶In fact, for large Δ the gap is much bigger as the measurable chromatic number of acyclic Δ -regular graphs is bounded by $O(\Delta/\log(\Delta))$ by [12, Theorem 3.9].

In the case of measurable edge colorings, Csóka, Lippner and Pikhurko [47] showed that $\chi'_\mu(\mathcal{G}) \leq \Delta + 1$ for a graphing $\mathcal{G} = (V, \mathcal{B}, E, \mu)$ that does not contain odd cycles and proved an upper bound of $\Delta + O(\sqrt{\Delta})$ colors for graphings in general. In a related result, Bernshteyn [9, Theorem 1.3] proved that $\Delta + o(\Delta)$ colors are enough (even for the so-called list-coloring version) provided that the graphing factors to the shift action $\Gamma \curvearrowright [0, 1]^\Gamma$ of a finitely generated group Γ . Answering a question of Abért, the first author and Pikhurko [71] proved a measurable version of Vizing's theorem for graphings, that is, $\chi'_\mu(\mathcal{G}) \leq \Delta + 1$ for any graphing $\mathcal{G} = (V, \mathcal{B}, \mu, E)$. As we discuss in Section 7, the technique developed in [71] was applied in the LOCAL model by Bernshteyn [11] and others [14, 36]. Finally, the full measurable analogue of Vizing's theorem was derived by the first author [69].

Theorem 4.11 (Measurable Vizing coloring, [69]). *Let $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ be a measured graph. Then $\chi'_\mu(\mathcal{G}) \leq \Delta(\mathcal{G}) + 1$.*

Combining Theorem 4.11 and Theorem 4.7 gives for every $\Delta \geq 3$ an example of an acyclic Δ -regular Borel graph \mathcal{G} such that $\chi'_\mu(\mathcal{G}) \leq \Delta + 1 < 2\Delta - 1 = \chi'_\mathcal{B}(\mathcal{G})$ for every Borel probability measure μ .

A Borel version of Vizing's theorem is possible under additional assumption on the structure of the graph. Most notably, Bowen and Weilacher [20] showed that $\chi'_\mathcal{B}(\mathcal{G}) \leq \Delta(\mathcal{G}) + 1$ under the assumption that \mathcal{G} is bipartite and has finite Borel asymptotic dimension, notion that generalizes the classical notion of asymptotic dimension introduced by Gromov [75] to the setting of measurable combinatorics by Conley, Jackson, Marks, Seward and Tucker-Drob [40], and Bernshteyn and Dhawan showed that $\chi'_\mathcal{B}(\mathcal{G}) \leq \Delta(\mathcal{G}) + 1$ under the assumption that \mathcal{G} has subexponential neighborhood growth. Stronger results in the special case of free Borel \mathbb{Z}^d -actions, that is, Borel edge coloring with $2d$ colors which is the analogue of König's line coloring theorem in this setting, were obtained independently around the same time in [7, 31, 74, 121].

Remark 4.12 (Baire measurable colorings). A topological relaxation of Borel colorings is formalized using the notion of Baire measurable functions and meager sets. Similarly to the measurable case, we are interested in coloring a given Borel graph off of a topologically negligible set for any compatible Polish topology that generates the underlying standard Borel structure. We refer the reader to the book of Kechris [84] for basic notions and to the survey of Kechris and Marks [86] for results in that direction.

Remark 4.13 (Derandomization). As illustrated in Figure 1 and by the distributed Brooks' theorem, in the regime $o(\log(n))$ randomness gives an advantage in the LOCAL model for solving LCL problems on graphs of degree bounded by $\Delta \geq 3$. This behavior might change when we restrict to a smaller class of graphs. For example, randomness does not help on graphs that look locally like the grid \mathbb{Z}^d , see [27], and the celebrated conjecture of Chang and Pettie [35] would imply the same for any graph class of subexponential growth.

On the side of measurable combinatorics, Csóka, Grabowski, Máthé, Pikhurko and Tyros [45] found a Borel version of the Lovász Local Lemma for graphs of

subexponential growth, which when combined with the correspondence of Bernshteyn shows that if an LCL problem has *randomized* LOCAL complexity $O(\log(n))$ then it can be solved in a *Borel* way on such graphs, see [12, Theorem 2.15]. This is a particular instance of a derandomization on graphs of subexponential growth, see also [16, 44, 117].

Naturally, it is tempting to ask about a higher analogue of derandomization. That is, for what class of Borel graphs does the classes of LCL problems that admit *measurable* and *Borel* solution coincide? While the answer is positive for Borel graphs that look locally like \mathbb{Z} even if we allow LCL problems with inputs [72], a counterexample for Borel graphs that look locally like \mathbb{Z}^2 was recently constructed by Berlow, Bernshteyn, Lyons and Weilacher [8].

A related question asks whether assuming that a given LCL problem can be solved measurably on *all* measured graphs of degree at most Δ implies the existence of a Borel solution on subexponential growth Borel graphs of degree at most Δ . The answer to this question turned out to be negative as well, in a recent work in progress a counterexample was found by Bowen, the first author and Rozhoň. Namely, it can be shown that the LCL problem that is a union of the sinkless orientation problem and “mark the line” problem can be solved measurably on all graphs, but does not admit a Borel solution on graphs that can have growth arbitrarily close to a linear growth.

5. Summary of coloring results and related topics

For the convenience of the reader, we summarize the results presented in previous sections in a table. The second column, that is, greedy coloring, refers to both a vertex coloring with $\Delta + 1$ colors and an edge coloring with $2\Delta - 1$ colors. Note that while the conjectured LOCAL complexity for Brooks coloring is $\Theta(\log(n))$ deterministic and $\Theta(\log(\log(n)))$ randomized, the randomized LOCAL complexity of Vizing coloring remains mysterious, see Problem 9.1.

Recall that we treat Δ as a constant, in particular, $O(n) = O(m)$, where n and m denote the sizes of the vertex and edge set of a given finite graph G . We denote as DLOCAL and RLOCAL the deterministic and randomized LOCAL model respectively.

5.1. Decidability and complexity. Complexity considerations are an important topic not only in classical graph theory but also in measurable combinatorics and distributed computing, where they give natural lower bounds on the existence of measurable solutions or efficient distributed algorithms.

We have seen that the Brooks and Vizing colorings can be constructed by a polynomial time algorithm. What if we want to find a “true” coloring that would match the chromatic number or index of a given graph? A classical result of Karp [83] says that deciding the existence of a vertex coloring with k colors of a given graph is **NP**-complete for every $k \geq 3$. The situation with chromatic numbers is more complicated, even approximate computations are known to be hard in various senses. For example, we have the following.

| | Graphs of degree bounded by Δ | | |
|----------------------|--------------------------------------|------------------------------|---|
| | Greedy coloring | Brooks coloring | Vizing coloring |
| DLOCAL upper bound | $O(\log^*(n))$ [38, 65, 91] | $O(\log^2(n))$ [63] | $\text{poly}(\log(\log(n))) \log^5(n)$ [14] |
| DLOCAL lower bound | $\Omega(\log^*(n))$ [91] | $\Omega(\log(n))$ [26] | $\Omega(\log(n))$ [26] |
| RLOCAL upper bound | $O(\log^*(n))$ [38, 65, 91] | $O(\log^2(\log(n)))$ [63] | $\text{poly}(\log(\log(n))) \log^2(n)$ [14] |
| RLOCAL lower bound | $\Omega(\log^*(n))$ [91] | $\Omega(\log(\log(n)))$ [34] | $\Omega(\log(\log(n)))$ [33] |
| Borel | yes [88] | no [96] | no [96] |
| measurable | yes [88] | yes [42] | yes [69] |
| augmenting subgraphs | $\Theta(1)$ | $\Theta(\log(n))$ [105] | $\Theta(\log(n))$ [33, 36] |
| deterministic | $O(n)$ | $O(n)$ | $O(n \log(n))$ |
| sequential algorithm | | [113] | [57] |
| randomized | $O(n)$ | $O(n)$ | $O(n)$ |
| sequential algorithm | | [113] | [14] |

Theorem 5.1 ([125]). *For every $\varepsilon > 0$ it is **NP**-hard to decide the chromatic number up to a factor of $n^{1-\varepsilon}$ of a graph of size n .*

Nevertheless, a combination of the remarkable theorems of Emden-Weinert, Hougardy and Kreuter [52], and Molloy and Reed [102] shows that there is a dichotomy for detecting large chromatic numbers via local obstacles, akin to Brooks' theorem. Let $\Delta \in \mathbb{N}$ and set k_Δ to be the maximum integer such that $(k+1)(k+2) \leq \Delta$, i.e., $k_\Delta \approx \sqrt{\Delta} - 2$.

Theorem 5.2. *Let Δ be large. Then we have the following.*

- [52] *It is **NP**-complete problem to check whether $\chi(G) \leq c$ for graphs of degree bounded by Δ , where $3 \leq c \leq \Delta - k_\Delta - 1$,*
- [102] *There is a linear time deterministic algorithm to check whether $\chi(G) \leq c$ for graphs of degree bounded by Δ , where $\Delta - k_\Delta \leq c$. Moreover, producing vertex coloring with c colors, if it exists, can be done in polynomial time.*

A recent work of Bamas and Esperet [4] showed that the same dichotomy⁷ holds in the LOCAL model as well. In particular, they showed that in the regime when c is above the threshold, there is a fast randomized distributed algorithm that produces vertex coloring with c colors or finds a local obstruction if such coloring does not exist. On the other hand, they show that if c is below the threshold, then the LOCAL complexity of any distributed algorithm that produces vertex coloring with c colors on graphs that have chromatic number equal to c is $\Omega(n)$. These results have a complete analogue in measurable combinatorics, see [12, Section 3.A].

For edge colorings, the situation is simpler in the sense that $\chi'(G) \in \{\Delta(G), \Delta(G)+1\}$ by Vizing's theorem. There are known sufficient conditions for $\chi'(G) = \Delta(G)$. For example, by König's theorem it is enough if G is bipartite, see the book of Stiebitz, Scheide, Toft, and Favrholt [114] for more details in this direction. Holyer [81], however, showed that distinguishing between the two cases in general is complicated.

⁷The threshold in their result is off by 1 compared to Theorem 5.2

Theorem 5.3 ([81]). *For every $\Delta \geq 3$. It is **NP**-complete to decide if a graph G of maximum degree at most Δ satisfies $\chi'(G) \leq \Delta$.*

In descriptive set theory, complexity of a collection of structures with a given property, for example, Borel graphs of a given Borel chromatic number, is measured using the complexity of the set of *codes* for these structures. These codes can be identified with real numbers, hence a collection of structures with a given property can be thought of as a subset of \mathbb{R} , and the complexity is measured using the *projective hierarchy* of subsets of \mathbb{R} , see [84, Section 37] for the definition and basic properties of projective sets. Roughly speaking, the complexity measures the number of alternating quantifiers over \mathbb{R} that are needed for the definition of the set. For instance, the standard definition of the set of Borel graphs of Borel chromatic number at most $k \in \mathbb{N}$, that is,

$$(5.1) \quad \{\mathcal{G} = (V, E, \mathcal{B}) : \exists \text{ Borel } c: V \rightarrow [k] \text{ such that } \forall \{x, y\} \in E \ c(x) \neq c(y)\},$$

can be expressed by first assuming $(V, \mathcal{B}) = [0, 1]$ and then encoding the parameters E and c as real numbers. In particular, the set from (5.1) is defined using $\exists\forall$ quantification over \mathbb{R} . Such sets are called Σ_2^1 sets. The class of Σ_2^1 sets gives a natural upper bound on the complexity of codes for the existence of graph colorings on Borel graphs. Analogously to **NP**-completeness, if both quantifiers $\exists\forall$ in the definition of a set are necessary, then we say that it is Σ_2^1 -complete. Surprisingly, there is an analogue of the aforementioned result of Karp in measurable combinatorics. In fact, the following result implies a strong failure of Brooks'-like theorems in the Borel context.

Theorem 5.4 ([25, 118]). *Let $k \geq 3$. The set of Borel graphs of Borel chromatic number at most k is Σ_2^1 -complete set. This is true even if we restrict to the class of k -regular acyclic Borel graphs. In contract, Borel graphs admitting a Borel 2-coloring form a Π_1^1 set.*

It is worth mentioning that in the past couple of years a number of further results have been shown concerning descriptive complexity of combinatorial notions, revealing interesting connections particularly to the complexity theory of constraint satisfaction problems (CSPs).

Note that graph coloring problems are special cases of homomorphism problems between relational structures (also called CSPs), for example, a k -coloring of a graph is just a homomorphism to the complete graph K_k . In the finite world, this kind of homomorphism problems are well investigated: a major recent breakthrough is that deciding the existence of a homomorphism to any H is either in P or NP-complete, [21, 29, 124]. In fact, a complete algebraic understanding of structures H for which the homomorphism problem is easy/hard has been reached.

Thornton [116] has initiated a systematic study of the complexity of Borel homomorphism problems of the following sort: given a finite relational structure H , decide the descriptive complexity of Borel structures (with the same signature) that admit a Borel homomorphism to H . Interestingly, if the H homomorphism problem is known to be NP-complete, then the corresponding Borel collection is Σ_2^1 -complete. However, recent results of the authors [73] show a stark contrast

with the finite world: if H corresponds to solving systems of linear equations over a finite field—this problem is in P in the finite case, by Gaussian elimination—the Borel problem is still Σ_2^1 -complete.

A different, wide open question is the investigation of the complexity of *hyperfinite* Borel graphs, see, e.g., [82, 85]. Note that establishing that hyperfinite Borel graphs form a Σ_2^1 set would yield a negative answer to some of the most important conjectures concerning this concept. In this direction, the first author and Higgins [70] have shown that deciding whether a Borel graph has *Borel asymptotic dimension* 1, which is a strengthening of hyperfiniteness, is Σ_2^1 -complete. Further, surprisingly, it turns out [55] that Σ_2^1 -completeness of hyperfiniteness also follows from a negative answer to the *increasing union conjecture* [82], that is, the statement that hyper-hyperfinite equivalence relations are hyperfinite.

Remark 5.5 (Deciding the Global Picture). It is an exciting question to understand for what graph classes is the membership question for LCL problems in Figure 1 itself decidable. That is, after we fix a graph class, is there an algorithm that would take an LCL problem as an input and output its LOCAL complexity? As it is possible to simulate the halting problem using LCL problems on \mathbb{Z}^2 , the membership problem is known to be *undecidable* on graphs that look locally like the grid \mathbb{Z}^d , whenever $d > 1$, see [103] and [27].

The same questions can be asked in the context of measurable combinatorics. That is, after we fix a class of Borel graphs, is there an algorithm that would take an LCL problem as an input and output whether it can be solved in a Borel way on every Borel graph in the class? Our current knowledge here is identical with the one in the LOCAL model. For example, undecidability for LCL problems on Borel graphs that look locally like \mathbb{Z}^d , where $d \geq 1$, has been independently established by Gao, Jackson, Krohne and Seward [59].

5.2. Formal connections. In this section we give a brief overview of the formal connections between distributed computing and measurable combinatorics that were first discovered in the seminal paper of Bernshteyn [12]. Throughout this section, we work with a fixed class of finite graphs that *approximate locally* a class of Borel graphs. One should think of the class of finite graphs of degree bounded by Δ and the class of Borel graphs of degree bounded by Δ , or the class of finite trees and the class of Borel acyclic graphs. Under these mild assumptions, Bernshteyn [12] proved the following, see Figure 1.

Theorem 5.6. *Let Π be a graph coloring problem, or an LCL problem in general.*

- (1) *If the deterministic LOCAL complexity of Π is $o(\log(n))$, then Π admits a Borel solution.*
- (2) *If the randomized LOCAL complexity of Π is $o(\log(n))$, then Π admits a measurable solution.*

Note that this result allows a black-box translations of upper bounds from distributed computing to measurable combinatorics and lower bounds in the opposite directions. We refer the reader to [12, Section 3] to see some of the striking applications of Theorem 5.6.

A natural question arises: can the implications in Theorem 5.6 be reversed? That is, can the existence of a Borel or a measurable solution imply the existence of an efficient distributed algorithm? It turns out that in the setting of deterministic distributed algorithms, the right analogue on the descriptive set theory sides are *continuous solutions*. Gao, Jackson, Krohne and Seward [59] combinatorially characterized what LCL problems can be solved in a continuous way on \mathbb{Z}^d . Their result immediately implies that this class coincide with LCL problems of deterministic LOCAL complexity $o(d\sqrt{n})$ on finite graphs that look locally like \mathbb{Z}^d , see [74, Section 7] for a high-level overview of the argument. This was generalized independently by Bernshteyn [13] and Seward [111] to Cayley graphs of finitely generated groups, and by Brandt, Chang, Grebík, Grunau, Rozhoň and Vidnyánszky [24] to regular trees.

We refer the reader to [17, 23, 58, 72, 74, 80, 92] for further information about the relationship between the complexity classes of LCL problems for particular graph classes and for connections to the theory of random processes, in particular, *finitary factors of iid*.

6. Marks' games

In this section we discuss the adaptation of seminal results of Marks, see Theorem 6.14 and [96], to the LOCAL model. This adaptation has been useful in several different ways. First, it resulted in new proofs of lower bounds for colorings and new results for general LCL problems on bounded degree trees in the LOCAL model. Second, a certain technical difficulty pointed towards introducing the main novelty to be discussed, namely the *homomorphism graph/ID graph technique* in the LOCAL model. Third, a transfer of this technique to descriptive combinatorics found a number of new applications there.

As discussed in Section 4, Marks has shown that the greedy upper bound for vertex and edge coloring of Δ -regular graphs is sharp in the Borel context, even in the case of acyclic graphs. His key idea was to define two player games, where, by strategy stealing arguments, the existence of winning strategies for any of the players contradicts the existence of Borel colorings. In what follows, we will only describe these techniques in the case of vertex colorings, nevertheless, they can be used in a much wider context, to rule out for example matchings or homomorphisms. These generalizations are often non-trivial and require additional technical work, e.g., considering hypergraphs instead of graphs and defining more complicated games, see [25, 96, 117].

Let us also mention that the first lower bounds for deterministic vertex colorings of trees of degree bounded by Δ were proved in [26], and later Brandt [22] has formalized these ideas in his celebrated *round elimination method*. Roughly speaking, what the round elimination method does is the following: assuming that a labeling problem Π on regular trees is solvable by an r -round algorithm, it constructs a new problem Π' solvable by an $r - 1$ -round algorithm. Now, if Π happens to be a fixed point of the operation $'$, then it must be solved by a 0-round algorithm showing that either Π is trivial, or the original assumption was false.⁸

⁸ See <https://github.com/olidennis/round-eliminator> for a fully automated computation of Π' .

This method has a striking similarity to Marks' technique in certain cases. For example, in the case of so called sinkless orientation problem, which is a fixed point of $'$, round elimination essentially produces strategies for the players in the games similar to what Marks considered. Hence, it would be very tempting to think that these methods are the same, say, for fixed points of the operation $'$. This is not the case in general: there are fixed point LCL problems on trees solvable in the Borel context but unsolvable in the LOCAL model. Nevertheless, understanding the exact nature of this connection is still an exciting open problem, see [92].

The adaptation of Marks method will be illustrated by proving the following.

Theorem 6.1. *There is no deterministic algorithm of complexity $o(\log(n))$ that produces a vertex coloring with Δ colors of acyclic graphs of degree $\leq \Delta$.*

There are three main ideas in the proof, all of which we will discuss in this section. As mentioned above, the most important one comes from [96]:

1. Games and strategy stealing to rule out the colorings.

As we will see, executing this idea runs into certain problems with the uniqueness of the IDs in the LOCAL model and with yielding acyclic graphs in the descriptive case. Let us also mention that this issue is present in Brandt's construction, as well; he overcame it by using probabilistic arguments that do not work in the Borel context.

In [23] the homomorphism graph/ID graph technique was invented, which works in both LOCAL and Borel realms.

2. Homomorphism graphs/ID graphs to enforce injectivity.

Subsequently, it became apparent that homomorphism graphs serve as a general tool to transfer properties of large degree graphs to bounded degree graphs. This has been widely utilized in [25] in the descriptive case, but not yet in the LOCAL model.

3. Formulating transfer principles and using games as voting systems.

Before turning to the proof, let us recall that by Bernshteyn's correspondence theorem (Theorem 5.6) inexistence of Borel solutions to a locally checkable problem automatically implies a corresponding result in the deterministic LOCAL model in $o(\log n)$ -many rounds, so formally, for the below results, we don't have to spell out the finitary adaptations. Despite this, we would like to describe them, since it has been developed in parallel to Bernshteyn's theorem, moreover, we believe that there should be models stronger than deterministic LOCAL (see Problem 9.15), for which these techniques still could be used.

Technical preliminaries. In the rest of this section, T_Δ will stand for the Δ -regular infinite tree and T_Δ^r for an r -neighborhood $B_{T_\Delta}(v, r)$ of any (every) vertex v in T_Δ . As discussed in Section 3 a deterministic r -round LOCAL algorithm on a fixed graph is the same as a map from r -neighborhoods with unique IDs of vertices of the graph. Therefore, sometimes we will consider T_Δ^r as being rooted in v (which does not change the isomorphism type of such graphs), and we will evaluate algorithms on such neighborhoods, yielding a color on the root. Also, we will see that it is possible to prove lower bounds already from the existence of algorithms that map

labeled copies of T_Δ^r to different colors, hence we could restrict our attention to only such maps. Moreover, in the case of vertex colorings, it suffices to check the validity of the algorithm on any pair of adjacent vertices, thus we will typically construct labelings of T_Δ^{r+1} to derive a contradiction.

It will be also convenient in the first part of our discussion to consider maps from labeled trees with labelings that are not necessarily injective. This motivates the following definitions (here “G” comes from generalized).

Definition 6.2. An r -round *GLOCAL* algorithm (on trees with degrees $\leq \Delta$) is a partial function \mathcal{A}_r from (isomorphism classes of) such labeled rooted trees of radius $\leq r$ to some set Σ .⁹

If G is a labeled tree with $\Delta(G) \leq \Delta$, we can *run* \mathcal{A}_r on G : at any vertex v of G we take the labeled isomorphism class of $B_G(v, r)$ and feed it into \mathcal{A}_r . This gives a partial vertex labeling of G .

Definition 6.3. Let \mathcal{F} be a collection of finite labeled trees of degrees $\leq \Delta$. An r -round *GLOCAL* algorithm \mathcal{A}_r *vertex colors* \mathcal{F} , if for any $G \in \mathcal{F}$, if \mathcal{A}_r is run on G in the above sense, it yields a vertex coloring of G (in particular, every vertex must get a color).

6.1. Games. We describe the most straightforward finitary adaptation of the original proof of Marks. This, in itself is not suitable to exclude colorings in the LOCAL model, as the labels come from the set $\{0, 1, 2, 3\}$, and hence results only in a less interesting statement, but still encompasses the main idea. For the sake of simplicity, we will consider labeled T_3^{r+1} ’s with labels $\in \{0, 1, 2, 3\}$, so that *neighboring vertices must have different labels*, i.e., the labeling is injective on the edges. Denote this collection by ET_3^{r+1} . This can be viewed as a simplified version of the LOCAL model, in which the IDs are not unique, and the algorithm does not know the size of the graph.

Theorem 6.4. Let r be arbitrary. There is no r -round *GLOCAL* algorithm to 3-color ET_3^{r+1} .

Proof. Towards contradiction, assume that such an r -round *GLOCAL* algorithm \mathcal{A}_r exists. For every $v \in \{0, 1, 2, 3\}$ and $i \in \{0, 1, 2\}$ define the game $\mathbb{G}(v, i)$ as follows. Two players, I and II label with numbers $\in \{0, 1, 2, 3\}$ the r -neighborhood of the rooted 3-regular tree from the root, according to Figure 2, in rounds, alternatingly.

The root is initially labeled by v . In the n th round I and II label all the vertices having distance n from the root. In the first round I labels one of the neighbors of the root, then II labels the other two neighbors. Then I continues with labelling the unlabelled neighbors of her first move (i.e., her “side”), then II with his, etc. The only rule which have to be maintained is that neighboring vertices must get different labels along the game. After r -rounds they label an r -neighborhood of the root, thus the following winning condition is well defined:

⁹ Note that in the definition of *GLOCAL*, we do not require the algorithm to always have an output. However, in practice, when we say that such an algorithm exists on some graph class, it necessarily outputs colors on vertices of those graphs. Thus, alternatively, we could require that \mathcal{A}_r is defined on every isomorphism class, but only correct on the class of graphs in question.

I wins in $\mathbb{G}(v, i)$ iff \mathcal{A}_r evaluated on the labeled r -neighborhood obtained by the end of the game gives a color different from i .

The basic strategy stealing observation is the following.

Lemma 6.4.1. *For every v there is an i so that I has no winning strategy in $\mathbb{G}(v, i)$.*

Proof. Playing these winning strategies, $(\sigma_i)_{i \in \{0,1,2\}}$, against each other (in the obvious manner, by “rotating them”, see Figure 3) would yield in a partial labeling of T_3^r on which \mathcal{A} does not output any color. \square

Now, by the pigeonhole principle, there are distinct $v, w \in \{0, 1, 2, 3\}$ and an i so that II has winning strategies in $\mathbb{G}(v, i)$ and $\mathbb{G}(w, i)$, denote these by σ and τ . Play σ and τ against each other so that for σ we pretend that I ’s first move was w ,




FIGURE 2. The game $\mathbb{G}(v, i)$




FIGURE 3. Strategy stealing

and for τ we pretend that it was v in the first round. Then, in the second round we use the response from τ as I 's play for σ , and similarly for τ we use σ 's response, etc (see again Figure 3).

This way we get a partial labeling of T_3^{r+1} , extend it in an arbitrary manner to an element of EI_3^{r+1} . Now, evaluating \mathcal{A}_r on the central nodes labeled by v and w , we must obtain the color i at both vertices, since σ and τ are winning strategies, contradicting that \mathcal{A}_r correctly solves the 3-coloring problem. \square

6.2. Homomorphism graphs. Of course, the key question now is how to use these ideas to obtain results about the ordinary LOCAL model. The reader can convince themselves that adding rules to the game in the most naïve way (e.g., requiring one of the players to use injective IDs) ruins the symmetry on which the strategy stealing arguments rely.

Observe that the above restrictions on the labeling can be stated equivalently as requiring the labeling to give a homomorphism from T_3^r to K_4 , the complete graph on 4 vertices, and this, restricted injectivity can be maintained throughout our games. One might replace K_4 with any other graph, motivating the following definition.

Definition 6.5. Let $\text{Hom}(T_\Delta^r, \mathcal{H})$ be the set of labelings of T_Δ^r that are homomorphisms to \mathcal{H} .

Now, of course, requiring homomorphisms alone cannot guarantee injectivity. However, adding edge labeling helps: instead of the Δ -regular tree consider it equipped with a proper edge Δ -coloring and assume that \mathcal{H} is equipped with a Δ edge labeling (which is typically not a proper coloring). To emphasize that we consider edge colored version of T_Δ , we denote it by T'_Δ . The next simple observation is the key.

Proposition 6.6. Assume that $\varphi: T'_\Delta \rightarrow \mathcal{H}$ an edge label preserving homomorphism and \mathcal{H} is acyclic. Then φ is injective. Similarly, for any $k > 0$ if the girth of \mathcal{H} is $> 2k + 2$ then φ is injective in every $k + 1$ -neighborhood.

Proof. To see the first statement, note that since its range is acyclic, if φ was not injective then there would be some (u, v) and (v, w) distinct edges in T'_Δ so that $\varphi(u) = \varphi(w)$. But as (u, v) and (v, w) are adjacent, their color is different, so the color of their image must also be different, contradicting $\varphi(u) = \varphi(w)$. The proof of the second statement is identical. \square

Thus, it is useful to consider the next definition.

Definition 6.7. Assume that the graph \mathcal{H} is equipped with a Δ edge labeling. Let $\text{Hom}^e(T_\Delta^r, \mathcal{H})$ be the set of labelings of $B_{T'_\Delta}(v, r)$ (for any v in T'_Δ) that are homomorphisms to \mathcal{H} and preserve edge labels.

A way to view the graph \mathcal{H} is as a “set of rules”, or a “game plan”, for the players in the game, as we will soon see.

6.3. Transfer principles. Now, we can summarize the intuition which guides us. We take a suitable graph \mathcal{H} , and consider all the possible homomorphism labelings of the tree T_Δ to \mathcal{H} . We would like to prove that there is no “local” coloring rule to Δ -color the resulting graphs. Assume that there is such a rule. For each $v \in V(\mathcal{H})$ we can label the tree in a lot of ways so that the root is labeled by v . Of course, all these labeled rooted trees get a color, according to the rule. We use the game to decide the most popular color i , that is, when II has a winning strategy in $\mathbb{G}(v, i)$. It turns out, that this transfers the original coloring to a coloring of \mathcal{H} .

Let us state the transfer principle for the distributed case.

Theorem 6.8 (Transfer principle 1, LOCAL). *Let \mathcal{H} be a finite graph. TFAE:*

- (1) \mathcal{H} admits a Δ -coloring.
- (2) There exists a 0-round GLOCAL algorithm to Δ -color $\mathbf{Hom}(T_\Delta^r, \mathcal{H})$, for all r .
- (3) There exist an r and an r -round GLOCAL algorithm to Δ -color $\mathbf{Hom}(T_\Delta^{r+1}, \mathcal{H})$.

Proof. To see that (1) implies (2), assume that \mathcal{H} admits a Δ -coloring. Given a vertex in some element of $\mathbf{Hom}(T_\Delta^r, \mathcal{H})$, it has a label from \mathcal{H} . Color the vertex with the color of its label.

Clearly, (2) \implies (3), so it suffices to show that (3) implies (1), so assume that $\mathbf{Hom}(T_\Delta^{r+1}, \mathcal{H})$ admits an r -round GLOCAL algorithm \mathcal{A} . For each color i and vertex $v \in V(\mathcal{H})$ consider the game $\mathbb{G}(v, i)$ defined as follows: two players build a neighborhood of a vertex of some element in $\mathbf{Hom}(T_\Delta^r, \mathcal{H})$, according to Figure 2, where the vertex is labeled by v and the two players have to maintain that the labeling is a homomorphism to \mathcal{H} at every step.

I wins iff \mathcal{A} evaluated on the labeled r -neighborhood obtained by the end of the game gives a color different from i . \square

Lemma 6.8.1. *For every v there is an i so that I has no winning strategy in $G(v, i)$.*

Proof. Playing these winning strategies against each other would yield in a labeling of T_Δ^r on which \mathcal{A} does not output any color. \square

The next lemma finishes the proof of the transfer principle.

Lemma 6.8.2. *For a $v \in V(\mathcal{H})$ let*

$$c(v) = \min\{i : \mathbb{G}(v, i) \text{ is won by } II\}.$$

Then c is Δ -coloring of \mathcal{H} .

Proof. Assume that this is not the case, i.e., there exists some $(v, w) \in \mathcal{H}$ so that $c(v) = c(w)$. Then II has winning strategies σ and τ in the games $\mathbb{G}(v, i)$ and $\mathbb{G}(w, i)$. Now play σ and τ against each other so that for σ we pretend that I ’s first move was w and for τ we pretend that it was v in the first round. Then in the second round we use the response from τ as I ’s play for σ and similarly for τ we use σ ’s response, etc.

This way we get a partial labeling of T_{Δ}^{r+1} , extend this in an arbitrary manner to obtain an element of $\mathbf{Hom}(T_{\Delta}^{r+1}, \mathcal{H})$. Now, evaluating \mathcal{A} on the central nodes labeled by v and w , we must obtain the color i at both vertices, since σ and τ are winning strategies, contradicting that \mathcal{A} solves the Δ -coloring problem. \square

This transfer principle alone is insufficient to get results for the LOCAL model, since the labelings typically are not injective. To remedy this, we have to consider edge labeled versions of the homomorphism graphs. For technical reasons (namely, in order for the players not to get stuck in the game), we assume that the edge labeling of \mathcal{H} is *nice*, that is, every vertex is adjacent to edges with all possible colors.

Definition 6.9. Assume that \mathcal{G} is a graph equipped with an edge Δ -labeling e . An *edge labeled Δ -coloring* of \mathcal{G} is a map $c: V(\mathcal{G}) \rightarrow \Delta$ so that there are no adjacent vertices v, w with $c(v) = c(w) = e(v, w)$. If such a coloring exists, we will say that the *edge labeled chromatic number of \mathcal{G} is $\leq \Delta$* , and denote this fact by $\chi^{el}(\mathcal{G}) \leq \Delta$.

Theorem 6.10 (Transfer principle 2, LOCAL). *Let \mathcal{H} be a finite graph together with a nice edge Δ -labeling. TFAE:*

- (1) $\chi^{el}(\mathcal{H}) \leq \Delta$.
- (2) *there is a 0-round GLOCAL algorithm producing a Δ -edge labeled coloring of $\mathbf{Hom}^e(T_{\Delta}^r, \mathcal{H})$, for every r .*
- (3) *there is an r and an r -round GLOCAL algorithm producing a Δ -edge labeled coloring of $\mathbf{Hom}^e(T_{\Delta}^{r+1}, \mathcal{H})$.*

Proof. The proof is a straightforward modification of the proof of Theorem 6.8, in particular, the implications (1) \implies (2) \implies (3) follow immediately.

To see that (3) \implies (1), for any $v \in \mathcal{H}$ and $i \leq \Delta$ define the game $\mathbb{G}(v, i)$ as in the proof of Theorem 6.8, with the additional assumption that the edge connecting the vertex labeled by v and I 's first move has to be colored i .

It is easy to check that the arguments in Lemma 6.8.1 and Lemma 6.8.2 still work, and c is an edge labeled Δ -coloring of \mathcal{H} . \square

6.3.1. Back to LOCAL. Now we are ready to go back to the ordinary LOCAL model, and deduce Theorem 6.1. Let us make a simple observation first.

Claim 6.10.1. Assume that \mathcal{F} is a collection of labeled trees of degree $\leq \Delta$ and $f: \mathbb{N} \rightarrow \mathbb{N}$ such that for each $G \in \mathcal{F}$ with $|V(G)| = n$

- (1) the set of labels has size $\leq n$,
- (2) if n is large enough, the labeling is injective on every $f(n) + 1$ -sized neighborhood of G .

Assume moreover that there is no $g \in o(f)$ so that there is a $g(n)$ -round GLOCAL algorithm to Δ -color every n -sized $G \in \mathcal{F}$. Then there is no $o(f)$ -round deterministic distributed algorithm to Δ -color every n -sized $G \in \mathcal{F}$.

Proof. If there was such a LOCAL algorithm, by the local injectivity of the labeling, we could interpret the labels as IDs and run the algorithms on the elements of \mathcal{F} . \square

The following Lemma guarantees the existence of graphs that satisfy simultaneously the negation of (1) in Theorem 6.10 and Proposition 6.6. The proof is based on a probabilistic method.

Lemma 6.10.1. *Let $f \in o(\log n)$. There exists a sequence of nicely Δ edge labeled graphs (\mathcal{H}_n) with the following properties:*

- (1) $|V(\mathcal{H}_n)| \leq n$.
- (2) $\chi^{el}(\mathcal{H}_n) > \Delta$.
- (3) the girth of \mathcal{H}_n is $> 2f(n) + 2$.

Proof sketch. Such a sequence can be constructed probabilistically, using the *configuration model* for regular random graphs, see [123]. Namely, on the vertex set $\{1, 2, \dots, n\}$ we independently choose Δ many d -regular random graphs and let \mathcal{H}_n to be the union of those, for a large constant d . Note that this gives a canonical edge labeling of \mathcal{H}_n with Δ colors with the property that every vertex is adjacent to exactly d edges of each of the colors.

By the results of Bollobás [19] on the independence ratio of random regular graphs and concentration inequalities of McDiarmid [100] it can be shown that $\chi^{el}(\mathcal{H}_n) > \Delta$ asymptotically almost surely as $n \rightarrow \infty$. Finally, seeing the sample \mathcal{H}_n as a random $d\Delta$ -regular graph, one can use [101, Corollary 1] to show that (3) is satisfied with positive probability. \square

Now we can prove the main theorem on colorings.

Proof of Theorem 6.1. Towards a contradiction, assume that there exists a LOCAL Δ -coloring algorithm for acyclic graphs of degree $\leq \Delta$ with running time $g \in o(\log n)$. We can assume that g is monotone. Let $(\mathcal{H}_n)_{n \in \mathbb{N}}$ be the sequence guaranteed by Lemma 6.10.1 for the function $f(n) = g(\Delta n)$.

Consider the family of graphs $\mathcal{F} = \bigcup_n \mathbf{Hom}^e(T_{\Delta}^{r_n}, \mathcal{H}_n)$, where r_n is chosen minimal so that $|V(T_{\Delta}^{r_n})| \geq n$. \square

Claim 6.10.2. For every graph in $\mathbf{Hom}^e(T_{\Delta}^{r_n}, \mathcal{H}_n)$ the labeling is injective in every $f(n) + 1$ -sized neighborhood.

Proof. Indeed, observe that, as the homomorphism defined by the labels has to preserve edge colors and the girth of \mathcal{H}_n is $> 2f(n) + 2$, it must be injective on every $f(n) + 1$ sized neighborhood by Proposition 6.6.

By the minimality of r_n , we have that $|V(T_{\Delta}^{r_n})| \leq \Delta n$. Thus, the labeling of every graph in $\mathbf{Hom}^e(T_{\Delta}^{r_n}, \mathcal{H}_n)$ is injective in every neighborhood of size

$$f(n) + 1 = g(\Delta n) + 1 \geq g(|V(T_{\Delta}^{r_n})|) + 1,$$

showing that the conditions of Claim 6.10.1 are satisfied. Hence, on the one hand, there is an $o(g)$ -time GLOCAL algorithm for the Δ -coloring of \mathcal{F} .

On the other hand, $r_n > c \log n$, for some $c > 0$ and, as $g \in o(\log n)$, for a large enough n we have $g(n) < c \log n < r_n$. Therefore, by assumption, there is a $g(n) < r_n$ round GLOCAL algorithm to Δ -color $\mathbf{Hom}^e(T_{\Delta}^{r_n}, \mathcal{H}_n)$, which is in particular an $r_n - 1$ -round GLOCAL algorithm. By Theorem 6.10 this implies $\chi^{el}(\mathcal{H}_n) \leq \Delta$, contradicting the choice of \mathcal{H}_n . \square

6.3.2. A localized transfer. As we will see in the descriptive case in the next section (Theorem 6.18), even descriptive combinatorial properties can be transferred through homomorphism graphs. In order to mimic this in the distributed version, we have to consider “localized” versions of the graph \mathcal{H} instead, i.e., instead of a fixed enumeration, we consider all the possible ID distributions on \mathcal{H} .

Definition 6.11. Let $\mathbf{Hom}(T_\Delta^r, \mathcal{H}^{ID})$ be the set of labelings of T_Δ^r that are the pullbacks of IDs (i.e., injective labelings with naturals $< |V(\mathcal{H})|$) on \mathcal{H} through a homomorphism.

Similarly, if \mathcal{H} is equipped with an edge Δ -labeling, let $\mathbf{Hom}^e(T_\Delta^r, \mathcal{H}^{ID})$ be the set of labelings of T_Δ^r that are the pullbacks of IDs on \mathcal{H} through an edge label preserving homomorphism.

Now we have the following.

Theorem 6.12 (Transfer principle 3, LOCAL). *Let \mathcal{H} be a finite graph.*

- (1) *If $\mathbf{Hom}(T_\Delta^{r+1}, \mathcal{H}^{ID})$ has an r -round Δ -coloring GLOCAL algorithm, then \mathcal{H} has an r -round Δ -coloring in the LOCAL model.*
- (2) *Assume that \mathcal{H} is equipped with a nice edge labeling. If $\mathbf{Hom}^e(T_\Delta^{r+1}, \mathcal{H}^{ID})$ has an r -round GLOCAL algorithm for edge labeled Δ -coloring, then \mathcal{H} has an r -round edge labeled Δ -coloring in the LOCAL model.*

Proof. The proofs are the same as the proofs of the (3) \implies (1) implications of Theorems 6.8 and 6.10. The only thing to observe is that the coloring defined as in Lemma 6.8.2 in fact only depends on the IDs in an r -round neighborhood of v in \mathcal{H} , hence it is doable in the LOCAL model. \square

The downside here that there is no hope for an equivalence in general, as the existence of a LOCAL algorithm on \mathcal{H} can happen for trivial reasons: for example if \mathcal{H} is star-like, with constant round LOCAL algorithm one can access global information. However, this information is not present on the homomorphism graph, because of the degree bound.

Remark 6.13. We conjecture that using this transfer principle one can give a proof of Theorem 6.1 without having to go through Lemma 6.10.1.

6.4. The descriptive case. To close our discussion of Marks’ method, let us mention that statements completely parallel to the ones described in the preceding section have been shown in the Borel context in [25, 96].

Marks initially proved the following. Consider the group

$$\Gamma_\Delta = \langle \alpha_1, \dots, \alpha_\Delta : \alpha_1^2 = \dots = \alpha_\Delta^2 = 1 \rangle,$$

that is, Γ_Δ is the free product of Δ -many involutions. Observe that $\text{Cay}(\Gamma_\Delta, S)$ is a Δ -regular tree, where $S = \{\alpha_1, \dots, \alpha_\Delta\}$.

Theorem 6.14. *Let $\Delta > 2$. Then there exists a Δ -regular acyclic Borel graph with $\chi_B(G) = \Delta + 1$. In fact, this holds for the Schreier graph of the free part of the action of Γ_Δ on 2^{Γ_Δ} .*

In the proof of this theorem, the players play on the infinite tree, games analogous to the ones described above. Interestingly, the existence of winning strategies is not at all clear: every acyclic Δ -regular Borel graph has a μ -measurable Δ -coloring by the results of Conley-Marks-Tucker-Drob [42] for any Borel measure μ , described above, showing that winning strategies might not exist even for Lebesgue-measurable sets. However, since the winning conditions in the games are Borel, to guarantee their existence, one can use Martin's Borel Determinacy theorem, one of the cornerstones of descriptive set theory.

The problem of non-injectivity of IDs in the Borel context corresponds to the players having to end up in the free part of the above action. Marks remedied this by arguing that there exists a Borel Δ -color assignment c' of the Schreier graph on the non-free part, with the property that for no x and i we have $c'(x) = c'(\gamma_i \cdot x) = i$. That is, χ_B^{el} (defined analogously to Definition 6.9) of the not-free part of the Schreier graphs is $\leq \Delta$. So, if there exists a Borel Δ -coloring of the free part then $\chi_B^{el}(Sch(\Gamma, 2^{\Gamma_\Delta})) \leq \Delta$. Now, we can derive a contradiction similarly to the proof of Theorem 6.10.

The homomorphism graph can be defined in the Borel context as follows: observe that if \mathcal{H} is a Borel graph, $V(\mathcal{H})^{\Gamma_\Delta}$ is a Borel space on which Γ_Δ acts by left-shift.

Definition 6.15. Let \mathcal{H} be a Borel graph. Let $\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})$ be the restriction of $Sch(\Gamma_\Delta, S, V(\mathcal{H})^{\Gamma_\Delta})$ to the set

$$\{h \in V(\mathcal{H})^{\Gamma_\Delta} : h \text{ is a graph homomorphism from } Cay(\Gamma_\Delta, S) \text{ to } V(\mathcal{H})\}.$$

Similarly to Lemma 6.8.2, one can define a coloring on \mathcal{H} by using a “voting system” based on the information about the winning strategies in the games analogous to the ones discussed above. Since this requires to know which player has a winning strategy, the resulting coloring will be slightly worse than Borel, namely, in the class $\mathfrak{D}\Delta_1^1$. Moreover, assuming extra set theoretic axioms and considering a slightly larger class, we can even obtain a full equivalence (below, χ_{pr} stands for projective chromatic number):

Theorem 6.16 ([25], transfer principle 1, Borel). *Let \mathcal{H} be a Borel graph.*

- (1) $\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})$ admits a Borel homomorphism into \mathcal{H} . Thus, for any n , if $\chi_B(\mathcal{H}) \leq n$ then $\chi_B(\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})) \leq n$.
- (2) If $\chi_{\mathfrak{D}\Delta_1^1}(\mathcal{H}) > \Delta$ then $\chi_B(\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})) > \Delta$.

Assuming suitable large cardinal hypotheses we also have

$$\chi_{pr}(\mathcal{H}) > \Delta \iff \chi_{pr}(\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})) > \Delta.$$

Finally, observe that $Cay(\Gamma_\Delta, S)$ is equipped with a proper edge coloring. Hence, the edge labeling trick is sufficient to produce acyclic graphs here as well.

Definition 6.17. Assume that the graph \mathcal{H} is equipped with a Borel edge Δ -labeling. Let $\mathbf{Hom}^e(\Gamma_\Delta, \mathcal{H})$ be the restriction of $\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})$ to the set

$$\{h \in V(\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})) : h \text{ preserves the edge labels}\}.$$

In case \mathcal{H} is acyclic, an argument similar to the proof of Proposition 6.6 yields that each $h \in \mathbf{Hom}(\Gamma_\Delta, \mathcal{H})$ is injective. Using this, one can show the following.

Theorem 6.18 (Transfer principle 2, Borel). *Assume that \mathcal{H} is a locally countable Borel graph and e is a Borel edge Δ -labeling of \mathcal{H} .*

- (1) $\mathbf{Hom}^e(\Gamma_\Delta, \mathcal{H})$ admits a label preserving Borel homomorphism into \mathcal{H} .
Thus, if $\chi_B^{el}(\mathcal{H}) \leq \Delta$ then $\chi_B(\mathbf{Hom}^e(\Gamma_\Delta, \mathcal{H})) \leq \Delta$.
- (2) *If \mathcal{H} is acyclic then so is $\mathbf{Hom}^e(\Gamma_\Delta, \mathcal{H})$.*
- (3) *If \mathcal{H} is acyclic and hyperfinite then so is $\mathbf{Hom}^e(\Gamma_\Delta, \mathcal{H})$.*
- (4) *If $\chi_{\Delta_1}^{el}(\mathcal{H}) > \Delta$ then $\chi_B(\mathbf{Hom}^e(\Gamma_\Delta, \mathcal{H})) > \Delta$.*

Assuming suitable large cardinal hypotheses we also have

$$\chi_{pr}^{el}(\mathcal{H}) > \Delta \iff \chi_{pr}^{el}(\mathbf{Hom}(\Gamma_\Delta, \mathcal{H})) > \Delta.$$

Since a suitable edge labeled version of $\mathcal{H} = \mathbb{G}_0$ (see [25, 88]) is hyperfinite and $\chi_{\Delta_1}^{el}(\mathcal{H}) > \Delta$, Theorem 6.18 yields Theorem 6.14 and can be also used to obtain the results in [39].

7. Versions of Vizing's theorem

In this section we describe the ideas behind the proof of the measurable version of Vizing's theorem and their applications to the LOCAL model of distributed computing.

The strategy to construct edge colorings both in measurable and distributed context is to gradually improve a partial coloring using augmenting subgraphs. To get a feeling about this type of construction in measurable combinatorics, we start by revisiting the result of Lyons and Nazarov [94], where they produced a *perfect matching on the Δ -regular tree as a factor of iid*. Their overall approach, that originated in the work of Nguyen and Onak [104] that was later formalized by Elek and Lippner [51], turned out to be extremely important not only for edge colorings, but also for applications that include for example the measurable circle squaring, or paradoxical decompositions, see [41, 66, 67].

The argument that Lyons and Nazarov employed to show that the final matching, after the augmenting procedure that might take an infinite number of steps, is well-defined almost surely is conceptually different than the one for the measurable Vizing's theorem. Namely, their argument uses the fact that the underlying measured graph, iid on Δ -regular tree, *expands in measure* which is a form of *global expansion*. Since not every acyclic graphing expands in measure, and interestingly Kun [89] found examples of Δ -regular acyclic graphings that do not admit measurable perfect matching, we rely on a different tool in the proof of measurable Vizing's theorem: a form of *local expansion* introduced in [71] that was later generalized in the notion of a *multi-step Vizing chain* by Bernshteyn [11]. This type of local expansion have proved to be particularly well-suited for the construction of distributed algorithms [11, 14, 36].

After discussing the important ideas behind the result of Lyons and Nazarov [94] in the next subsection, we formally introduce the concept of multi-step Vizing chains. In the last part, we summarize the applications that include measurable Vizing's theorem and the construction of distributed algorithms of LOCAL complexity $\text{poly}(\log(n))$.

7.1. Perfect matching on regular trees. Recall that a *matching* in a graph $G = (V, E)$ is a collection of edges $M \subseteq E$ such that for every $v \in V$ there is at most one $e \in M$ such that $v \in e$. We denote as U_M the set of *unmatched vertices of M* , that is, $v \in U_M$ if and only if $v \notin e$ for every $e \in M$. Vertices that are not unmatched are *covered by M* . We say that M is a *perfect matching* if $U_M = \emptyset$.

A classical result of Kőnig states that a bipartite Δ -regular graph G satisfies $\chi'(G) = \Delta$. In particular, G admits a perfect matching. Analogues of Kőnig's theorem in measurable combinatorics are simply false. First, Laczkovich [90] gave an example of a 2-regular measurably bipartite graphing that does not admit a measurable perfect matching. Then Marks [96] showed that there are acyclic Δ -regular Borel bipartite Borel graphs that do not admit Borel perfect matching for every $\Delta \geq 3$. Finally, Kun [89] strengthened Marks' result by finding an acyclic Δ -regular measurably bipartite graphing that does not admit measurable perfect matching for every $\Delta \geq 3$.

On the positive side, Lyons and Nazarov [94] showed that there is a factor of iid perfect matching on the Δ -regular tree, which implies, by the result of Hatami, Lovász and Szegedy [78], that *every* acyclic Δ -regular graphing admits a measurable matching such that the measure of unmatched vertices is arbitrarily small.

In our exposition of the arguments from [94] we follow the approach of Elek and Lippner [51]. Basically all measurable constructions that are based on augmenting subgraph technique use some version of Theorem 7.2 in the background.

Definition 7.1. Let M be a matching in $G = (V, E)$. We say that

$$P = (v_1, v_2, \dots, v_k) \subseteq V$$

is an *(M -)augmenting path* if $v_1, v_k \in U_M$ and $\{v_{2i}, v_{2i+1}\} \in M$ for every $i \in \{1, \dots, k/2 - 1\}$.

Flipping the edges on an augmenting path $P = (v_1, v_2, \dots, v_k)$ modifies the matching M to a matching M' that agrees with M outside of P and such that $\{v_{2i-1}, v_{2i}\} \in M'$ for every $i \in \{1, \dots, k/2\}$.

Note that the length of every augmenting path P is even and that all vertices of P are covered after flipping the edges.

We state a version of the result of Elek and Lippner [51], which is suitable in our context.

Theorem 7.2 (Elek–Lippner [51]). *Let $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ be a graphing of degree bounded by $\Delta \in \mathbb{N}$, M be a measurable matching in \mathcal{G} and $k \in \mathbb{N}$. Then there is a measurable matching $M' \subseteq E$ such that*

$$(7.1) \quad \mu(\{v \in V : \exists w \in N_{\mathcal{G}}(v) \{v, w\} \in M' \setminus M\}) \leq k\mu(U_M)$$

and every M' -augmenting path has length at least k .

Proof sketch. Let $\mathcal{G}^{(2k+1)}$ denote the $2k+1$ st power of \mathcal{G} , that is, $x, y \in V$ form an edge in $\mathcal{G}^{(2k+1)}$ if their graph distance in \mathcal{G} is at most $2k+1$. It is easy to check that $\mathcal{G}^{(2k+1)}$ is a graphing of bounded maximum degree. By a straightforward application of Theorem 4.6, there is a sequence $(B_i)_{i \in \mathbb{N}}$ of Borel $\mathcal{G}^{(2k+1)}$ -independent sets such that $\{i \in \mathbb{N} : v \in B_i\}$ is infinite for every $v \in V$.

Set $M_0 = M$ and inductively define a sequence of measurable matchings $(M_i)_{i \in \mathbb{N}}$ as follows. Suppose that M_i has been defined. If possible pick an M_i -augmenting path P_v in \mathcal{G} of length at most k that starts at v for every $v \in B_i$. This assignment can be done in Borel way. Note that if $v \neq w \in B_i$, then P_v and P_w are edge disjoint. In particular, we may define the measurable matching M_{i+1} by flipping the edges on all of these M_i -augmenting paths simultaneously.

We claim that $M' = \lim_{i \rightarrow \infty} M_i$ is a well-defined measurable matching that satisfies the claim of the theorem. Indeed, the limit is well-defined as every edge is flipped only finitely many times. This follows from the combination of the facts that every vertex lies on finitely many paths of length at most k and that once a vertex is covered by M_i , then it is covered by M_j for every $i \leq j \in \mathbb{N}$. Consequently, if P is a finite M' -augmenting path starting at $v \in V$, then there is an $i_0 \in \mathbb{N}$ such that P is an M_i -augmenting path for every $i \geq i_0$. In particular, the length of P has to be at least k , as otherwise it would be flipped for any $i \geq i_0$ such that $v \in B_i$.

Finally, (7.1) follows from the mass transport principle as every vertex from U_M is responsible for flipping edges at k -many other vertices. \square

Let $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ be an acyclic Δ -regular graphing and consider the following process. Start with an empty matching M_0 and define inductively M_{k+1} from M_k using Theorem 7.2 so that every M_{k+1} -augmenting path has length at least $2k$ (as our main application, Theorem 7.4, is for bipartite graphs, we use $2k$ instead of just k). This produces a sequence of measurable matchings $(M_k)_{k \in \mathbb{N}}$.

We would like to claim now that the sequence $(M_k)_{k \in \mathbb{N}}$ admits almost surely a pointwise limit $M = \lim_{k \rightarrow \infty} M_k$ that is the desired measurable perfect matching. Recall that the pointwise limit is defined by viewing each M_k as its characteristic function. However, such a limit need not exist as it might be the case that a positive fraction of edges is flipped in the sequence $(M_k)_{k \in \mathbb{N}}$ infinitely often.

Observe that a sufficient condition for the pointwise limit $M = \lim_{k \rightarrow \infty} M_k$ to exist almost surely is if

$$(7.2) \quad \sum_{k=0}^{\infty} 2k\mu(U_{M_k}) < \infty.$$

This follows from a combination of Theorem 7.2 and Borel–Cantelli lemma. Hence, it is enough to find a condition that guarantees that $\mu(U_{M_k}) \in o(1/k^2)$. This is the place where expansion in measure kicks in. To keep things simple, we introduce the concept of measure expansion for measurably bipartite graphs only.

Definition 7.3. Let $\eta > 0$ and $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ be an acyclic Δ -regular graphing that admits a measurable bipartition $V = A \sqcup B$. We say that \mathcal{G} is an η -expander if $\mu(N_{\mathcal{G}}(X)) \geq (1 + \eta)\mu(X)$ for every $X \subseteq V$ such that $X \subseteq A$ or $X \subseteq B$, and $\mu(X) \leq 1/4$.

Now we are ready to state a weak version of the result of Lyons and Nazarov [94].

Theorem 7.4 (Lyons–Nazarov [94]). *Let $\eta > 0$ and $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ be an acyclic Δ -regular measurably bipartite graphing that is an η -expander. Then \mathcal{G} admits a measurable perfect matching.*

Proof sketch. It is enough to verify that the sequence of measurable matching $(M_k)_{k \in \mathbb{N}}$ defined after Theorem 7.2 satisfies $\mu(U_{M_k}) \in o(1/k^2)$.

Observe that as \mathcal{G} is regular and satisfies the mass transport principle, we have that $\mu(A) = \mu(B) = 1/2$. In fact, we have $\mu(U_{M_k} \cap A) = \mu(U_{M_k} \cap B)$ for every $k \in \mathbb{N}$. Set $X_0^k = U_{M_k} \cap A$ and define inductively a sequence of measurable sets $(X_\ell^k)_{\ell \in \mathbb{N}}$ as follows. If ℓ is even, then set $X_{\ell+1}^k = N_{\mathcal{G}}(X_\ell^k) \subseteq B$, and if ℓ is odd, then set $X_{\ell+1}^k = \{v \in V : \exists w \in X_\ell^k \{v, w\} \in M_k\} \subseteq A$. Let $\ell_0 \in \mathbb{N}$ be the first (odd) number such that either $\mu(X_{\ell_0}^k) > 1/4$ or $\mu(X_{\ell_0}^k \cap U_{M_k}) > 0$. This is well-defined as \mathcal{G} is an η -expander.

Observe that for every $0 \leq j \leq (\ell_0 + 1)/2$, we have that

$$(1 + \eta)^{j+1} \mu(X_0^k) \leq \mu(X_{2j+1}^k) = \mu(X_{2j+2}^k)$$

as \mathcal{G} is a graphing and an η -expander. We define analogously a sequence $(Y_\ell^k)_{\ell \in \mathbb{N}}$ and $\ell'_0 \in \mathbb{N}$ for $Y_0^k = U_{M_k} \cap B$. In particular,

$$\mu(U_{M_k}) = \mu(X_0^k) + \mu(Y_0^k) \leq 2(1 + \eta)^{-(\max\{\ell_0, \ell'_0\} + 1)/2}.$$

We finish the proof by showing that $k \leq \max\{\ell_0, \ell'_0\} + 1$

First observe that if $\mu(X_{\ell_0}^k \cap U_{M_k}) > 0$, then there are M_k -augmenting paths of length $\ell_0 + 1$. By symmetry the same holds for $Y_{\ell'_0}^k$. In particular, if $\max\{\ell_0, \ell'_0\} + 1 < 2k$ it must be the case that

$$\mu(X_{\ell_0}^k), \mu(Y_{\ell'_0}^k) > 1/4.$$

But then we conclude that $\mu(X_{\ell_0+1}^k \cap Y_{\ell'_0}^k) > 0$ as $X_{\ell_0+1}^k, Y_{\ell'_0}^k \subseteq A$, $\mu(A) = 1/2$ and $\mu(X_{\ell_0+1}^k) = \mu(X_{\ell_0}^k) > 1/4$. Consequently, this guarantees the existence of M_k -augmenting paths of length $\ell_0 + 1 + \ell'_0 + 1$. Indeed, consider a concatenation of any paths from X_0^k to $X_{\ell_0+1}^k \cap Y_{\ell'_0}^k$ and from $X_{\ell_0+1}^k \cap Y_{\ell'_0}^k$ to Y_0^k . Altogether, this shows $k \leq \max\{\ell_0, \ell'_0\} + 1$ and the proof is finished. \square

Remark 7.5 (Approximate König's line coloring theorem). It was observed, for example, by Lyons [93] that the iterative application of these arguments yield a factor of iid edge coloring with $\Delta + 1$ colors of the Δ -regular tree such that one of the color classes can have arbitrarily small measure. In fact this holds for all bipartite measured graphs of degree bounded by Δ , see [68, 69, 119].

7.2. Multi-step Vizing chains. Our goal is to define the notion of a multi-step Vizing chain and discuss the connections to local expansion. We start by recalling some notation from [71].

Let $G = (V, E)$ be a finite (or countable) connected graph and set $\Delta = \Delta(G) \in \mathbb{N}$. For the rest of this section we fix a partial edge coloring $d: E \rightarrow [\Delta + 1]$. We write $m_d(v)$ for the set of colors missing at $v \in V$, that is, $\alpha \notin m_d(v)$ if there is an $e \in \text{dom}(d)$ such that $v \in e$ and $d(e) = \alpha$. Observe that $|m_d(v)| > 0$ for every $v \in V$.

A *chain* is a sequence $P = (e_0, e_1, \dots) \subseteq E$ of edges such that for every $i \in \mathbb{N}$ with $e_i, e_{i+1} \in P$ we have $e_i \cap e_{i+1} \neq \emptyset$. Let $l(P) = |P|$ denote the number of edges in P , we allow $l(P)$ to be infinite. We say that P is *edge injective* if every

edge appears at most once in P . If P is a finite chain with the last edge e and Q is a chain with the first edge f and $e \cap f \neq \emptyset$, then we write $P \frown Q$ for the chain that is the *concatenation* of P and Q . If a chain R is of the form $P \frown Q$, then we say that P is a *prefix* of R , denoted by $P \sqsubseteq R$.

We say that a chain $P = (e_i)_{i < l(P)}$ is *d-shiftable* if $l(P) \geq 1$, P is edge injective, $e_0 \in U_d$ and $e_j \in \text{dom}(d)$ for every $1 \leq j < l(P)$. A *d-shiftable* chain P is called *d-proper-shiftable* if $d_P: E \rightarrow [\Delta + 1]$, the *P-shift* of d , defined as

- (1) $\text{dom}(d_P) = \text{dom}(d) \cup \{e_0\} \setminus \{e_{l(P)-1}\}$ where we put $\{e_{l(P)-1}\} = \emptyset$ if $l(P) = \infty$,
- (2) $d_P(e_i) = d(e_{i+1})$ for every $i + 1 < l(P)$,
- (3) $d_P(f) = d(f)$ for every $f \in \text{dom}(d) \setminus P$;

is a partial edge coloring. Finally, we say that a *d-proper-shiftable* chain is *d-augmenting* if either $l(P) = \infty$ or P is finite with $m_{d_P}(x) \cap m_{d_P}(y) \neq \emptyset$ where $x \neq y$ are the vertices of the last edge $e_{l(P)-1}$ of P .

A *fan around $x \in V$ starting at $e \in E$* is a finite edge injective chain $F = (e_0, e_1, \dots, e_k)$ such that $x \in e_j$ for every $0 \leq j \leq k$ and $e_0 = e$. If we don't want to specify $x \in V$ and $e \in E$, we simply say that F is a *fan*.

The α/β -*path starting at $x \in V$* , denoted by $P_d(x, \alpha/\beta)$, is the unique maximal chain $P = (e_i)_{i < l(P)}$ such that $d(e_i) = \alpha$ (resp. $d(e_i) = \beta$) for every $i < l(P)$ that is even (resp. odd). A *truncated α/β -path* is a prefix of any α/β -path $P_d(x, \alpha/\beta)$. If we don't want to specify the colors in the path, we simply say that $P = P_d(x, \alpha/\beta)$ is an *alternating path*.

Definition 7.6 (Vizing Chain). Let $e \in U_d$ and $x \in e$. A *Vizing chain* $W(x, e)$ is a *d-proper-shiftable* chain of the form $W(x, e) = F_e \frown P_{\alpha_e/\beta_e}$, where F is a fan at x starting at e and P is an α_e/β_e -path for some $\alpha_e \neq \beta_e \in [\Delta + 1]$, see Figure 4.

Recall that the standard proof of Vizing's theorem can be viewed as a sequential algorithm that at each step finds a Vizing chain $W(x, e)$, for a given uncolored edge e , that is *d-augmenting* and improves the current coloring d by passing to $d_{W(x, e)}$




FIGURE 4. An example of a Vizing chain $W(x, e) = F_e \frown P_{\alpha_e/\beta_e}$.

and coloring the last edge of $W(x, e)$. In particular, we get that a d -augmenting Vizing chain exists for every uncolored edge.

Next, we present a general definition of a multi-step Vizing chain and discuss some of the properties related to local expansion. The high-level idea is to turn the algorithm for finding a Vizing chain into a branching process by repeatedly truncating the alternating path and growing a new Vizing chain after shifting the colors.

Definition 7.7 (Multi-step Vizing chain, [11, 36]). Let $e \in U_d$, $x \in e$ and $i \in \mathbb{N}$. An i -step Vizing chain $V(x, e)$ is a d -proper-shiftable chain of the form

$$V(x, e) = F_1 \cap P_1 \cap \dots \cap F_i \cap P_i,$$

where F_1 is a fan at x starting at e and for every $1 \leq j \leq i$ we have that F_j is a fan and P_j is a truncated alternating path, see Figure 5.

The concept was first introduced for $i = 2$ by the first author and Pikhurko in [71] and for general $i \in \mathbb{N}$ by Bernshteyn [11]. Note that while the high-level idea described before the definition provides some intuition on how to build multi-step Vizing chains, it is the d -proper shiftable requirement on the chains that makes the analysis complicated. A successful approach to remedy this issue is to count the number of *non-overlapping* multi-step Vizing chains, these might be defined for example by demanding that $P_j \cap P_{j'} = \emptyset$ for every $1 \leq j < j' \leq i$. A precise analysis of the number of i -step Vizing chains for various values of i was done in [11, 14, 36, 71]. In particular, we refer the reader to [14, Section 2 and 5.1] for the description and analysis of the randomized Multi-Step Vizing Algorithm that finds a short augmenting subgraph with high-probability.

In the remainder of this section, we discuss the connections to local expansion to complement the discussion in Section 7.1.




FIGURE 5. An example of a 3-step Vizing chain. Fans are depicted by black thick edges, while dotted curly edges indicate how the alternating paths continue after they get truncated.

Consider the following problem. Let G be a finite graph and suppose that for every $e \in U_d$ there is no augmenting subgraph of size at most $\ell \gg \Delta$. In particular, every (1-step) Vizing chain $W(x, e) = F_1 \cap P_1$ assigned to $e \in U_d$ and $x \in e$ has size at least ℓ . What can we say about the fraction of uncolored edges $|U_d|/m$?

The *key property* of multi-step Vizing chains, which is essential in all the arguments in [11, 14, 36, 71], is that every edge $f \in \text{dom}(d)$ belongs to at most $\text{poly}(\Delta)^i$ -many distinct i -step Vizing chains. Indeed, detecting an uncolored edge $e \in U_d$ such that $f \in V(x, e)$ for some i -step Vizing chain basically boils down to following iteratively at most i -many alternating paths starting at f . Importantly, this bound does not depend on ℓ . On the other hand, we expect, at least when i is a small constant, that the number of i -step Vizing chains that we can grow from a given $e \in U_d$ should be of the order $(\ell/i)^i$. The heuristic here is that we run the “branching chain” process i -times and in each step $1 \leq j \leq i$ we can truncate the currently last alternating path, that must have length at least ℓ/i , at almost every of the first ℓ/i -many edges to create P_j .

To derive an estimate on $|U_d|/m$, we formalize these observations in a double counting argument assuming for simplicity that $i = 1$. Define an auxiliary bipartite graph $H_d = U_d \sqcup \text{dom}(d)$, where (e, f) forms an edge if there is $x \in e$ such that $f \in W(x, e)$. By our assumption we have that $\deg_{H_d}(e) \geq \ell - 1$ for every $e \in U_d$, and by the discussion in the previous paragraph, we have that $\deg_{H_d}(f) \leq \text{poly}(\Delta)$ for every $f \in \text{dom}(d)$. A double counting argument in H_d then yields that

$$(7.3) \quad (\ell - 1)|U_d| \leq \text{poly}(\Delta) |\text{dom}(d)| \leq \text{poly}(\Delta)m.$$

In particular, the fraction of uncolored edges is $O(1/\ell)$ if we view Δ as a constant.

We remark that an analogous computation for the perfect matching problem, assuming that the size of every augmenting path is large, does not give any estimate on the size of unmatched edges as the degrees in both parts of the analogously defined auxiliary bipartite graph might be comparable.

7.3. Applications. We start by recalling the remarkable result of Christiansen Theorem 2.5.

Theorem 7.8 (Christiansen [36]). *Let G be a graph, $d: E \rightarrow [\Delta(G)+1]$ be a partial edge coloring and $e \in U_d$ be an uncolored edge. Then there exists an augmenting subgraph H for e such that $|E(H)| \in O(\Delta^7 \log(n))$.*

The strategy to derive Theorem 7.8 is to fix $i, \ell \in \mathbb{N}$ and count the number of vertices reachable by i -step Vizing chains with the property that each of its truncated alternating paths have size at most ℓ . This strategy originated in [11], where Bernshteyn described a randomized algorithm for growing multi-step Vizing chains and showed that *there are many* uncolored edges that admit an augmenting subgraph of size $O(\log^2(n))$. We remark that for his analysis, he needed both parameters to satisfy $i, \ell \in \Theta(\log(n))$.

The main novelty in the approach of Christiansen is a highly non-trivial argument that even with $\ell \in \text{poly}(\Delta)$ the number of reachable points from every uncolored edge grows exponentially in i unless one of these chains is d -augmenting. This

implies, in particular, that after $\Theta(\log(n))$ -steps we either exhaust the whole graph or find a d -augmenting chain.

In order to derive a fast distributed algorithm for edge colorings based on augmenting subgraphs of size $\Omega(\log(n))$, one needs to show not only that such subgraphs exist, but also that they can be assigned to a large portion of the uncolored edges in an edge disjoint way. Results along this line are essential in all the constructions from [11, 14, 36].

We formulate the currently best version of such result that is from [14] and that combines the idea of multi-step Vizing chains with the *entropy compression method*. We remark that the existence of the randomized distributed algorithm from Theorem 3.7 is a direct consequence of this result, see [14] for more details.

Theorem 7.9 (Many disjoint augmenting subgraphs, Theorem 2.4 in [14]). *Let $d: E \rightarrow [\Delta + 1]$ be a partial edge coloring. There exists a randomized distributed algorithm of LOCAL complexity $\text{poly}(\Delta) \log(n)$ that outputs a set $W \subseteq U_d$ with expected size $\mathbb{E}(|W|) \geq |U_d| / \text{poly}(\Delta)$ and an assignment of augmenting subgraphs H_e for $e \in W$ such that*

- (1) *the graphs H_e , for $e \in W$, are vertex disjoint,*
- (2) *$|E(H_e)| \leq \text{poly}(\Delta) \log(n)$.*

Next, we turn our attention to applications in measurable combinatorics. Combining the ideas and concepts introduced in the previous parts of this section allows us to sketch a formal proof of the measurable Vizing's theorem for graphings.

Theorem 7.10 (Vizing's theorem for graphings, [71]). *Let $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ be a graphing of degree bounded by $\Delta \in \mathbb{N}$. Then $\chi'(\mathcal{G}) \leq \Delta + 1$.*

Proof sketch. As a first step we pass to a Borel probability measure ν on E that satisfies the mass transport principle on the Borel line graph $L(\mathcal{G})$ and has the property that $\mu(\{v : \exists e \in A \text{ } v \in e\}) = 0$ whenever $\nu(A) = 0$ for every measurable set $A \subseteq E$, see [71, Section 3].

Following the same line of reasoning as in Theorem 7.2, see [71, Section 4], we construct a sequence of partial measurable edge colorings $(d_n)_{n \in \mathbb{N}}$ with the property that

- (a) every d_n -augmenting 2-step Vizing chain has length at least $\ell_n = 2^n$,
- (b) $U_{d_n} \supseteq U_{d_{n+1}}$,
- (c) $\nu(\{e \in E : d_{n+1}(e) \neq d_n(e)\}) \leq 8\ell_n \nu(U_{d_n})$

for every $n \in \mathbb{N}$.

Define the auxiliary bipartite graphing $\mathcal{H}_{d_n} = U_{d_n} \sqcup \text{dom}(d_n)$ by setting (e, f) to be an edge if there is a 2-step Vizing chain $V(x, e)$ of length at most ℓ_n such that $f \in V(x, e)$. The computation in [71, Section 3] confirms the intuition that $\deg_{\mathcal{H}_{d_n}}(e) \in \Omega(\ell_n^2)$ for every $e \in U_{d_n}$ while the bound $\deg_{\mathcal{H}_{d_n}}(f) \in \text{poly}(\Delta)$ holds for every $f \in \text{dom}(d_n)$ by the argument from Section 7.2. As ν satisfies the mass transport principle, we conclude that

$$\nu(U_{d_n}) \in O(1/\ell_n^2).$$

Combined with the property (c) above, we see that the condition analogous to (7.2) is satisfied as

$$\sum_{n=0}^{\infty} 8 \ell_n \mu(U_{d_n}) = \sum_{n=0}^{\infty} 2^{n+3} / 2^{2n} < \infty.$$

In particular, $d = \lim_{n \rightarrow \infty} d_n$ is a well-defined measurable edge coloring of \mathcal{G} by the Borel–Cantelli lemma. \square

We finish this section with a discussion about the additional ingredients that are needed in the proof of the full measurable Vizing’s theorem for general measured graphs, Theorem 4.11.

It is a standard fact in measurable combinatorics, see for example [87, Chapter 9], that if $\mathcal{G} = (V, E, \mathcal{B}, \mu)$ is a measured graph, then we may assume that μ satisfies a *weighted version of the mass transport principle*. That is, every pair of points $(v, w) \in V \times V$ is assigned a value $\rho_\mu(v, w)$ that measures the ratio of the infinitesimal mass at w compared to v and satisfies the weighted version of the equality (4.1). The function ρ_μ is known as the *Radon–Nikodym cocycle* of \mathcal{G} (and μ). Mimicking the argument from Theorem 7.10 and replacing the length of the 2-step Vizing chains by their relative weight given by ρ_μ , we run into a trouble when estimating the size of $\nu(U_{d_n})$ for every $n \in \mathbb{N}$. This is because short chains might have big weight which works against us in the double counting argument on \mathcal{H}_{d_n} .

The way how to remedy this issue in [69] is to “smooth out” the measure μ so that the value of the Radon–Nikodym cocycle on every pair of adjacent vertices $\{v, w\} \in E$ satisfies

$$(7.4) \quad 1/(2\Delta) \leq \rho(v, w) \leq 2\Delta.$$

This is roughly done as follows. First, we use Theorem 4.6 to assign in a Borel way to μ -almost every $x \in V$ a finite measure κ_x supported on the (countable) connectivity component of x with the property that if $\text{dist}_{\mathcal{G}}(x, y) = k$, then $\kappa_x(y) = \rho_\mu(x, y)/(4\Delta)^k$. Then we define a new measure ν to be the normalization of the convolution of the collection $\{\kappa_x\}_{x \in V}$ and μ . That is, $\nu(A)$ is equal, up to re-normalization, to

$$\mu \star \{\kappa_x\}_{x \in V}(A) = \int_{x \in V} \sum_{y \in A} \kappa_x(y) d\mu(x)$$

for every $A \in \mathcal{B}$. A direct computation shows that ν and μ are equivalent and ρ_ν satisfies (7.4).

The estimate in (7.4) allows to compare the size and weight of any, for technical reasons, 3-step Vizing chain in a way that works in our favor for the double counting argument. We refer the reader to [69] for more details about the arguments.

8. Brooks’ theorem for subexponential growth graphs

In this section we describe a deterministic distributed algorithm of LOCAL complexity $O(\log^*(n))$ that produces Brooks coloring on a class of finite graphs of a fixed subexponential growth that satisfy the assumptions from Brooks’ theorem.

This result was likely known to experts in the field but we were unable to find a reference, or a written proof.

By Bernshteyn's correspondence [12], Theorem 5.6, this implies that Brooks' theorem holds in the Borel context for Borel graphs of subexponential growth. It was communicated to us by Bernshteyn that this result also follows from the proofs in [42] as the graph induced on the set of the vertices where the described construction fails has to be of exponential growth. This is in stark contrast to Theorem 6.14; see also Remark 4.13, which discusses a derandomization perspective for subexponential growth graphs.

Recall that $f: \mathbb{N} \rightarrow \mathbb{N}$ is *subexponential* if $f \in o((1 + \varepsilon)^n)$ for every $\varepsilon > 0$.

Definition 8.1 (Graphs of subexponential growth). A function $f: \mathbb{N} \rightarrow \mathbb{N}$ *bounds the growth* of a graph G if for every $v \in V(G)$ we have $|B_G(x, r)| < f(r)$. We denote by \mathbb{F}_f the class of graphs of growth bounded by f .

We say that class of graphs \mathbf{F} , or a graph G , is of a *subexponential growth*, if there is subexponential $f: \mathbb{N} \rightarrow \mathbb{N}$ such that $\mathbf{F} \subseteq \mathbb{F}_f$, or $G \in \mathbb{F}_f$.

We prove the following.

Theorem 8.2. *Let $\Delta \geq 3$ and $\mathbf{F} \subseteq \mathbb{F}_f$ be a class of graphs of a subexponential growth and degree bounded by Δ that do not contain the complete graph on $\Delta + 1$ vertices. Then there is a deterministic distributed algorithm for Brooks coloring on \mathbf{F} of LOCAL complexity $O(\log^*(n))$.*

As an immediate corollary, we get a Borel version of Brooks theorem.

Remark 8.3. Note that the below theorem also follows from [12, Theorem 2.15] combined with the distributed upper bound for Brooks' coloring from [61]: indeed, by Bernshteyn's theorem, $O(\log n)$ randomized complexity algorithms directly yield Borel solutions on subexponential growth graphs.

Theorem 8.4 (Borel Brooks coloring for graphs of subexponential growth, [12, 42]). *Let $\mathcal{G} = (V, E, \mathcal{B})$ be a Borel graph of subexponential growth that does not contain a complete graph on $\Delta(\mathcal{G}) + 1$ vertices and satisfies $\Delta(\mathcal{G}) \geq 3$. Then $\chi_{\mathcal{B}}(\mathcal{G}) \leq \Delta(\mathcal{G})$.*

Proof. Let f be the subexponential function that bounds the growth of \mathcal{G} . Define \mathbf{F} to be the class of finite subgraphs of \mathcal{G} and observe that we may apply Theorem 8.2 for \mathbf{F} . The rest follows from Theorem 5.6. \square

Our approach to prove Theorem 8.2 is identical to the one in [16], where Bernshteyn and Dhawan proved an analogous theorem for Vizing coloring. Nevertheless, we provide the details for completeness. The only difference is that in the place where they use the result of Christiansen, Theorem 2.5, to bound the size of an augmenting subgraph for edge colorings, we use the result of Panconesi and Srinivasan [105] that we recall next.

Theorem 8.5 (Small augmenting subgraphs for vertex colorings, [105]). *Let G be a finite graph that does not contain a complete graph on $\Delta + 1$ vertices, where $\Delta \geq \Delta(G)$, $v \in V(G)$ and c be a partial vertex coloring with Δ colors such that*

$U_c = \{v\}$. Then there is an augmenting subgraph for v that is a path of length $O(\log n)$.

Proof of Theorem 8.2. Let f be the subexponential function that bounds the growth of the graphs from \mathbf{F} . It is clearly enough to prove the claim for the inclusion maximal class \mathbf{F} .

The main observation, see [16, Lemma 3.3], that crucially uses the subexponential bound on the growth is that we can always find an augmenting subgraph of constant size.

Proposition 8.6. *There exists an $R \in \mathbb{N}$ so that for every $G \in \mathbf{F}$, every partial vertex coloring c of G with Δ colors and every vertex $v \in U_c$ there is an augmenting subgraph H for v contained in $B_G(v, R)$.*

Proof. Let C be a constant so that Theorem 8.5 holds with a path of size $C \log n$ for each graph of size n . Find some $\varepsilon > 0$ with $C \log(1 + \varepsilon) < 1$ and an R large enough so that $f(R) < (1 + \varepsilon)^R$.

Now let c be an arbitrary partial vertex coloring of an element $G \in \mathbf{F}$ with Δ colors and $v \in U_c$. Consider the graph G' which is the connected component of v in the graph G restricted to $B_G(v, R) \cap (\text{dom}(c) \cup \{v\})$. Let H be an augmenting subgraph in G' for c returned by Theorem 8.5. Then we have

$$|V(H)| < C \log |V(G')| \leq C \log |B_G(v, R)| \leq C \log f(R) < C \log ((1 + \varepsilon)^R) < R.$$

Recall that an augmenting subgraph is connected by the definition, so, if the size of H is $\leq R - 1$, it has to be contained in the ball $B_G(v, R - 1)$. But, then H must be an augmenting subgraph for c in the entire graph G . \square

Now, to prove Theorem 8.2, take any graph $G \in \mathbf{F}$. Recall that $G^{(k)}$ stands for the k th power of G , that is, the graph where two vertices are connected if their graph distance in G is at most k . Clearly, the graph $G^{(2R+2)}$ has degrees bounded by Δ^{2R+2} , so, in $O(\log^* n)$ rounds we can compute a proper vertex coloring of $G^{(2R+2)}$ with $k = \Delta^{2R+2} + 1$ colors by Theorem 3.4. Enumerate the color classes as $(C_i)_{i \leq k}$. Then, in k -many steps inductively, define partial vertex colorings $(c_i)_{i \leq k}$ as follows: if c_{j-1} is defined on $\bigcup_{i < j} C_i$, for every $v \in C_j$ by Proposition 8.6 we can find an augmenting subgraph in $B_G(v, R)$ for c_{j-1} and v . Recolor the vertices using this graph to obtain c_j . Observe that, as the distance between any two elements of C_j is $> 2R + 2$, c_j will still be a partial vertex coloring. This algorithm terminates after k -many steps, yielding a vertex coloring with Δ colors. Finally, observe that as k and R are constants, the overall complexity of the algorithm remains $O(\log^* n)$, as desired. \square

9. Open problems

We finish the survey by listing some exciting open problems that are directly connected to the results discussed in this survey.

Edge colorings. We start with a couple of questions about Vizing colorings. The main problem here concerns the existence of a fast distributed algorithm for Vizing's theorem, see also [14, Question 1.9] and [11, Question 1.2].

Problem 9.1. *Is the randomized LOCAL complexity of Vizing coloring $o(\log(n))$?*

In [36], Christiansen proved a *local version of Vizing's theorem*, that is, every finite graph G admits an edge coloring c with $\Delta(G) + 1$ colors with the property that

$$c(e) \leq \max\{\deg_G(v), \deg_G(w)\} + 1$$

for every edge $e = \{v, w\} \in E(G)$. The approach from [36] does not seem to automatically generalize to the measurable setting.

Problem 9.2. *Is there a measurable version of local Vizing's theorem?*

In Remark 4.12, we briefly mentioned the topic of *Baire measurable colorings*. Qian and Weilacher [108] showed that $\Delta + 2$ colors are enough for a Baire measurable edge coloring of a graph of degree bounded by Δ .

Problem 9.3. *Is there a Baire measurable version of Vizing's theorem?*

In Section 7.1, we have seen that there is a factor of iid perfect matching on the Δ -regular tree T_Δ . The following question is still open. Note that Kun's result [89] implies that König's line coloring theorem fails for measurably bipartite acyclic Δ -regular measured graphs in general, for every $\Delta \geq 3$.

Problem 9.4. *Is there a factor of iid edge coloring with Δ colors on the Δ -regular tree T_Δ ? In other words, does measurable König's line coloring theorem hold for the iid graph on T_Δ ?*

An intriguing question related to the results in Section 7.1 asks about an optimal algorithm for perfect matching on the Δ -regular tree T_Δ . This can be formalized using the notion of *finitary factors of iid*, which provides a strictly finer framework to measure complexity of LCLs than the one given by the LOCAL model, see [72, 74, 80] for the exact definition and related results. On a high-level, finitary factors of iid are randomized deterministic algorithms that do not know the size of the graph, hence, the locality into which every vertex needs to look is a random variable that is almost surely finite. The complexity is measured by the tail decay of this random variable.

A recent result of Berlow, Bernshteyn, Lyons and Weilacher gives the first example of an LCL problem (on \mathbb{Z}^2) that admits a factor of iid solution but does not admit a finitary factor of iid solution. It would be nice to show that the perfect matching problem on T_Δ has the same property. We remark that by [23], there is no finitary factor of iid perfect matching on T_Δ of subexponential tail decay.

Problem 9.5. *Is there a finitary factor of iid perfect matching on the Δ -regular tree T_Δ ?*

LCL problems. The class of measured graphs that satisfy the mass transport principle is much more studied than the class of general measured graphs. It is,

however, not known if the mass transport principle gives any advantage for solving LCL problems. To the following problem the discussion at the end of Section 7.3 about Radon–Nikodym cocycles might be relevant.

Problem 9.6 (Question 6.10 in [8]). *Is there an LCL problem that could be solved on measured graphs that satisfy the mass transport principles but not on general measured graphs? The same question is open on any subclass of measured graphs that look locally like a fixed graph, e.g., grid, regular tree etc.*

A possible candidate to answer Problem 9.6 is the *unfriendly coloring problem*. In its most basic form it asks for a coloring of vertices of a graph G with 2 colors so that every vertex receives a color that is different from the color of at least half of its neighbors. Conley and Tamuz [44] showed that a measured graph \mathcal{G} admits a measurable unfriendly coloring provided that the Radon–Nikodym cocycle satisfies $1 - 1/\Delta \leq \rho(x, y) \leq 1 + 1/\Delta$ for every edge $\{x, y\} \in \mathcal{G}$. Note that this is much stronger than (7.4).

Problem 9.7. *Does every measured graph of bounded degree admits a measurable unfriendly coloring? In fact, it is also open whether every Borel graph of bounded degree admits a Borel unfriendly coloring.*

In Remark 4.13, we discussed versions of derandomization in the context of measurable combinatorics. While the recent counterexamples discussed there answered in negative most of the basic questions, the following two problems are still open.

Problem 9.8. *Let \mathbf{F} be a class of graphs of bounded degree and linear growth. Is it true that an LCL problem on \mathbf{F} can be solved measurably on measured graphs that are locally in \mathbf{F} if and only if it can be solved in a Borel way?*

The following problem is a special case of the celebrated conjecture of Chang and Pettie [35].

Problem 9.9. *Let \mathbf{F} be a class of graphs of bounded degree and of subexponential growth. Is the LOCAL complexity of the distributed Lovász Local Lemma on \mathbf{F} equal to $\Theta(\log^*(n))$? That is, is there a speed-up result from randomized LOCAL complexity $O(\log(n))$ to deterministic LOCAL complexity $O(\log^*(n))$ for \mathbf{F} ?*

Complexity. In Section 5.1, we mentioned in Theorem 5.4 that Brooks' theorem strongly fails in the Borel context. In particular, the set of Δ -regular acyclic Borel graphs that admit a Borel vertex coloring with Δ colors is as complicated as it gets. Related is the question about the complexity of edge colorings. Thornton [116] showed that distinguishing between Borel chromatic index Δ and $\Delta + 1$ for Borel graphs of degree bounded by Δ is Σ_2^1 -complete. As the graphs in his construction are not acyclic and Marks' result shows that Borel chromatic index of Δ -regular acyclic graphs can be anything between Δ and $2\Delta - 1$, it is natural to ask the following.

Problem 9.10. *Is it Σ_2^1 -complete to distinguish between any k and $k + 1$ Borel chromatic index on Δ -regular acyclic graphs, where $\Delta \leq k \leq 2\Delta - 2$?*

As mentioned above, the projective complexity in the Borel context reflects to some extent the finite experience: deciding Borel 2-colorability of a Borel graph is strictly simpler than to decide its 3-colorability (see [30, 116]). More precisely, the codes of the collections of Borel graphs that admit a Borel 2-coloring form a Π_1^1 set, while for k -coloring, where $k \geq 3$, the corresponding set is Σ_2^1 -complete, see also Section 5.1.

Problem 9.11. *Let $\Delta \geq 3$. Do Δ -regular acyclic Borel hyperfinite Borel graphs that have Borel chromatic number at most Δ form a Σ_2^1 -complete?*

To our knowledge, surprisingly, there are no results about complexity in the measurable context.

Problem 9.12. *Let $k \geq 3$ and fix a Borel measure μ on some standard Borel space X . What is the complexity of the codes¹⁰ of locally countable Borel graphs \mathcal{G} with $\chi_\mu(\mathcal{G}) \leq k$?*

Is there a precise way in which deciding the measurable 2-colorability of an acyclic bounded degree graphing is easier than to decide its 3-colorability?

The straightforward upper bound for both 2- and 3-colorings in this situation is Σ_1^1 , so it would be extremely interesting if the latter was Σ_1^1 -hard and the former was Δ_1^1 .

A different complexity related question arises from the fact that all the examples considered in complexity considerations in [25, 118] follow the same pattern: constructing a family of Borel graphs parametrized by the reals, so that each of them admits a Borel 3-coloring, but altogether it is hard to decide whether such a coloring exists. This suggests that there could be a positive result, once one excludes this type of examples. The following question formalizes the most optimistic scenario.

Problem 9.13. *Assume that \mathcal{G} is a Δ -regular acyclic Borel graph so that there is no smooth Borel superequivalence relation¹¹ E of $E_{\mathcal{G}}$ so that \mathcal{G} restricted to each E class admits a Borel Δ -coloring, for $\Delta \geq 3$. Does Marks' example, i.e., the Schreier graph of the free part of the left-shift action of Γ_Δ on 2^{Γ_Δ} , Borel homomorph to \mathcal{G} ?*

Furthermore, the produced examples use non-compactness of the underlying space in an essential way. Thus, another positive result could be possible when one restricts to nicer Borel graphs, such as continuous graphs on compact spaces.

Problem 9.14. *Assume that \mathcal{G} is a compact subshift of the free part of 2^{Γ_Δ} that does not admit a Borel Δ -coloring, for $\Delta \geq 3$. Does Marks' example Borel homomorph to \mathcal{G} ? What is the projective complexity of such subshifts?*

Generalizations. One of the most exciting problems is to transfer the full strength of Martin's theorem to the finite context. Indeed, excluding deterministic algorithms of complexity $o(\log n)$ for Δ vertex colorings on trees of degree $\leq \Delta$ only

¹⁰ In some fixed encoding, where all the codes form a Borel set. This is possible if one discards a μ measure zero set.

¹¹ See [85] for the definition.

requires the determinacy of clopen games, the proof of which is straightforward, see [84, Section 20.B]. In contrast, full Borel Determinacy is one of the most involved and exciting arguments of classical descriptive set theory.

A key feature that distinguishes the Borel case from the LOCAL model is that “one can see in an arbitrary distance”. It would be nice to have a version of the LOCAL model, which incorporates this, but still non-trivial. Let us pose the following -admittedly vague- problem.

Problem 9.15. *Is there a stronger variant of the LOCAL model (e.g., for which all the LCL problems that are known to admit a Borel solution are efficiently solvable, see [23, 72]) so that 3-coloring of acyclic graphs of degree at most 3 is still not possible, because of Marks’ method?*

More generally, instead of colorings, it is reasonable to consider homomorphism problems. The following problems are variants of [86, Problem 8.12].

Problem 9.16. *Characterize the following families of finite graphs:*

- (1) *Graphs H so that there is a deterministic distributed algorithm of complexity $O(\log^* n)$ on the class of trees of degree $\leq \Delta$ that produces a homomorphism to H .*
- (2) *Graphs H so that every Δ -regular acyclic Borel graph admits a Borel homomorphism to H .*
- (3) *Graphs H so that there is a factor of iid homomorphism to H from T_Δ .*
- (4) *Graphs H so that every Δ -regular acyclic Borel graph equipped with a measure μ admits a Borel homomorphism to H modulo a μ null set.*

Is membership in any of these families algorithmically undecidable?

Let us mention that undecidability phenomena are known to be present in the case of grid graphs, see [59, Section 4.3]. However, it is also conceivable that one can have a full understanding of the solvability of such homomorphism problems in the special case of trees, as it is the case in the Baire measurable context, see [23].

It seems that even to prove a negative result in these cases, one would need to significantly extend the collection of examples H to which there are definable homomorphisms of the above sort, see [48], as, so far all the examples about which we are aware of are very close to $K_{\Delta+1}$.

Acknowledgments. We are very grateful to Jukka Suomela for letting us use Figure 1. We are also grateful to Anton Bernshteyn, Yi-Jun Chang, Clinton Conley, Alexander Kechris, Andrew Marks, Oleg Pikhurko, Václav Rozhoň, Felix Weilacher, and the anonymous referee for their comments on the first version of this manuscript.

J. G. was supported by MSCA Postdoctoral Fellowships 2022 HORIZON-MSCA-2022-PF-01-01 project BORCA grant agreement number 101105722. ZV was supported by Hungarian Academy of Sciences Momentum Grant no. 2022-58 and National Research, Development and Innovation Office (NKFIH) grants no. 113047, 129211.

References

- [1] D. Aldous, R. Lyons, *Processes on unimodular random networks*, Electron. J. Probab. **12** (2007), 1454–1508.
- [2] A. Balliu, S. Brandt, D. Olivetti, J. Suomela, *Almost global problems in the local model*, Distrib. Comput. **34**(4)(2021), 259–281.
- [3] A. Balliu, J. Hirvonen, J. H. Korhonen, T. Lempiäinen, D. Olivetti, J. Suomela, *New classes of distributed time complexity*, In: I. Diakonikolas (ed.) et al., *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing, STOC '18*, Los Angeles, CA, USA, June 25–29, 2018. Association for Computing Machinery (ACM), New York, NY, 2018, 1307–1318.
- [4] É. Bamas, L. Esperet, *Distributed coloring of graphs with an optimal number of colors*, In: R. Niedermeier (ed.) et al., *36th international symposium on theoretical aspects of computer science, STACS 2019*, March 13–16, 2019, Berlin, Germany, Schloss Dagstuhl – Leibniz Zentrum für Informatik, Wadern, LIPIcs – Leibniz Int. Proc. Inform. **126**, 2019, 10.
- [5] L. Barenboim, M. Elkin, *Sublogarithmic distributed mis algorithm for sparse graphs using nash-williams decomposition*, In: *Proceedings of the 27th annual ACM symposium on principles of distributed computing, PODC '08*, Toronto, Canada, August 18–21, 2008, Association for Computing Machinery (ACM), New York, NY, 2008, 25–34.
- [6] L. Barenboim, M. Elkin, *Distributed Graph Coloring: Fundamentals and Recent Developments*, Morgan & Claypool Publishers, 2013.
- [7] F. Bencs, A. Hrušková, L. Tóth, *Factor-of-IID schreier decorations of lattices in euclidean spaces*, Discrete Math. **347**(9) (2024), 114056.
- [8] K. Berlow, A. Bernshteyn, C. Lyons, F. Weilacher, *Separating complexity classes of lcl problems on grids*, arXiv:2501.17445 (2025).
- [9] A. Bernshteyn, *Measurable versions of the Lovász local lemma and measurable graph colorings*, Adv. Math. **353** (2019), 153–223.
- [10] ———, *Descriptive combinatorics and distributed algorithms*, Notices Am. Math. Soc. **69**(9) (2022), 1496–1507.
- [11] ———, *A fast distributed algorithm for $(\Delta + 1)$ -edge-coloring*, J. Comb. Theory, Ser. B **152** (2022), 319–352.
- [12] ———, *Distributed algorithms, the Lovász Local Lemma, and descriptive combinatorics*, Invent. Math. **233** (2023), 495–542.
- [13] ———, *Probabilistic constructions in continuous combinatorics and a bridge to distributed algorithms*, Adv. Math. **415** (2023), 108895.
- [14] A. Bernshteyn, A. Dhawan, *Fast algorithms for Vizing’s theorem on bounded degree graphs*, arXiv:2303.05408 (2023).
- [15] ———, *A linear-time algorithm for $(1 + \epsilon)\Delta$ -edge-coloring*, arXiv:2407.04887 (2024).
- [16] ———, *Borel Vizing’s theorem for graphs of subexponential growth.*, Proc. Am. Math. Soc. **153** (2025), 7–14.
- [17] A. Bernshteyn, F. Weilacher, *Borel versions of the Local Lemma and LOCAL algorithms for graphs of finite asymptotic separation index*, arXiv:2308.14941 (2023).
- [18] S. Bhattacharya, D. Carmon, M. Costa, S. Solomon, T. Zhang, *Faster $(\Delta + 1)$ -edge coloring: Breaking the $m\sqrt{n}$ time barrier*, arXiv:1907.03201 (2024).
- [19] B. Bollobás, *The independence ratio of regular graphs*, Proc. Am. Math. Soc. **83**(2) (1981), 433–436.
- [20] M. Bowen, F. Weilacher, *Definable König theorems*, Proc. Am. Math. Soc. **151** (2023), 4991–4996.
- [21] Z. Brady, *Notes on CSPs and polymorphisms*, arXiv:2210.07383 (2022).
- [22] S. Brandt, *An automatic speedup theorem for distributed problems*, In: T. Nowak (ed.), *Proceedings of the 38th ACM symposium on principles of distributed computing, PODC '19*, Toronto, ON, Canada, July 29 – August 2, 2019, Association for Computing Machinery (ACM), New York, NY, 2019, 379–388.

- [23] S. Brandt, Y. Chang, J. Grebík, C. Grunau, V. Rozhoň, Z. Vidnyánszky, *Local problems on trees from the perspectives of distributed algorithms, finitary factors, and descriptive combinatorics*, arXiv:2106.02066 (2021), A short version of the paper was presented at the 13th Innovations in Theoretical Computer Science conference (ITCS 2022).
- [24] S. Brandt, Y. Chang, J. Grebík, C. Grunau, V. Rozhoň, Z. Vidnyánszky, *Deterministic distributed algorithms and descriptive combinatorics on Δ -regular trees*, arXiv:2204.09329 (2022).
- [25] S. Brandt, Y. Chang, J. Grebík, C. Grunau, V. Rozhoň, and Z. Vidnyánszky, *On homomorphism graphs*, Forum Math. Pi **12** (2024), e10.
- [26] S. Brandt, O. Fischer, J. Hirvonen, B. Keller, T. Lempinen, J. Rybicki, J. Suomela, J. Uitto, *A lower bound for the distributed Lovász local lemma*, In: D. Wichs (ed.) et al., *Proceedings of the 48th annual ACM SIGACT symposium on theory of computing, STOC '16*, Cambridge, MA, USA, June 19–21, 2016, Association for Computing Machinery (ACM), New York, NY, 2016, 479–488.
- [27] S. Brandt, J. Hirvonen, J. H. Korhonen, T. Lempinen, P. R. J. Östergård, C. Purcell, J. Rybicki, J. Suomela, P. Uznański, *LCL problems on grids*, In: *Proceedings of the 2017 ACM symposium on principles of distributed computing, PODC '17*, Washington, DC, USA, July 25–27, 2017, Association for Computing Machinery (ACM), New York, NY, 2017, 101–110.
- [28] R. L. Brooks, *On colouring the nodes of a network.*, Math. Proc. Camb. Philos. Soc. **37** (1941), 194–197.
- [29] A. A. Bulatov, *A dichotomy theorem for nonuniform CSPs*, In: *IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, 2017, 319–330.
- [30] R. Carroy, B. D. Miller, D. Schritterser, Z. Vidnyánszky, *Minimal definable graphs of definable chromatic number at least three*, Forum Math. Sigma **9** (2021), e7.
- [31] N. Chandgotia, S. Unger, *Borel factors and embeddings of systems in subshifts*, arxiv:2203.09359 (2022).
- [32] Y. Chang, *The Complexity Landscape of Distributed Locally Checkable Problems on Trees*, In: H. Attiya (ed.), *34th international symposium on distributed computing, DISC 2020*, virtual conference, October 12–16, 2020, Proceedings, Schloss Dagstuhl – Leibniz Zentrum für Informatik, Wadern, LIPIcs – Leibniz Int. Proc. Inform. **179**, 2020, 18.
- [33] Y. Chang, Q. He, W. Li, S. Pettie, J. Uitto, *Distributed edge coloring and a special case of the constructive Lovász local lemma*, ACM Trans. Algorithms **16**(1) (2019), 8.
- [34] Y. Chang, T. Kopelowitz, S. Pettie, *An exponential separation between randomized and deterministic complexity in the local model*, SIAM J. Comput. **48**(1) (2019), 122–143.
- [35] Y. Chang, S. Pettie, *A time hierarchy theorem for the local model*, SIAM J. Comput. **48**(1) (2019), 33–69.
- [36] A. B. G. Christiansen, *The power of multi-step vizing chains*, In: B. Saha (ed.) et al., *Proceedings of the 55th annual ACM SIGACT symposium on theory of computing, STOC '23*, Orlando, FL, USA, June 20–23, 2023, Association for Computing Machinery (ACM), New York, NY, 2023, 1013–1026.
- [37] K. Chung, S. Pettie, H. Su, *Distributed algorithms for the Lovász local lemma and graph coloring*, Distrib. Comput. **30**(4) (2017), 261–280.
- [38] R. Cole, U. Vishkin, *Deterministic coin tossing with applications to optimal parallel list ranking*, Inf. Control **70**(1) (1986), 32–53.
- [39] C. T. Conley, S. Jackson, A. S. Marks, B. Seward, R. Tucker-Drob, *Hyperfiniteness and Borel combinatorics*, J. Eur. Math. Soc. (JEMS) **22**(3) (2020), 877–892.
- [40] ———, *Borel asymptotic dimension and hyperfinite equivalence relations*, Duke Math. J. **172**(16) (2023), 3175–3226.
- [41] C. T. Conley, S. C. Jackson, D. Kerr, A. S. Marks, B. Seward, R. D. Tucker-Drob, *Følner tilings for actions of amenable groups*, Math. Ann. **371** (2018), 663–683.
- [42] C. T. Conley, A. S. Marks, R. D. Tucker-Drob, *Brooks' theorem for measurable colorings*, Forum Math. Sigma **4** (2016), e16.

- [43] C. T. Conley, B. D. Miller, *Measure reducibility of countable Borel equivalence relations*, Ann. Math. (2) **185**(2) (2017), 347–402.
- [44] C. T. Conley, O. Tamuz, *Unfriendly colorings of graphs with finite average degree*, Proc. Lond. Math. Soc. (3) **122**(2) (2021), 229–233.
- [45] E. Csóka, L. Grabowski, M. András, O. Pikhurko, K. Tyros, *Moser-tardos algorithm with small number of random bits*, arXiv:2203.05888 (2022).
- [46] E. Csóka, L. Grabowski, A. Máthé, O. Pikhurko, K. Tyros, *Borel version of the Local Lemma*, arXiv:1605.04877 (2017).
- [47] E. Csóka, G. Lippner, O. Pikhurko, *Invariant Gaussian processes and independent sets on regular graphs of large girth*, Random Struct. Algorithms **47**(2) (2015), 284–303.
- [48] E. Csóka, Z. Vidnyánszky, *FIID homomorphisms and entropy inequalities*, arXiv:2407.10006 (2024).
- [49] R. Dougherty, M. Foreman, *Banach-Tarski paradox using pieces with the property of Baire*, Proc. Natl. Acad. Sci. USA **89**(22) (1992), 10726–10728.
- [50] G. Elek, *Qualitative graph limit theory. Cantor dynamical systems and constant-time distributed algorithms*, arXiv:1812.07511 (2018).
- [51] G. Elek, G. Lippner, *Borel oracles. An analytical approach to constant-time algorithms*, Proc. Am. Math. Soc. **138**(8) (2010), 2939–2947.
- [52] T. Emden-Weinert, S. Hougardy, B. Kreuter, *Uniquely colourable graphs and the hardness of colouring graphs of large girth*, Comb. Probab. Comput. **7**(4) (1998), 375–386.
- [53] P. Erdős, *Graph theory and probability*, Can. J. Math. **11** (1959), 34–38.
- [54] M. Fischer, M. M. Halldórsson, Y. Maus, *Fast distributed Brooks’ theorem*, In: N. Bansal (ed.) et al., *Proceedings of the 34th annual ACM-SIAM symposium on discrete algorithms, SODA 2023*, Florence, Italy, January 22–25, 2023, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA; Association for Computing Machinery (ACM), New York, NY, 2023, 2567–2588.
- [55] J. Frisch, F. Shinko, Z. Vidnyánszky, *Hyper-hyperfiniteness and complexity*, arXiv:2409.16445 (2024).
- [56] D. Gaboriau, *Coût des relations d’équivalence et des groupes*, Invent. Math. **139**(1) (2000), 41–98.
- [57] H. N. Gabow, T. Nishizeki, O. Kariv, D. Leven, O. Terada, *Algorithms for edge-coloring graphs*, Technical Report TR-41/85, Dept. of Computer Science, Tel Aviv University, 1985.
- [58] S. Gao, S. Jackson, E. Krohne, B. Seward, *Borel combinatorics of countable group actions*, arXiv:2401.13866 (2024).
- [59] S. Gao, S. Jackson, E. Krohne, B. Seward, *Continuous combinatorics of abelian group actions*, arXiv:1803.03872 (2018).
- [60] M. Ghaffari, *Distributed maximal independent set using small messages*, In: T. M. Chan (ed.), *Proceedings of the 30th annual ACM-SIAM symposium on discrete algorithms, SODA 2019*, San Diego, CA, USA, January 6–9, 2019, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA; New York, NY: Association for Computing Machinery (ACM), 2019, 805–820.
- [61] M. Ghaffari, C. Grunau, V. Rozhoň, *Improved deterministic network decomposition*, In: *Proc. of the 32. ACM-SIAM Symp. on Discrete Algorithms (SODA) (USA)*, Society for Industrial and Applied Mathematics, 2021, 2904–2923.
- [62] M. Ghaffari, D. G. Harris, F. Kuhn, *On derandomizing local distributed algorithms*, In: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, 2018, 662–673.
- [63] M. Ghaffari, J. Hirvonen, F. Kuhn, Y. Maus, *Improved distributed delta-coloring*, In: *Proceedings of the 37th ACM symposium on principles of distributed computing, PODC ’18*, Egham, UK, July 23–27, 2018, Association for Computing Machinery (ACM), New York, NY, 2018, 427–436.
- [64] M. Ghaffari, F. Kuhn, Y. Maus, *On the complexity of local distributed graph problems*, In: H. Hatami (ed.) et al., *Proceedings of the 49th annual ACM SIGACT symposium on*

- theory of computing, STOC '17*, Montreal, QC, Canada, June 19–23, 2017, Association for Computing Machinery (ACM), New York, NY, 2017, 784–797.
- [65] A. V. Goldberg, S. A. Plotkin, G. E. Shannon, *Parallel symmetry-breaking in sparse graphs*, SIAM J. Discrete Math. **1**(4) (1988), 434–446.
 - [66] Ł. Grabowski, A. Máthé, O. Pikhurko, *Measurable circle squaring*, Ann. Math. (2) **185** (2017), 671–710.
 - [67] Ł. Grabowski, A. Máthé, O. Pikhurko, *Measurable equidecompositions for group actions with an expansion property*, J. Eur. Math. Soc. **24** (2022), 4277–4326.
 - [68] J. Grebík, *Approximate Schreier decorations and approximate König’s line coloring theorem*, Ann. Henri Lebesgue **5** (2022), 303–315.
 - [69] ———, *Measurable Vizing’s theorem*, Forum Math. Sigma **13** (2025), e32.
 - [70] J. Grebík, C. Higgins, *Complexity of finite borel asymptotic dimension*, arXiv:2411.08797 (2024).
 - [71] J. Grebík, O. Pikhurko, *Measurable versions of vizing’s theorem*, Adv. Math. **374** (2020), 107378.
 - [72] J. Grebík, V. Rozhoň, *Classification of local problems on paths from the perspective of descriptive combinatorics*, In: J. Nešetřil, G. Perarnau, J. Rué, O. Serra (eds.), *Extended Abstracts EuroComb 2021: European Conference on Combinatorics, Graph Theory and Applications*, Trends in Mathematics **14**, Birkhäuser, Cham, 2021, 553–559.
 - [73] J. Grebík, Z. Vidnyánszky, *Complexity of linear equations and infinite gadgets*, arXiv:2501.06114 (2025).
 - [74] J. Grebík, V. Rozhoň, *Local problems on grids from the perspective of distributed algorithms, finitary factors, and descriptive combinatorics*, Adv. Math. **431** (2023), 109241.
 - [75] M. Gromov, *Asymptotic invariants of infinite groups, geometric group theory, vol. 2*, Lond. Math. Soc. Lecture Note Ser. **182**, Cambridge University Press, 1993.
 - [76] C. Grunau, V. Rozhoň, S. Brandt, *The landscape of distributed complexities on trees and beyond*, In: A. Milani (ed.) et al., *Proceedings of the 41st ACM symposium on principles of distributed computing, PODC '22*, Salerno, Italy, July 25–29, 2022, Association for Computing Machinery (ACM), New York, NY, 2022, 37–47.
 - [77] R. P. Gupta, *Studies in the theory of graphs*, Ph.D. Thesis, Tata Institute of Fundamental Research, 1967.
 - [78] H. Hatami, L. Lovász, B. Szegedy, *Limits of locally-globally convergent graph sequences*, Geom. Funct. Anal. **24**(1) (2014), 269–296.
 - [79] J. Hirvonen, J. Suomela, *Distributed algorithms 2020*, Aalto University, Finland, 2021.
 - [80] A. E. Holroyd, O. Schramm, D. B. Wilson, *Finitary coloring*, Ann. Probab. **45**(5) (2017), 2867–2898.
 - [81] I. Holyer, *The NP-completeness of edge-coloring*, SIAM J. Comput. **10**(4) (1981), 718–720.
 - [82] S. Jackson, A. S. Kechris, A. Louveau, *Countable Borel equivalence relations*, J. Math. Log. **2**(1) (2002), 1–80.
 - [83] R. M. Karp, *On the computational complexity of combinatorial problems*, Networks **5**(1) (1975), 45–68.
 - [84] A. S. Kechris, *Classical descriptive set theory*, Grad. Texts Math. **156**, Springer-Verlag, New York, 1995.
 - [85] A. S. Kechris, *The theory of countable borel equivalence relations*, Cambridge University Press, 2024.
 - [86] A. S. Kechris, A. S. Marks, *Descriptive graph combinatorics*, <http://www.math.caltech.edu/~kechris/papers/combinatorics20book.pdf>, 2020.
 - [87] A. S. Kechris, B. D. Miller, *Topics in orbit equivalence*, Lect. Notes Math. **1852**, Springer-Verlag, Berlin, 2004.
 - [88] A. S. Kechris, S. Solecki, S. Todorćević, *Borel chromatic numbers*, Adv. Math. **141**(1) (1999), 1–44.
 - [89] G. Kun, *The measurable Hall theorem fails for treeings*, arXiv:2106.02013 (2021).

- [90] M. Laczkovich, *Equidecomposability and discrepancy; a solution of Tarski's circle-squaring problem*, J. Reine Angew. Math. **404** (1990), 77–117.
- [91] N. Linial, *Locality in distributed graph algorithms*, SIAM J. Comput. **21**(1) (1992), 193–201.
- [92] C. R. Lyons, *Descriptive combinatorics on trees, grids, and non-amenable graphs*, Ph.D. Thesis, UCLA, 2024.
- [93] R. Lyons, *Factors of IID on trees*, Comb. Probab. Comput. **26**(2) (2017), 285–300.
- [94] R. Lyons, F. Nazarov, *Perfect matchings as IID factors on non-amenable groups*, Eur. J. Comb. **32**(7) (2011), 1115–1125.
- [95] A. Marks, S. Unger, *Baire measurable paradoxical decompositions via matchings*, Adv. Math. **289** (2016), 397–410.
- [96] A. S. Marks, *A determinacy approach to Borel combinatorics*, J. Am. Math. Soc. **29**(2) (2016), 579–600.
- [97] A. S. Marks, *Measurable graph combinatorics*, In: D. Beliaev (ed.) et al., *International congress of mathematicians 2022, ICM 2022*, Helsinki, Finland, July 6–14, 2022, European Mathematical Society (EMS), 2022, 1488–1502.
- [98] A. S. Marks, S. T. Unger, *Borel circle squaring*, Ann. Math. (2) **186**(2) (2017), 581–605.
- [99] D. A. Martin, *Borel determinacy*, Ann. Math. (2) **102**(2) (1975), 363–371.
- [100] C. McDiarmid, *Concentration for independent permutations*, Comb. Probab. Comput. **11**(2) (2002), 163–178.
- [101] B. D. McKay, N. C. Wormald, B. Wysocka, *Short cycles in random regular graphs*, Electron. J. Comb. **11**(1) (2004), 66.
- [102] M. Molloy, B. Reed, *Colouring graphs when the number of colours is almost the maximum degree*, J. Comb. Theory, Ser. B **109** (2014), 134–195.
- [103] M. Naor, L. Stockmeyer, *What can be computed locally?*, SIAM J. Comput. **24**(6) (1995), 1259–1277.
- [104] H. N. Nguyen, K. Onak, *Constant-time approximation algorithms via local improvements*, In: *49th Annual IEEE Symposium on Foundations of Computer Science*, 2008, 327–336.
- [105] A. Panconesi, A. Srinivasan, *The local nature of Δ -coloring and its algorithmic applications*, Combinatorica **15** (1995), 255–280.
- [106] A. Panconesi, A. Srinivasan, *Improved distributed algorithms for coloring and network decomposition problems*, In: *Proc. 24th ACM Symp. on Theory of Computing (STOC)*, 1992, 581–592.
- [107] O. Pikhurko, *Borel combinatorics of locally finite graphs*, In: K. K. Dabrowski (ed.) et al., *Surveys in combinatorics 2021. Based on plenary lectures given at the 28th British combinatorial conference*, Durham, UK, July 5–9, 2021. Cambridge University Press, Cambridge, Lond. Math. Soc. Lect. Note Ser. **470**, 2021, 267–320.
- [108] L. Qian, F. Weilacher, *Descriptive combinatorics, computable combinatorics, and asi algorithms*, arXiv:2206.08426 (2022).
- [109] V. Rozhoň, *Invitation to local algorithms*, arXiv:2406.19430 (2024).
- [110] V. Rozhoň, M. Ghaffari, *Polylogarithmic-time deterministic network decomposition and distributed derandomization*, In: K. Makarychev (ed.) et al., *Proceedings of the 52nd annual ACM SIGACT symposium on theory of computing, STOC '20*, Chicago, IL, USA, June 22–26, 2020, Association for Computing Machinery (ACM), New York, NY, 2020, 350–363.
- [111] B. Seward, *Personal communication*.
- [112] C. Sinnamón, *Fast and simple edge-coloring algorithms.*, arXiv:1907.03201 (2021).
- [113] S. Skulrattanakulchai, *Δ -list vertex coloring in linear time*, Inf. Process. Lett. **98**(3) (2006), 101–106.
- [114] M. Stiebitz, D. Scheide, B. Toft, L. M. Favrholdt, *Graph edge coloring*, Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2012.
- [115] H. Su, H. T. Vu, *Towards the locality of Vizing's theorem*, In: M. Charikar (ed.) et al., *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing, STOC '19*, Phoenix, AZ, USA, June 23–26, 2019, Association for Computing Machinery (ACM), New York, NY, 2019, 355–364.

- [116] R. Thornton, *An algebraic approach to Borel CSPs*, arXiv:2203.16712 (2022).
- [117] ———, *Orienting Borel graphs*, Proc. Am. Math. Soc. **150**(4) (2022), 1779–1793.
- [118] S. Todorčević, Z. Vidnyánszky, *A complexity problem for Borel graphs*, Invent. Math. **226** (2021), 225–249.
- [119] L. M. Tóth, *Invariant Schreier decorations of unimodular random networks*, Ann. Henri Lebesgue **4** (2021), 1705–1726.
- [120] V. G. Vizing, *Critical graphs with given chromatic class*, Metody Diskret. Analiz. **5** (1965), 9–17.
- [121] F. Weilacher, *Borel edge colorings for finite dimensional groups*, arXiv:2104.14646 (2021).
- [122] R. Wilson, *Four Colors Suffice: How the Map Problem Was Solved - revised color edition*, Princeton university press, 2014.
- [123] N. C. Wormald, *Models of random regular graphs*, Lond. Math. Soc. Lect. Note Ser. **267** (1999), 239–298.
- [124] D. Zhuk, *A proof of the CSP dichotomy conjecture*, J. ACM **67**(5) (2020), 1–78.
- [125] D. Zuckerman, *Linear degree extractors and the inapproximability of max clique and chromatic number*, In: J. M. Kleinberg (ed.), *Proceedings of the 38th annual ACM symposium on theory of computing, STOC 2006* Seattle, WA, USA, May 21–23, 2006. ACM Press, New York, NY, 2006, 681–690.