

Invitation to Local Algorithms

Václav Rozhoň

December 2023

Introduction

The field of local algorithms is an area on the border of theoretical computer science and discrete mathematics where a lot of progress has happened in the past decade. This text is trying to serve as an introductory material presenting a certain view of the field. It aims to be helpful to beginning researchers in the area or researchers working in adjacent areas.

There are already many resources on various aspects of local algorithms: the classical book of Peleg [Pe100], survey of Suomela [Suo13], book of Barenboim and Elkin [BE13], introductory text of Suomela [Suo20], or a recent book by Hirvonen and Suomela [HS20]. Unlike other texts, this one primarily explores the field's conceptual framework and complexity-theoretical aspects, rather than delving into individual problems.

If you find errors in the text, please let me know – this is the first version of it so there will be many! Several researchers generously gave me feedback on a preliminary version of this text, including Yi-Jun Chang, Mohsen Ghafari, Yannic Maus, Seth Pettie, and Jukka Suomela.

Contents

1	Local Complexity Fundamentals	3
1.1	First Example	3
1.2	Formal Definitions	4
1.3	Sequential vs Distributed Local Complexity	7
1.4	Derandomization	9
1.5	Network Decompositions	11
1.6	Bounds for Concrete Problems	15
2	The Sublogarithmic Regime	17
2.1	The Symmetry-Breaking Regime	17
2.1.1	Fast Coloring Algorithm and its Implications	17
2.1.2	Lower bound for coloring via round elimination	20
2.2	The Lovász Local Lemma Regime	22
2.2.1	Fast Algorithms for the Local Lemma	23
2.2.2	Lower Bound via Round Elimination	26
2.3	Speedups and Slowdowns	28
2.3.1	Slowdowns	29
2.3.2	Speedups	30
2.4	Classification of Local Problems	33
2.4.1	Sequential local complexities	34
2.5	Classification of Local Problems for Concrete Graph Classes	36
3	Applications	39
3.1	Distributed Computing (CONGEST)	39
3.2	Local Computation Algorithms and the Volume Model	39
3.3	PRAM	40
3.4	Massively Parallel Computing (MPC)	41
3.5	Descriptive Combinatorics	41
3.6	Other Models	42

Chapter 1

Local Complexity Fundamentals

In the first chapter, we introduce local algorithms and local problems in [Section 1.1](#). We then carefully discuss the appropriate formal definitions in [Section 1.2](#). The following sections [Sections 1.3 to 1.5](#) discuss the basic theory of local algorithms and aim to convey that we are after a very clean, fundamental, and robust concept. Finally, [Section 1.6](#) surveys some known results for concrete local problems.

1.1 First Example

Consider a very long oriented cycle that we want to properly color with as few colors as possible (see [Figure 1.1](#)). Two colors are enough if the number of vertices n is even, otherwise we need three colors. However, there is something uneasy about the 2-coloring solution even when it is possible – the solution lacks any flexibility. A decision to color any particular vertex with one of the two colors already implies how all the other vertices are going to be colored.

This lack of flexibility can be undesirable for all kinds of reasons, typically when we want to design a coloring algorithm that is in some way parallel or distributed. If we enlarge our palette to three colors, the problem seems to go away though: Now, coloring one vertex red still implies that its neighbors are not red, but other vertices can have an arbitrary color.

Imagine that there is a computer in every vertex of the cycle and neighboring computers can communicate. The computers are trying to solve the coloring problem together. How many rounds of communication are needed until each computer outputs its color? A message-passing algorithm of this sort is known as a *local algorithm* and we define it formally in [Section 1.2](#).

It is possible to convince oneself that in the case of 2-coloring, at least around $n/4$ rounds are necessary to solve the problem, even if n is divisible by two.¹ But what about the 3-coloring scenario? Can we solve that problem in 10 rounds of communication? Or $O(\log n)$? Or is it similarly hard to 2-coloring?

There indeed is a simple randomized local algorithm that solves our 3-coloring problem after $O(\log n)$ rounds of exchanging messages. Let's describe it next. The algorithm should serve as an example that nontrivial local algorithms are indeed possible, though we will see a better one for this problem later.

The 3-coloring algorithm has two phases. In the first phase, every computer flips a coin and selects itself with probability $1/2$ (top picture in [Figure 1.1](#)). Subsequently, the vertex asks its neighbors whether they are also selected. If at least one neighbor is selected, the vertex unselects itself (middle picture in [Figure 1.1](#)).

In the second phase of the algorithm, every selected vertex first colors itself red. Then, it is responsible for coloring the yet uncolored vertices until the next red vertex. The red node sends a message in the direction of edges, asking the subsequent vertices to color themselves by alternating the two remaining colors (see the bottom picture in [Figure 1.1](#)). This algorithm properly colors the oriented cycle with 3 colors and the number of communication rounds that it needs is asymptotically at most as large as the length of the longest run of non-red vertices in our coloring.

¹Consider two opposing nodes u, v in the cycle graph: In less than $n/4$ rounds of communication, there is no third node that could send a message to both u and v . Intuitively, the two vertices then cannot know whether their distance is even or odd. This argument can be made into a proper lower bound after the model of local algorithms is properly defined in [Section 1.2](#).

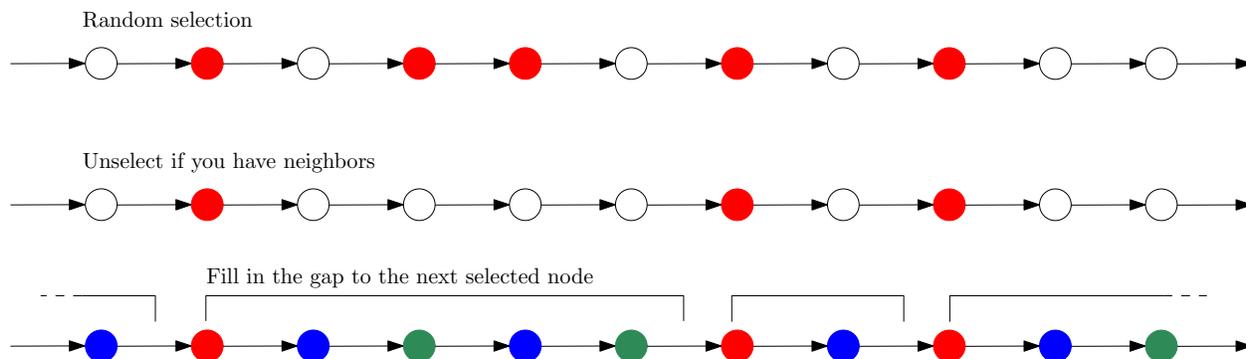


Figure 1.1: An example local algorithm that uses three colors to color a long cycle, a small part of which is shown. First, every vertex flips a coin and selects itself with probability $1/2$. Second, a vertex unselects itself whenever a neighbor is selected. Third, selected vertices color themselves red and each selected vertex is then responsible for coloring the subsequent vertices until the next selected one with alternating colors.

To understand this quantity, consider any run of ℓ consecutive nodes and let us upper bound the probability that the run does not contain any red node. This is done by splitting the run into consecutive triples of vertices. For every triple, we know that with probability $1/8$, its middle node is initially selected, while its two neighbors are not. The selected middle node then remains selected after the end of the first phase and is colored red. Making this argument for every triple and using that the appropriate events are independent, we conclude that the probability of no red node in the run is at most $(7/8)^{\lceil \ell/3 \rceil}$. Taking $\ell = O(\log n)$ and union bounding over all n different runs of vertices of length ℓ , we conclude that all of them contain at least one red node with $1 - 1/\text{poly}(n)$ probability. We will call this guarantee “with high probability” later on. That is, our algorithm finishes after $O(\log n)$ communication rounds, with high probability.

Surprisingly, the fastest local algorithm for the 3-coloring problem has a much better, albeit not constant complexity of $O(\log^* n)$.² However, instead of focusing on specific algorithms, this text is trying to give a bit more general understanding of what is going on here. For example, it turns out that if you construct any local algorithm with complexity $O(\log n)$ for any reasonable problem defined on the cycle as we just did, there is a general theory that can turn this algorithm into a faster, $O(\log^* n)$ -round, algorithm for the very same problem ([Theorem 2.30](#)). Clearly, something interesting is going on here!

1.2 Formal Definitions

In this section, we formally define local problems and algorithms.

Local problems: We will be mostly interested in the so-called local problems. These are the problems on graphs such that if the solution is incorrect, we can find out by looking at a small neighborhood of one vertex.

Given a graph G and its node $u \in V(G)$, the *ball* $B_G(u, r)$ ³ around u of radius r is the subgraph of nodes around u up to distance r . More generally, an *r -hop neighborhood* is a graph with one highlighted node v such that the radius of that graph measured from v is at most r .

Definition 1.1 (A local problem). *Local problem*⁴ Π with checkability radius r is formally a triplet (S, r, \mathcal{P}) . Here, S is a finite set of allowed labels and each \mathcal{P} is a set of “allowed” S -colored r -hop neighborhoods. A solution to Π in a graph G is an assignment of a color from S to every vertex of G such that for every $u \in V(G)$ we have $B_G(u, r) \in \mathcal{P}$.

²The function $\log^* n$ measures how many times we need to take the logarithm of n until we get a value of size at most 2, i.e., $\log^* 2^2 = 1$, $\log^* 2^{2^2} = 2$ and so on.

³We write $B(u, r)$ when G is clear from the context.

⁴Our definition is a simplified variant of the definition of the so-called locally checkable labeling problem by Naor and Stockmeyer [[NS95](#)], discussed later in [Section 2.5](#).

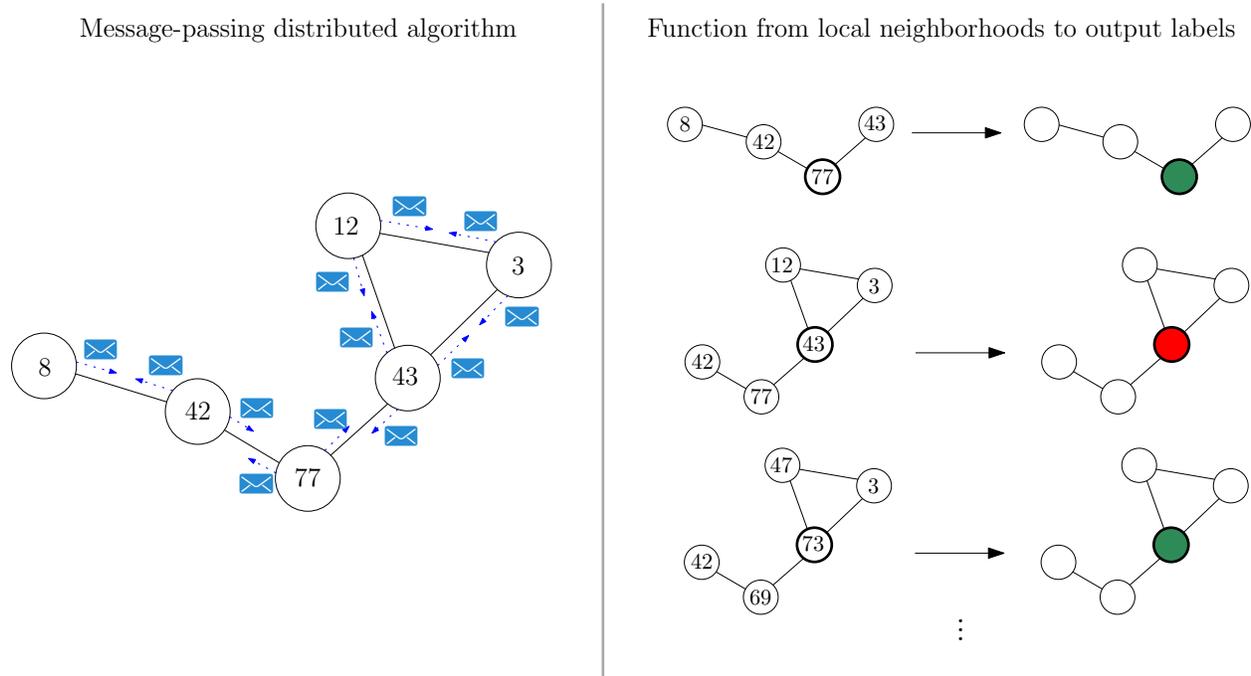


Figure 1.2: This picture shows the two fundamentally different ways of understanding local algorithms. Left: A $t(n)$ -round local algorithm is a distributed protocol where in each round, each node can send any message to any of its neighbors. The computers start with the knowledge of their unique identifier (or a random string). Right: A local algorithm with round complexity $t(n)$ is a function that maps all possible $t(n)$ -hop neighborhoods into output labels. Applying this function to every vertex of the input graph always has to solve our problem. For example, if our problem is a coloring problem, the first two local neighborhoods in the above table need to map the two vertices with labels 77 and 43 to different colors, since the two labeled 2-hop neighborhoods could be a part of the same graph (which is, in fact, shown on the left).

For example, 3-coloring is a local problem for $S = \{R, G, B\}$, $r = 1$, and \mathcal{P} containing all properly colored 1-hop neighborhoods. On the other hand, a non-example of a local problem is coloring an input graph on n vertices with n colors: the local problem should not have different constraints for graphs of different size. Of course, while the theory of local algorithms is simplest for local problems as we defined them, the applications of local algorithms are not limited to local problems.

Local algorithms: There are two equivalent ways of thinking about local algorithms⁵ and both of them are important (see Figure 1.2). An intuitive, algorithmic definition was already sketched in Section 1.1: We assume that there is a computing device at every node. For simplicity, these devices are assumed to have unbounded computational power, thus excluding Turing machines from the definitions. A $t(n)$ -round local algorithm is a protocol where these devices communicate for $t(n)$ synchronous message-passing rounds using the edges of the input graph to send messages. At the beginning, the device in each node starts only with the information about its identifier/random string and the size of the graph, n . When the protocol finishes, each device outputs its part of the solution (e.g., its color).

It will be also helpful to understand an alternative, equivalent definition that extracts the essence of what we are measuring with local algorithms. In this alternative definition, a local algorithm with round complexity $t(n)$ is simply a function that we can apply to every ball $B(u, t(n))$ of the input graph to compute the output at a given node u . Let us state it now formally.

Definition 1.2 (Local algorithm). *A local algorithm \mathcal{A} with a round complexity of $t(n)$ is a function that accepts two inputs: firstly, the value n , and secondly, a labeled $t(n)$ -hop neighborhood of a certain node u .*

⁵In the literature, these algorithms are often referred to as “distributed algorithms in the LOCAL model of computing”. We use the shorter and less formal term “local algorithm” for better readability.

When we use this second definition, *running* a local algorithm on an input graph G simply means coloring each node $u \in V(G)$ with the output of $\mathcal{A}_n(B_G(u, t(n)))$ where we used the first parameter of a local algorithm, the size of the underlying graph n , as a subscript. *Solving* a problem Π on G simply means that after running \mathcal{A} on G , the output colors satisfy constraints \mathcal{P} on all vertices of G .

Moreover, in the case of *deterministic* local algorithms, we assume that the nodes of the input graph are additionally labeled with unique identifiers from the range $[n^{O(1)}] = \{1, 2, \dots, n^{O(1)}\}$.⁶ In the case of *randomized* local algorithms, we assume that the nodes of the input graph are labeled with infinite bit strings. Solving a problem then means solving it with overall error probability at most $1/n^{O(1)}$, if the bit strings are sampled independently randomly.⁷

We notice that if there is a deterministic local algorithm solving some problem with round complexity $t(n)$, there is also a randomized local algorithm solving the same problem with the same round complexity. This is because any randomized algorithm can start by each node generating a random identifier from the range $[n^C]$: The probability that these identifiers are not unique, i.e., some two nodes have the same identifier, is at most $n^2 \cdot \frac{1}{n^C}$. Choosing C large enough, this error probability can be made as small as any polynomial function of n .

Finally, we remark that we can talk about local algorithms solving problems on graphs with additional structure (e.g. directed graphs) or on concrete graph classes. For example, in our introductory example from [Section 1.1](#), it makes sense to think of all definitions relative not to the class of all graphs but to the class of graphs that are oriented paths. One interesting setup that we discuss mostly in [Chapter 2](#) is the class of bounded-degree graphs where we fix some constant Δ and analyze the class of graphs of degree at most Δ . Notice that on these graphs, the set \mathcal{P} from the definition of local problems, as well as the support of the function \mathcal{A} from the definition of local algorithms, are finite.

Equivalence of the two definitions: Let's see a proof sketch of why the two definitions are equivalent. On the one hand, let's say we are given a function \mathcal{A} that maps $t(n)$ -hop neighborhoods to output labels and we want to construct a $t(n)$ -round message-passing protocol. Consider the protocol where in the i -th round, each vertex u sends its neighbors everything there is to know about the ball $B(u, i)$: How the graph looks like and what are the identifiers/random strings at every node. Each node v can then internally use this information from its neighbors to learn everything there is to know about the ball $B(u, i + 1)$. After $t(n)$ rounds of communication, each vertex v knows its $t(n)$ -hop neighborhood $B(v, t(n))$. The vertices then stop communicating, and each one applies the function \mathcal{A} locally to its ball.

On the other hand, assume that we have a $t(n)$ -round communication protocol and want to turn it into a function \mathcal{A} that takes $t(n)$ -hop neighborhood as inputs. We notice that if we know the $t(n)$ -hop neighborhood $B(u, t(n))$ of a node u , we can simulate the first round of the protocol in that ball and get to know the state of all vertices in $B(u, t(n) - 1)$ after the first round. Continuing like this inductively, we conclude that starting with the knowledge of $B(u, t(n))$, we can learn the state of the protocol at the center node u after $t(n)$ rounds.

There are two different but equivalent ways of understanding local algorithms.

1. *They are message-passing protocols running for some number of rounds.*
2. *The output at each node is a function of its local neighborhood.*

Importantly, it will be very helpful for us to keep *both* definitions in mind: When we design local algorithms, the message-passing definition is more helpful as it is natural to think as "first, we run the protocol \mathcal{A}_1 , then the protocol \mathcal{A}_2 ". On the other hand, when we prove lower bounds, the formal definition of [Definition 1.2](#) is easier to use. When we think of applications to distributed/parallel algorithms, the protocol-design definition is preferable since this is how the actual parallel/distributed algorithms are implemented. In some other applications, like applications to descriptive combinatorics, the formal definition is perhaps a bit more natural.

⁶While assuming polynomial-range identifiers may look a bit arbitrary, we will see in [Chapter 2](#) that the notion of deterministic algorithms is very robust. We simply need a way of breaking the potential symmetry of the input graph.

⁷Formally-minded readers may feel uneasy about the definitions not specifying the constant in the $n^{O(1)}$ expressions. We will see later in [Theorem 2.20](#) that the exact constant in the definition typically does not matter. Formally, when we say that there is a local algorithm, it means that for every C there is an algorithm in the setup where the size of identifiers/error probability is n^C .

1.3 Sequential vs Distributed Local Complexity

This section presents one of the most fundamental results for local algorithms. Currently, it may be very unclear what kind of problems can be solved with a local algorithm of round complexity, say, $\text{poly log } n$. This will become much clearer, since we will next see that, up to $\text{poly log } n$, the model of local algorithms is the same as the model of so-called *sequential local algorithms* that are much easier to understand.

The case of maximal independent set: As an example, let's think of a concrete local problem known as the *maximal independent set* problem. In this problem, every node must be labeled either selected or unselected. The constraint is that each selected node should not neighbor any other selected node, while each unselected node should neighbor at least one selected node⁸.

Is there a local algorithm constructing a maximal independent set in a polylogarithmic number of rounds? This is not clear at all! The answer to this question is positive, and perhaps the simplest algorithm is the randomized algorithm of Luby [Lub86; ABI86]. This algorithm in fact served as the foundational example that later led Linial [Lin92] to define local algorithms. But Luby's algorithm is a non-trivial algorithm⁹ and just by staring at the maximal independent set problem, it is quite unclear whether a fast local algorithm exists, or not.

This stands in stark contrast with the "sequential" world: If we do not care about all vertices outputting the answer "at once", we can compute a maximal independent set with the following simple algorithm: We choose any order of vertices and iterate over them in that order. Whenever we consider a vertex u , we look at its neighbors, and if at least one of them is already selected, we mark u as unselected. Otherwise, we mark u as selected.

Here is a curious property of the above algorithm: We can still think of it as a "local" algorithm. Indeed, each node makes its decision by examining its 1-hop neighborhood. The only difference is that the algorithm is a *sequential local algorithm* where we iterate over nodes in an arbitrary order, not a *distributed local algorithm*¹⁰ as we defined it in Definition 1.2 where all nodes have to output the answer at once.

A fundamental result of local complexity is the fact that these two definitions are equally powerful, up to polylogarithmic factors. Hence, in the concrete example of the maximal independent set, we can think of this problem as being "easy" not because of a clever algorithm like Luby's, but because of the above simple sequential algorithm.

The distributed round complexity of any local problem equals its sequential local complexity, up to $\text{poly log}(n)$.

Formal definition of sequential local algorithms: We next make this principle formal. We need to start by defining a general sequential local algorithm. Here is a definition of a deterministic sequential local algorithm, made slightly more powerful than the maximal-independent-set algorithm by allowing the output of each node to be not just the final color, but also additional information that future vertices can read off the node.

Definition 1.3 (Sequential local algorithms). *A sequential local algorithm of local complexity $t(n)$ is a function \mathcal{A} defined on labeled $t(n)$ -hop neighborhoods. Its output for a neighborhood $B(u, t(n))$ around a node u is a pair (s, t) , with s being the output at u and t being additional information stored at u . An input neighborhood to \mathcal{A} has some nodes labeled by these pairs.*

Running a sequential local algorithm means iterating over the vertices in some order and each time applying \mathcal{A} to produce the answer at that vertex. When we run \mathcal{A} on a node v , the algorithm has access to all already produced pairs (s, t) at the vertices in $B(u, t(n))$ on which \mathcal{A} has already been run. Solving a problem with a sequential local algorithm means that regardless of the order in which we choose the vertices, this process results in a solution to the problem.

⁸This is a much easier problem than the *maximum independent set* problem where we additionally maximize the number of selected nodes.

⁹We can briefly describe the algorithm: It runs in $O(\log n)$ rounds and in each round, every vertex chooses a random number. If its number is the largest among its neighbors, the vertex goes in the independent set and is removed from future iterations, together with its neighbors. After $O(\log n)$ rounds, all vertices are removed with high probability and the algorithm terminates.

¹⁰We sometimes use the name *distributed local algorithm* to stress that we are talking about a local algorithm and not a sequential local one.

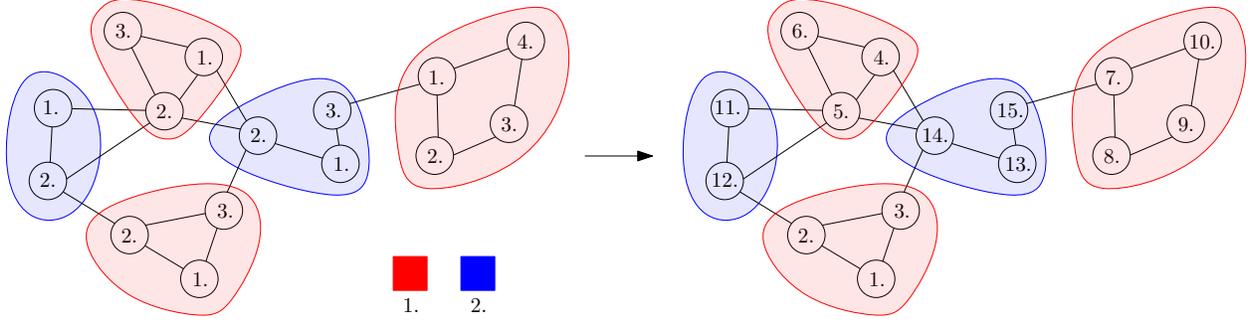


Figure 1.3: This picture shows a network decomposition with $c = 2$ color classes and $d = 2$ diameter. It also shows how network decomposition is used to convert an input sequential local algorithm (of local complexity 1) into a distributed local algorithm in [Theorem 1.4](#).

Left: The color classes of the network decompositions are ordered as (red, blue). We iterate over the color classes and in one iteration, we consider each cluster separately and simulate an input sequential local algorithm in it (see the node ordering inside each cluster). When the algorithm is simulated in blue clusters, it has access to the output of neighboring red vertices.

Right: The partial simulations of the sequential local algorithm in each cluster are consistent with a single run of that algorithm over all vertices.

Notice that we do not require unique identifiers in the definition; we will see later in [Theorem 2.7](#) that they are not needed for local problems. We can also define (oblivious) randomized algorithms where first an adversary chooses an order in which we iterate over vertices; then we sample random bits in each vertex and run our sequential local algorithm. We will next prove the following theorem by Ghaheri, Kuhn, and Maus [\[GKM17\]](#).¹¹

Theorem 1.4 (Ghaheri, Kuhn, and Maus [\[GKM17\]](#)). *Let \mathcal{A} be a deterministic (randomized) sequential local algorithm with local complexity $t(n)$. Then, there is a deterministic (or randomized, respectively) distributed local algorithm simulating \mathcal{A} with round complexity $\tilde{O}(t(n) \cdot \log^3(n))$.*¹²

We note that it is known that there are local problems such that their sequential and distributed local complexity differ by a factor of $\Omega(\log n / \log \log n)$. [\[Gav+09\]](#)

Network decompositions: A crucial tool that we will rely on in this section and the next one is the concept of a *network decomposition*. Network decomposition is a clustering of the input graph into clusters of small diameter¹³ (see [Figure 1.3](#)).

Definition 1.5 (Network decomposition). *A (c, d) -network decomposition of a graph G is a coloring of G with c colors. We require that vertices of each color induce a graph such that each of its connected components (that we call clusters) has diameter at most d .*

We defer the discussion about the existence of network decompositions to [Section 1.5](#). For now, we will simply state the guarantees of the currently best deterministic network decomposition construction.

Theorem 1.6 ([\[Gha+23\]](#)). *There is a deterministic local algorithm that outputs a $(O(\log n), O(\log n \cdot \log \log \log n))$ -network decomposition in $\tilde{O}(\log^3 n)$ rounds.*

Proof of Theorem 1.4: Armed with the above algorithm for network decomposition, let us prove [Theorem 1.4](#).

Proof of Theorem 1.4. An important concept employed throughout this text is working within the power graph: Given a graph G and a parameter r , we define the power graph G^r to be the graph with $V(G^r) = V(G)$ where two vertices are connected if their distance in G^r is at most r .

¹¹Some version of this theorem was implicitly understood in the late 80s and led to the development of algorithms for network decompositions [\[Awe+89; LS93\]](#).

¹²We use $\tilde{O}(t(n))$ to denote the complexity $O(t(n) \cdot \log^{O(1)} t(n))$.

¹³A diameter of a graph G is defined as $\max_{u,v} d_G(u, v)$ where $d_G(u, v)$ is the distance between u and v in G .

We start with a sequential local algorithm \mathcal{A} of complexity $t(n)$. We will work in the power graph $G^{t(n)}$ and construct a network decomposition in it with $c, d' = \tilde{O}(\log n)$ via [Theorem 1.6](#). Consider this network decomposition in the context of the original graph G : We constructed clusters of actual diameter $d \leq t(n) \cdot d' = \tilde{O}(t(n) \cdot \log n)$ in G . Moreover, two clusters from the same color class have distance at least $t(n) + 1$ in G . Finally, since every communication round in G^r can be simulated in r communication rounds in G , the round complexity of constructing our network decomposition is $\tilde{O}(t(n) \cdot \log^3 n)$ by [Theorem 1.6](#).

We will now use our clustering to simulate \mathcal{A} . We will simulate an order of iterating over the vertices where we first iterate over all the vertices in the first color class, then all the vertices in the second color class, and so on. For a fixed color class, we will arbitrarily simulate the algorithm \mathcal{A} in each cluster C independently of all other clusters of the same color (see [Figure 1.3](#)).

Notice that every two clusters C_1, C_2 of the same color i are far enough so that the $t(n)$ -hop neighborhood of any vertex $u \in C_1$ never contains a vertex $u' \in C_2$. Hence, simulations in different clusters of the i -th color do not interact. Our simulation is thus a faithful simulation of iterating over all the vertices of G in a certain order and applying the sequential algorithm \mathcal{A} to them.

Finally, let us discuss how the simulation of \mathcal{A} is implemented with a local algorithm. Think of a message-passing protocol with c phases where in the i -th phase, each cluster C of color i first chooses a leader node, e.g., the node with the smallest identifier. This node collects all information about the output of \mathcal{A} so far up to the distance $t(n)$ from C . Then, the leader node internally simulates \mathcal{A} on C and sends the result of that simulation back to all nodes in C . All this can be done in a number of rounds proportional to the diameter of C and $t(n)$. Thus, the overall round complexity of the simulation is $c \cdot \tilde{O}(t(n) \cdot \log n) = \tilde{O}(t(n) \cdot \log^2 n)$. \square

1.4 Derandomization

By now, we understand that the sequential local complexity is closely related to the distributed round complexity. However, we still do not understand the power of randomness. There might be scenarios where a problem's randomized (sequential or distributed) local complexity is significantly lower than its deterministic (sequential or distributed) complexity. Interestingly, this never happens for local problems. Distributed local algorithms for them can be derandomized with $\text{poly} \log(n)$ slowdown in round complexity.¹⁴

Any local problem has the same deterministic and randomized round complexity, up to $\text{poly} \log(n)$.

Before proving this result, let us contemplate how it fits into the big picture. Thus far, we have seen *six* plausible definitions of how to measure the local complexity of a problem. There are the following three different ways of thinking about it, and for each one of them, we can define both the deterministic and the randomized complexity:

1. (*distributed protocol*) There are computers at nodes, we design a message-passing protocol, and we measure the number of rounds of this protocol.
2. (*distributed local complexity*) Output at each node is a function of its local neighborhood.
3. (*sequential local complexity*) We iterate over the nodes in an arbitrary order and settle each output at a node by looking at its local neighborhood.

We now understand that all of these definitions are equivalent, up to $\text{poly} \log(n)$ and for local problems.

Formal statement of derandomization: Formally, we will prove the following statement by Ghahari, Harris, and Kuhn [[GHK18](#)], using the derandomization method of conditional expectations.

Theorem 1.7 (Ghahari, Harris, and Kuhn [[GHK18](#)]). *Let Π be any local problem of randomized round complexity $t(n)$. Then, its deterministic sequential local complexity is $O(t(n))$.*

¹⁴It is important to restrict ourselves to local problems. Otherwise, consider the following silly counterexample problem: We are to mark some vertices of the input graph so that at least $n/3$ but at most $2n/3$ vertices are marked. Using randomness, this problem can be solved in 0 round complexity: every vertex simply flips a coin. But imagine trying to solve the problem deterministically: If the input graph has no edges, we are pretty screwed!

Proof. We are given a randomized local algorithm \mathcal{A} with round complexity $t(n)$ for a local problem $\Pi = (S, \mathcal{P})$ with checkability radius r . We will next describe a deterministic sequential local algorithm that writes an infinite sequence of bits into each node u of the input graph G so that if we then simulate \mathcal{A} with these bits, it solves Π .

For a vertex $u \in V(G)$, define the failure indicator $X(u)$ as the indicator random variable for the event that if we run \mathcal{A} with truly random bits, it fails at u at solving the local problem $\Pi = (S, \mathcal{P})$. By failure at u we mean that $B(u, r) \notin \mathcal{P}$. We notice that $X(u)$ depends only on the output of \mathcal{A} at an r -hop neighborhood of u , and thus it ultimately depends only on random bits in the $(r + t(n))$ -hop neighborhood of u . Moreover, the probability of failure at u is less than $1/n$ by \mathcal{A} being a randomized algorithm solving Π . This implies that $\mathbb{E} \left[\sum_{u \in V(G)} X(u) \right] < 1$.

Next, we will consider the following sequential local algorithm. We iterate over the nodes in an arbitrary order, and whenever it is a node u_k 's turn, we fix the random bits $B(u_k)$ at this node to a value $b(u_k)$ such that

$$\begin{aligned} & \mathbb{E} \left[\sum_{v \in V(G)} X(v) \mid \forall i \in [k] : B(u_i) = b_i \right] \\ & \leq \mathbb{E} \left[\sum_{v \in V(G)} X(v) \mid \forall i \in [k-1] : B(u_i) = b_i \right]. \end{aligned}$$

In words, we set the random bits so that the expected number of errors does not increase.

First, such a value $b(u_k)$ of random bits at u_k definitely exists, since the right-hand side of the above inequality simply averages over many possible instantiations of $B(u_k)$ (that is, we rely on the law of total expectation). Second, we can compute this value of random bits by looking only at the $(r + t(n))$ -hop neighborhood of u_k , since the values X_v for v outside of $B(u_k, r + t(n))$ are independent of the choice of random bits at u_k .

After this sequential algorithm with local complexity $(r + t(n))$ finishes, we have set the random bits at every vertex $u \in V(G)$ in a way that makes

$$\mathbb{E} \left[\sum_{u \in V(G)} X(u) \mid \forall i \in [n] : B(u_i) = b_i \right] < 1.$$

But all values $X(u)$ are now deterministic, so we conclude that no failure occurs if we run \mathcal{A} with these bits.

Finally, after this derandomization procedure is run, we also have to run \mathcal{A} . We will defer the discussion of how to combine two sequential local algorithms into one to [Lemma 1.9](#) that shows how to construct a sequential local algorithm with local complexity $O(t(n))$ that simulates first running the derandomization procedure and then running \mathcal{A} with the bits computed by that procedure. \square

Putting [Theorems 1.4](#) and [1.8](#) together, we get the following derandomization theorem for (distributed) local algorithms.

Theorem 1.8 (Gha ari, Harris, and Kuhn [[GHK18](#)]). *Let Π be any local problem of checkability r and randomized round complexity $t(n)$. Then, its deterministic distributed local complexity is $\tilde{O}(t(n) \cdot \log^3 n)$.*

Conversely, it's known that for some local problems, the gap between randomized and deterministic local complexity can be as large as $\Omega(\log n / \log \log n)$ [[Bal+20c](#)].

Leftover: composing sequential algorithms: We will briefly discuss how two sequential local algorithms run one after the other can be composed into a single one of larger local complexity.

Lemma 1.9 (Gha ari, Kuhn, and Maus [[GKM17](#), Observation 2.1, Lemma 2.2]). *Let $\mathcal{A}_1, \mathcal{A}_2$ be two deterministic (randomized) sequential local algorithms with local complexities $t_1(n), t_2(n)$. Then, there is a single deterministic (randomized, respectively) sequential local algorithm \mathcal{A} of complexity $2(t_1(n) + t_2(n))$ ¹⁵ that simulates the output of first running \mathcal{A}_1 , and then running \mathcal{A}_2 on the output of \mathcal{A}_1 .*

¹⁵In general, k local sequential algorithms can be simulated with complexity $2 \sum_{i=1}^k t_i(n)$.

Proof. We would like \mathcal{A} to work as follows: We iterate over the nodes and when it is the turn of a node u , we first simulate \mathcal{A}_1 for all nodes in $B(u, t_2(n))$. Then, we use the computed information to simulate \mathcal{A}_2 at u to compute the final output at u .

The only difficulty is that once u simulates \mathcal{A}_1 for a node $v \in B(u, t_2(n))$, we cannot simulate \mathcal{A}_1 at v again in the future since we want to have the guarantee that all simulations of \mathcal{A}_1 taken together correspond to a consistent iteration over the nodes of the input graph and running \mathcal{A}_1 on them.

To implement this idea, \mathcal{A} additionally stores at u the output of simulations of \mathcal{A}_1 within $B(u, t_2(n))$. The final local complexity of \mathcal{A} will be $\max(t_1(n) + t_2(n), 2t_2(n))$: \mathcal{A} starts by looking at its $2t_2(n)$ -hop neighborhood and fixing the answers of \mathcal{A}_1 for nodes $v \in B(u, t_2(n))$ at which \mathcal{A}_1 was already simulated in the past. Only then we simulate \mathcal{A}_1 for the rest of the nodes in $B(u, t_2(n))$ and run \mathcal{A}_2 after that. \square

1.5 Network Decompositions

Let us recall the definition of a network decomposition:

Definition 1.5 (Network decomposition). *A (c, d) -network decomposition of a graph G is a coloring of G with c colors. We require that vertices of each color induce a graph such that each of its connected components (that we call clusters) has diameter at most d .*

Let us now discuss constructions of network decompositions¹⁶: the missing piece in proof of [Theorems 1.4](#) and [1.8](#).

Existence of network decompositions: First of all, it is unclear whether network decomposition of the input graph, say with parameters $c, d = O(\log n)$ always exists. Let's confirm this by constructing it using the folklore sequential *ball-carving* algorithm.

Theorem 1.10 (Ball-carving algorithm). *An $(O(\log n), O(\log n))$ -network decomposition exists for any graph G .*

Proof. We will show how to construct the first color class of the network decomposition. In particular, we will find a family of vertex-sets $\mathcal{C} = \{C_1, C_2, \dots, C_t\}$, where we call each $C_i \subseteq V(G)$ a *cluster*, such that:

1. The clusters are not adjacent; i.e., there is no edge $uv \in E(G)$ with $u \in C_i, v \in C_j, i \neq j$,
2. each cluster C_i has diameter $O(\log n)$,
3. at least $n/2$ vertices are clustered, i.e., $|\bigcup_{i=1}^t C_i| \geq n/2$.

We can then construct the desired network decomposition by simply repeating this process $\log_2 n$ times and always removing the clustered vertices from the graph until every vertex is clustered.

To construct the clustering \mathcal{C} , we iterate over the vertices of G in an arbitrary order. Each time we are at a vertex u , we consider the balls $B(u, 0), B(u, 1), \dots$ of increasing radii around it. In particular, think of gradually growing larger and larger balls around u , and once the size of the ball does not at least double, i.e., once we have $|B(u, i+1)| \leq 2 \cdot |B(u, i)|$ for the first time, we let $B(u, i)$ be the next cluster of \mathcal{C} . Additionally, we remove all vertices from the boundary $B(u, i+1) \setminus B(u, i)$ from G . After the ball around u is grown, we continue this procedure with an arbitrary next vertex of G that was not explored yet.

To analyze this algorithm, first notice that no two clusters are adjacent, since after growing each cluster, we delete its boundary ([Item 1](#)).

Moreover, all balls have diameter $O(\log n)$: This is because the volume of each ball $B(u, i)$ grows as $|B(u, i)| \geq 2^i$ until we finish growing. This implies that if we are not finished in iteration $1 + \log_2 n$, the ball $B(u, 1 + \log_2 n)$ has to contain more than n vertices ([Item 2](#)), a contradiction.

Finally, we cluster at least half of the vertices, because by definition, whenever a ball stops growing, we have $|B(u, i+1) \setminus B(u, i)| \leq |B(u, i)|$, so the number of vertices unclustered because of the ball around u can be charged to the number of vertices clustered around u ([Item 3](#)). \square

¹⁶In the literature, one can very often encounter numerous variants of network decompositions with names like low-diameter clusterings, padded decompositions, or sparse neighborhood covers.

Deterministic distributed algorithms: There is a long line of work on deterministic algorithms for constructing network decompositions [Awe+89; PS92; RG20; GGR21; CG21; Roz+22a; RHG23; Gha+23]. The currently fastest deterministic algorithm for network decomposition is the algorithm by Gha ari, Grunau, Haeupler, Ilchi, and Rozho [Gha+23]. We stated its guarantees in [Theorem 1.6](#).

Perhaps the simplest $\text{poly log } n$ -round algorithm is the algorithm of Rozho and Gha ari [RG20] and its variant by Rozho , Haeupler, and Grunau [RHG23]. They both need $O(\log^7 n)$ rounds and construct clusters with diameter $O(\log^3 n)$. We will next explain the construction from [RG20] that in fact outputs clusters that only have a weaker guarantee of small *weak-diameter*. A weak-diameter of a cluster C is defined as $\max_{u,v \in C} d_G(u, v)$, as opposed to the (strong-) diameter which is defined as $\max_{u,v \in C} d_{G[C]}(u, v)$. That is, a cluster with a small weak diameter can be even disconnected. A small weak-diameter is sufficient for our applications such as the proof of [Theorem 1.4](#) and we will discuss at the end of this section how one can convert the weak-diameter guarantee to the strong-diameter one in a black box way.

Theorem 1.11 (Distributed ball-carving algorithm). *There is a local algorithm with round complexity $O(\log^7 n)$ that constructs a network decomposition with $c = O(\log n)$ colors and $d = O(\log^3 n)$ weak-diameter.*

Proof. Similarly to [Theorem 1.10](#), we show how to construct a family of clusters $\mathcal{C} = \{C_1, C_2, \dots, C_t\}$ such that:

1. No two clusters are adjacent,
2. each cluster C_i has weak-diameter $O(\log^3 n)$,
3. at least $n/2$ vertices are clustered.

Recall that in deterministic local algorithms, every vertex starts with a unique $b = O(\log n)$ -bit identifier. Our algorithm has b phases. At the beginning of the algorithm, we start with a clustering \mathcal{C}_0 that contains each node of G as a trivial one-vertex cluster. Each cluster C in our clustering is assigned an identifier $id(C)$ which is a unique b -bit string. At the beginning, the identifier of a cluster is set to be the identifier of its unique vertex.

The clustering \mathcal{C}_i evolves during the following b phases, with some vertices changing which cluster they belong to and some vertices being deleted from \mathcal{C}_i . We will prove that after each phase $i \in [b]$, the clustering \mathcal{C}_i has the following guarantees:

1. Consider the graph $G_i = G[\bigcup_{C \in \mathcal{C}_i} C]$, i.e., the graph induced by vertices that were not deleted yet. Consider any connected component of G_i . All clusters in any connected component of G_i agree on the first i bits of their unique identifier.
2. Each cluster $C \in \mathcal{C}_i$ has weak-diameter $O(i \cdot \log^2 n)$.
3. $|V(G_i)| \geq n - i \cdot \frac{n}{2b}$.

Plugging $i = b$, the inductive guarantees of [Items 1](#) to [3](#) reduce to the final desired guarantees of [Items 1](#) to [3](#) from the beginning of the proof. Thus, we only need to show how one phase of the algorithm makes sure that if we start with the guarantees for some i , they also hold for $i + 1$.

At the beginning of each phase $i + 1 \in [b]$, we classify each cluster in \mathcal{C}_i as either *active* or *inactive*, where a cluster C is active if and only if the $(i + 1)$ -th bit in $id(C)$ is equal to zero.

Next, we run a variant of the ball-carving algorithm from [Theorem 1.7](#) with $t = 10b \log_2 n$ steps. In each step of this algorithm, each vertex u from some inactive cluster first checks whether it neighbors with a vertex that is currently present in an active cluster. If there are more such neighboring vertices, u chooses an arbitrary such vertex in C and then u *proposes to join* C .

Next, each cluster C collects how many vertices proposed to join C . If there are at least $|C|/(2b)$ such vertices, then C decides to grow: All the vertices that proposed to join C leave their original cluster and join C . On the other hand, if less than $|C|/(2b)$ proposed to join C , all these vertices are *deleted*; they will not be present in \mathcal{C}_{i+1} and they will not participate in the rest of the algorithm. Note that after an active cluster C decides to delete the proposing vertices, it is neighboring only with nodes from other active clusters and thus it does not grow anymore in the rest of the phase. This finishes the description of one phase of the algorithm (see [Figure 1.4](#) for an illustration).

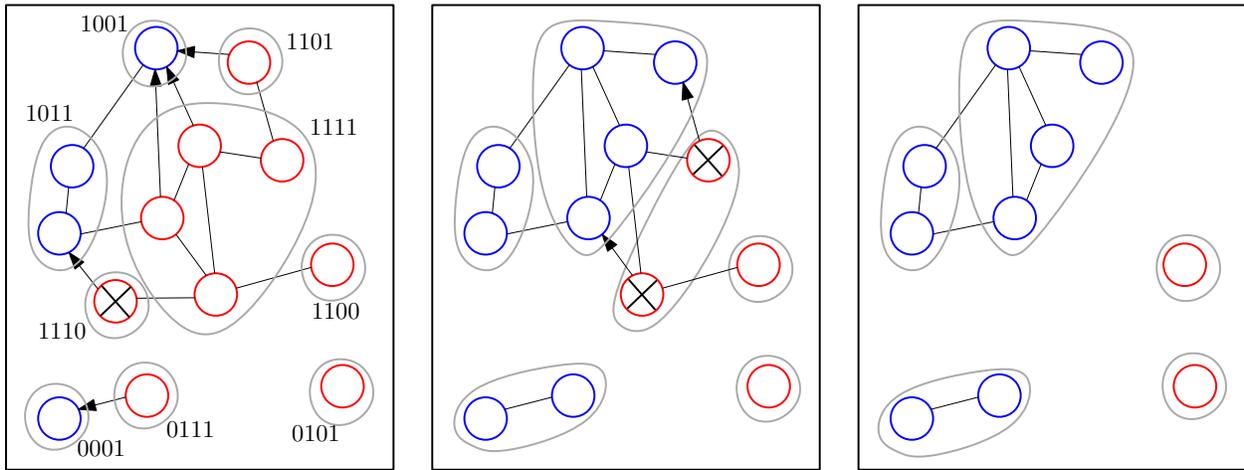


Figure 1.4: This picture illustrates the second phase of the algorithm from [Theorem 1.11](#), in a simple example graph. Left: At the beginning of the second phase, the clusters are already separated according to their first bit. In particular, if the identifier starts with 1, the cluster is in the top connected component, while if the identifier starts with 0, the corresponding cluster is in one of the two bottom connected components. In the second phase, a cluster is active (blue) if its second bit in the identifier is 0, and inactive (red) otherwise. In the first step of this phase, each red vertex proposes to join an arbitrary neighboring blue cluster, provided that there is one (the arrows in the picture). The active cluster then either decides to grow (cluster 1001) or the vertices that proposed to it are deleted (cluster 1011). Note that if a cluster decides to delete its boundary, it does not neighbor any red nodes from that point on. Middle: Note that the inactive cluster with identifier 1111 became disconnected but its weak-diameter remains the same as at the beginning of the phase. In the second step of this phase, each red vertex again proposes to join an arbitrary neighboring blue cluster. Right: The picture shows the resulting clustering after the second phase is finished. Note that the only two adjacent clusters with identifiers 1011 and 1001 will be separated in the following, third, phase, as their identifiers differ on the third position.

To analyze the phase $i + 1$, we first check that each cluster stops growing during the phase. This is true because otherwise, at the end of the phase, the cluster would have to contain more than

$$(1 + 1/(2b))^{10b \log_2 n} > n$$

nodes. In particular, the weak-diameter of each active cluster grows additively by at most $10b \log_2 n = O(\log^2 n)$, while the weak-diameter of each inactive cluster does not increase. This proves [Item 2](#).

Next, consider any connected component K of G_i . The clusters in K already agree on the first i bits of their identifier at the beginning of the phase $i + 1$. Recall that all active clusters stop growing during the phase and delete the inactive nodes on their boundary. Thus in G_{i+1} , the component K further splits into connected components K_1, K_2, \dots such that for each component K_i we have that either all of its clusters are active or all are inactive. We conclude that clusters in each connected component of G_{i+1} agree on the first $i + 1$ bits in their identifier as needed in [Item 1](#).

Finally, whenever a cluster C stops growing, we delete at most $|C|/(2b)$ nodes. Thus, during the phase we delete at most $n/(2b)$ nodes which proves [Item 3](#).

Let us discuss the running time of the algorithm. The clusters have weak-diameter $O(\log^3 n)$. Hence, each growing step can be implemented with that round complexity. There are $t = O(\log^2 n)$ steps in one phase, $b = O(\log n)$ phases, and we need to repeat the overall algorithm $c = \log_2 n$ times. We conclude that the overall round complexity is $O(\log^7 n)$. \square

Next, let us show how we can use our current knowledge of local algorithms to construct a local algorithm for network decomposition with $c, d = O(\log n)$.

Corollary 1.12. *There is a local algorithm with $O(\log^9 n)$ round complexity that constructs a network decomposition with $c = O(\log n)$ colors and $d = O(\log n)$ diameter.*

Proof. We observe that the algorithm from [Theorem 1.10](#) can be seen as a sequence of $O(\log n)$ sequential local algorithms per [Definition 1.3](#), each one with local complexity $O(\log n)$. Using [Lemma 1.9](#), we can thus view it as a single sequential local algorithm of local complexity $O(\log^2 n)$.

Thus, we can use [Theorem 1.4](#) to convert this algorithm into a local algorithm. While we phrased the guarantees of [Theorem 1.4](#) using the best available network decomposition algorithm, we can plug in the network decomposition from [Theorem 1.11](#) instead; the reduction holds even if we use a network decomposition with a weak-diameter guarantee. \square

Let us briefly go back to the observation from the proof of [Corollary 1.12](#) that the algorithm from [Theorem 1.10](#) can be viewed as a sequential local algorithm. This observation helps us to appreciate network decomposition as the “complete” problem for turning sequential local algorithms into distributed ones: On the one hand, network decomposition allows us to do this task via [Theorem 1.4](#); on the other hand, *any* way of proving that theorem leads to a distributed algorithm for network decomposition via the proof of [Corollary 1.12](#).

Randomized distributed algorithms: There is a classic randomized algorithm by Linial and Saks [[LS93](#)] that constructs an $(O(\log n), O(\log n))$ -network decomposition in $O(\log^2 n)$ distributed rounds. Let us sketch its beautiful variant by Miller, Peng, and Xu [[MPX13](#)] also known as the MPX algorithm:

In their algorithm, every vertex independently chooses a random *head start* sampled from an exponential distribution; that is, a head start of each node is 0 with probability $1/2$, 1 with probability $1/4$, and so on. Next, we run the breadth-first-search algorithm from all nodes at once with those head starts. That is, we first choose some number $T = O(\log n)$ such that with high probability, no node samples a head start larger than T . Then, we simulate a run of breadth-first search from an additional virtual node u_0 which is connected to every other node u with an oriented edge of length $T - \text{head start}(u)$. All vertices reached by the breadth-first search from the same starting node u form one cluster. The output of the algorithm are only those vertices such that all their neighbors are from the same cluster.

One can prove that the output clustering from this algorithm contains at least a constant fraction of all vertices and uses disjoint clusters of diameter $O(\log n)$, with high probability. We repeat this process $c = O(\log n)$ times to construct a network decomposition.

1.6 Bounds for Concrete Problems

In the last section of this chapter, we briefly survey known results for some concrete well-known local problems. While the technology developed in this section often automatically provides a $\text{poly log } n$ -round deterministic algorithm for a given problem, we can try to improve the complexity further: For many problems, we can get randomized algorithms with complexities closer to $\text{poly log log } n$. Typically, we want to understand the complexity as the function of both the number of vertices n , and the maximum degree Δ .

Many concrete results use certain general techniques. Randomized local algorithms with $\text{poly log log } n$ dependency typically use the shattering framework (see [Section 2.2](#)) [[Bar+16](#); [FG17](#); [CLP18](#)]. Deterministic local algorithms with $\text{poly log } n$ rounds often rely on local rounding techniques [[HKP01](#); [FGK17](#); [Har19](#); [GK22](#); [Fao+23](#)]. Many state-of-the-art lower bounds are proven using the technique of round elimination (see [Sections 2.1](#) and [2.2](#) or the survey of Suomela [[Suo13](#)]) [[Bra19](#); [Bra+16](#); [Bal+19b](#); [Bal+21a](#); [BO20](#); [Bal+22c](#); [BBO20](#)].

Maximal independent set: The maximal independent set problem is one of the most prominent problems in the area of local, distributed, parallel, or sublinear algorithms. The seminal randomized Luby’s algorithm [[Lub86](#); [ABI86](#)] with $O(\log n)$ round complexity initiated a long line of work on the problem [[Lub86](#); [ABI86](#); [Nao91](#); [BE08](#); [Gha16](#); [Bal+19b](#); [GG23](#)]. The currently fastest deterministic algorithm by Gha ari and Grunau [[GG23](#)] combines ideas related to network decompositions and local rounding techniques to achieve round complexity $\tilde{O}(\log^2 n)$. On the lower bound side, Balliu, Brandt, Hirvonen, Olivetti, Rabie, and Suomela [[Bal+19b](#)] used round elimination to show that the deterministic local complexity of maximal independent set is $\Omega(\log n / \log \log n)$. Thus, there is still a gap in our understanding:

Problem 1.13. *Is there a deterministic $\tilde{\Theta}(\log n)$ -round algorithm for maximal independent set?*¹⁷

On the randomized side, there is a randomized local algorithm for the problem by Gha ari [[Gha16](#)] using the shattering approach and the best deterministic algorithm for the maximal independent set. With the currently best deterministic algorithm, its round complexity is $O(\log \Delta) + \tilde{O}(\log^3 \log n)$.

This complexity is very close to optimal: Balliu, Brandt, Hirvonen, Olivetti, Rabie, and Suomela [[Bal+19b](#)] use round elimination to give a lower bound of $\Omega(\frac{\log \log n}{\log \log \log n})$ for randomized algorithms, on a graph with $\Delta = \Omega(\frac{\log \log n}{\log \log \log n})$. Next, a lower bound of Kuhn, Moscibroda, and Wattenhofer [[KMW16](#)] says that the local complexity of maximal independent set is $\Omega(\frac{\log \Delta}{\log \log \Delta})$. Thus, one cannot hope for an algorithm with round complexity $O(\log \Delta) + o(\frac{\log \log n}{\log \log \log n})$, or even $o(\Delta) + o(\frac{\log \log n}{\log \log \log n})$.

Maximal Matching: Maximal matching can be seen as a special case of the maximal independent set problem on the line graph, i.e., the graph constructed from an input graph G that has $E(G)$ as the vertex set and it contains an edge whenever two original edges of G share a vertex. Thus the upper bounds for the above maximal independent set apply also to this problem. On the other hand, the lower bounds of Balliu, Brandt, Hirvonen, Olivetti, Rabie, and Suomela [[Bal+19b](#)] for maximal independent set also apply to matching. There is a long line of work on the maximal matching problem and its variants [[HS12](#); [GHS14](#); [Gha16](#); [Bal+19b](#); [GG23](#); [Bal+23a](#); [DS24](#)].

Vertex coloring: Linial’s coloring algorithm [[Lin92](#)] that we will discuss more in-depth later in [Section 2.1](#) provides a coloring of graphs of degree at most Δ with $\Delta + 1$ many colors in $O(\Delta^2 + \log^* n)$ time. There is a long line of work trying to improve this complexity [[Lub86](#); [ABI86](#); [CV86](#); [GPS88](#); [Lin92](#); [SV93](#); [Joh99](#); [Nao91](#); [PR01](#); [KW06](#); [Kuh09](#); [SW10](#); [BE11](#); [BEK15](#); [Bar15](#); [Bar+16](#); [HSS16](#); [CLP18](#); [BEG18](#); [MT20](#); [Mau21](#); [GK22](#); [GG23](#); [Kuh20](#)]. The fastest known deterministic algorithm in the regime of large Δ is either the algorithm of Gha ari and Kuhn [[GK22](#)] with $O(\log n \cdot \log^2 \Delta)$ complexity, or the $\tilde{O}(\log^2 n)$ -round algorithm of Gha ari and Grunau [[GG23](#)]. The fastest randomized algorithm in this regime is by Chang, Li, and Pettie [[CLP18](#)] and uses the shattering framework: Its randomized round complexity is $\tilde{O}((\log \log n)^2)$. If we want the dependency on n to be only $O(\log^* n)$, then the fastest algorithm is the algorithm of Fraigniaud, Heinrich, and Kosowski [[FHK16](#)] with complexity $O(\sqrt{\Delta} \log \Delta + \log^* n)$. The following problem is well-known in the local algorithms community.

Problem 1.14. *Is there a $O(\text{poly log } \Delta + \log^* n)$ -round local algorithm for $\Delta + 1$ coloring?*

¹⁷This text lists some open problems in the area. For a different list, see [[Suo23](#)].

There is a long line of work on numerous variants of the coloring problem, in particular on Δ -coloring (i.e., local Brooks' theorem) [Gha+18a], defective colorings [CPS17b; FK23], frugal coloring [CPS17b], fractional colorings [BKO22], arboricity-dependent coloring [BE11], weak coloring [Kuh09; Bal+19c; Bal+19d] and other variants.

Edge coloring: When allowing $2\Delta - 1$ colors, the proper edge coloring problem becomes a special case of $\Delta + 1$ vertex coloring on line graphs, making the results for that problem applicable. There is a long line of work on $2\Delta - 1$ edge coloring trying to improve upon the state of the art for $\Delta + 1$ vertex coloring [BKO20; Bal+22d; Gha+18b]. As a highlight, Balliu, Brandt, Kuhn, and Olivetti [Bal+22d] provided a $O(\text{poly log } \Delta + \log^* n)$ -round algorithm for it, thus resolving [Problem 1.14](#) in this special case.

Another natural target is $\Delta + 1$ edge coloring guaranteed by Vizing's theorem. There is a long line of work on the problem [Cha+18; GP20; Gre23; Ber22a; BD23b; BD23a]; the fastest algorithms of Bernshteyn [Ber22a] and Bernshteyn and Dhawan [BD23b] have deterministic $\text{poly}(\Delta \log n)$ local complexity.

Spanners: Given a graph G , its subgraph H is a spanner with *stretch* α if for every two nodes u, v we have $d_G(u, v) \leq d_H(u, v) \leq \alpha d_G(u, v)$. It is known that for any k there is always a spanner with stretch $2k - 1$ and $n^{1+1/k}$ edges [Alt+93]. The task is to construct a spanner with similar parameters with a local algorithm.

This problem is a nice example of a non-local problem (certifying whether a subgraph H is a spanner has checkability radius $\Omega(n)$ in general) where it is still meaningful to ask about its local complexity. It is known that the deterministic local complexity of constructing a spanner with stretch $2k - 1$ with $O(kn^{1+1/k})$ edges is $O(k)$ [BS07; Der+08]. For $k = \text{poly log } n$, we can deterministically construct spanners with $\text{poly log } n$ stretch and $n + n/\text{poly log } n$ edges in $\text{poly log } n$ deterministic rounds [EN17; Bez+22].

Other problems: There are many other problems studied in the literature: Approximation to covering or packing integer linear programs [GKM17], approximate maximum matching and flows [NO08; Åst+10; FMU22; CP23], dominating sets and set covers [Suo10; GHS13; DKM19], low out-degree orientations [BE10; GS17], and many others.

Chapter 2

The Sublogarithmic Regime

In the previous chapter, we gained a good understanding of the fundamentals of local complexity – in essence, if we care about the complexities up to $\text{poly } \log n$, there are several equivalent definitions of it with sequential local algorithms being a particularly helpful model for designing algorithms.

In this chapter, we will restrict our attention to bounded degree graphs and local problems of sublogarithmic complexity. Something remarkable is going to happen: We will see that while, a priori, we would expect all kinds of problem complexities, there are only three distinctive classes of local problems. This *sharp threshold phenomenon* is a consequence of remarkable results known as *speedup theorems*.

As a rough roadmap for the rest of the chapter, we are going to give a more or less self-contained proof of the following [Theorem 2.1](#) in it, with [Section 2.1](#) focusing on the symmetry-breaking regime, [Section 2.2](#) focusing on the Lovász-local-lemma regime, and [Section 2.3](#) focusing on showing that there are gaps in between the regimes.

Theorem 2.1 (Classification of local problems with $o(\log n)$ complexity on bounded degree graphs). *Let us fix any Δ and the class of graphs of degree at most Δ . Then, any local problem with randomized round complexity $o(\log n)$ has one of the following three round complexities.*

1. Order-invariant regime: *The problem has $O(1)$ deterministic and randomized round complexity.*
2. Symmetry-breaking regime: *The deterministic and randomized round complexity of the problem lies between $\Omega(\log \log^* n)$ and $O(\log^* n)$ (both the deterministic and the randomized complexity is the same function).*
3. Lovász-local-lemma regime: *The problem has deterministic round complexity between $\Omega(\log n)$ and $\tilde{O}(\log^4 n)$. Its randomized round complexity is between $\Omega(\log \log n)$ and $\tilde{O}(\log^4 \log n)$.*

2.1 The Symmetry-Breaking Regime

In the bounded-degree regime, the problems of maximal independent set or the closely related problem of $(\Delta + 1)$ -coloring play a bit similar role to network decomposition, as we will see in [Theorem 2.5](#). These problems are known as *symmetry-breaking* problems. This section shows that the round complexity of these problems is $\Theta(\log^* n)$ on bounded-degree graphs.

Basic symmetry breaking problems like maximal independent set and $\Delta + 1$ coloring are closely related. Their round complexity on bounded degree graphs is $\Theta(\log^ n)$.*

2.1.1 Fast Coloring Algorithm and its Implications

We will first discuss a classical local coloring algorithm of Linial [[Lin92](#)] improving upon earlier work from [[CV86](#); [GPS88](#)]. This algorithm colors a graph of degree at most Δ with $O(\Delta^2)$ many colors in $O(\log^* n)$

rounds. We will later see in [Theorem 2.5](#) that this implies that on bounded-degree graphs, the local complexity of constructing the maximal independent set and $\Delta + 1$ coloring is $O(\log^* n)$.

The main idea behind Linial's algorithm is as follows: At the beginning of the algorithm, we will think of the unique identifiers at the vertices as an input proper coloring with colors from the (very large) range $[n^{O(1)}]$. The heart of the proof is to provide a 1-round local algorithm that takes as input a coloring with colors from k , and outputs a coloring with colors from a much smaller range $O_\Delta(\log k)$. Repeating this process $O(\log^* n)$ many rounds, we end up with a coloring with $O_\Delta(1)$ many colors.

It remains to construct a suitable 1-round algorithm. To do this, we will use the so-called cover-free families defined below.

Definition 2.2. *Given a ground set $[k']$, a family of sets $S_1, \dots, S_k \subseteq [k']$ is called a Δ -cover free family if for each set of indices $i_0, i_1, \dots, i_\Delta \in [k]$, we have $S_{i_0} \setminus \left(\bigcup_{j=1}^\Delta S_{i_j}\right) \neq \emptyset$. That is, no set in the family is a subset of the union of some Δ other sets.*

Such families can be constructed with the following parameters.

Lemma 2.3 ([KSS81; EFF82] or see [Gha20, Lemma 1.19, 1.20]). *For any k, Δ , there exists a Δ -cover free family of size k on a ground set of size $k' = O(\Delta^2 \log k)$. Moreover, if $k \leq \Delta^3$, it exists for $k' = O(\Delta^2)$.*

We can now state and analyze Linial's algorithm.

Theorem 2.4 (Linial [Lin92]). *The deterministic round complexity of constructing a coloring with $O(\Delta^2)$ many colors is $O(\log^* n)$.*

Proof. We will show a one-round algorithm that turns an input proper k -coloring into a proper $O(\Delta^2 \log k)$ coloring. Moreover, in case $k \leq \Delta^3$, we can turn the k -coloring into an $O(\Delta^2)$ coloring. We then interpret the input unique identifiers as proper coloring and apply the one-round algorithm $O(\log^* n)$ many times. A calculation shows that the number of colors then drops to $O(\Delta^2 \log \Delta)$. After one more round and using the case $k \leq \Delta^3$, the number of colors becomes $O(\Delta^2)$, as needed.

The one-round algorithm simply interprets each input color from $[k]$ as a set in a Δ -cover free family over the ground set $[k']$ for $k' = O(\Delta^2 \log k)$ (and $k' = O(\Delta^2)$ in case $k \leq \Delta^3$). Such a family exists by [Lemma 2.3](#). After each vertex u with a color S_u learns the colors S_{v_1}, \dots, S_{v_d} of its neighbors, it simply chooses any color in the set $S_u \setminus \left(\bigcup_{i=1}^d S_{v_i}\right)$ as its new color. This set is non-empty by the definition of a cover-free family, and the new coloring is proper. ¹ □

Simulation of sequential local algorithms: We notice that graph coloring can be seen as a special case of network decomposition discussed in [Chapter 1](#) where each cluster has diameter 0. In particular, coloring with a small number of colors allows us to turn sequential local algorithms into distributed local ones similarly to [Theorem 1.4](#).

Theorem 2.5. *Let \mathcal{A} be a sequential local algorithm with local complexity $t(n)$. Then, we can turn it into a distributed local algorithm of round complexity $O(\Delta^{O(t(n))} + t(n) \cdot \log^* n)$. If the sequential algorithm is deterministic, so is the distributed one.*

Proof. We follow the outline of the proof of [Theorem 1.4](#) that simulated sequential local algorithms using network decompositions. To simulate a sequential local algorithm \mathcal{A} of local complexity $t(n)$, we first construct a coloring of $G^{t(n)}$ with $O(\Delta^2(G^{t(n)}))$ colors² using Linial's algorithm from [Theorem 2.4](#). That

¹If we do not care about how the resulting number of colors depends on Δ , there is a simple one-round algorithm going back to Cole and Vishkin [CV86] that does not rely on cover-free families and leads to a coloring with $2^{O(\Delta \log \Delta)}$ many colors in $O(\log^* n)$ rounds. The algorithm turns a coloring with 2^k colors into a coloring with $2^{O(\Delta \log k)}$ colors as follows: Think of the input color of each vertex u as a k -bit string s_u . The vertex sends this color to all its neighbors. After u receives $d \leq \Delta$ strings of its neighbors, s_1, s_2, \dots, s_d , it does the following. For each received s_i , the node u identifies an index j_i where s_i differs from its own bit string s_u . Such an index exists since we assumed that we started with proper coloring. The new color of u is defined as the concatenation $s'_u = j_1 \circ s_u[j_1] \circ \dots \circ j_d \circ s_u[j_d]$. That is, the vertex u simply remembers only the values of bits of s_u on the positions where those values show that s_u differs from its neighbors. We observe that no two neighboring vertices can end up with the same string. Moreover, if we write s'_u in binary, it has only $O(\Delta \cdot \log k)$ many bits.

²The notation $\Delta(G)$ stands for the maximum degree of G .

algorithm needs $O(t(n) \cdot \log^* n)$ rounds where we multiply by $t(n)$ because one round of communication in $G^{t(n)}$ can be simulated in $t(n)$ rounds in G . Subsequently, we iterate over all colors and simulate the sequential algorithm \mathcal{A} as in [Theorem 1.4](#). That simulation takes additional $O(\Delta^2(G^{t(n)})) = \Delta^{O(t(n))}$ rounds. \square

As a corollary of [Theorem 2.5](#), we get that maximal independent set or $\Delta + 1$ -coloring can be constructed in $O(\log^* n)$ rounds on bounded-degree graphs.

Corollary 2.6. *Sequential local algorithms with local complexity $O(1)$ (such as algorithms for the maximal independent set or $\Delta + 1$ coloring) can be simulated with round complexity $O(\log^* n)$ on bounded degree graphs.*

Understanding unique identifiers: Let us now discuss the unique identifiers from the range $[n^{O(1)}]$ in the definition of deterministic algorithms. We will use our understanding of coloring to see that the strength of the model of deterministic algorithms remains the same even if the identifiers come from a much larger range like $[2^n]$ or $[2^{2^n}]$. Moreover, we will understand why identifiers are not needed in the definition of sequential local algorithms.

The following theorem will use an instance of a *fooling argument*, variants of which we will enjoy employing later on. To motivate it, notice that it is a bit awkward that the definition of a *local* algorithm talks about *globally* unique identifiers. We will next “fool” a given deterministic local algorithm solving a local problem by supplying to it a distance coloring (i.e., coloring of the power graph) with $n^{O(1)}$ many colors instead of unique identifiers. The algorithm still has to work since a failure of the algorithm for input distance coloring would imply a failure for input unique identifiers.

Theorem 2.7. *Let \mathcal{A} be a deterministic local algorithm of round complexity $t(n)$ for a local problem Π . Then, given any $s = n^{\omega(1)}$, there is a deterministic local algorithm \mathcal{A}' solving Π in $O(t(n) \cdot \log_n^*(s))$ rounds³ and assumes that the identifiers are from range $[s]$.*

Similarly, let \mathcal{A} be a sequential local algorithm of local complexity $t(n)$ for a local problem Π in a model of sequential local algorithms where each node has additionally a unique identifier from $[n^{O(1)}]$. Then there is a sequential local algorithm \mathcal{A}' of complexity $O(t(n))$ in the standard model of sequential local algorithms without any identifiers.

Proof. Let \mathcal{A} be a deterministic algorithm of round complexity $t(n)$ solving a local problem Π with local checkability r in the model where the unique identifiers are from $[n^{O(1)}]$ (the constant in the exponent is assumed to be large enough). We construct a new algorithm \mathcal{A}' that works in the less powerful model with identifiers from $[s]$ as follows. We first compute a coloring of the power graph $G^{2(t(n)+r)}$ with $n^{O(1)}$ many colors, then we simulate \mathcal{A} with that coloring as identifiers.

In particular, the coloring is constructed by iterating color reductions of Linial’s algorithm of [Theorem 2.4](#). Recall that each color reduction reduces the range of color exponentially, thus after $\log_n^*(s)$ rounds, we reduce the input coloring with colors from $[s]$ to a coloring with colors of size at most $O(\Delta^2(G^{2(t(n)+r)})) = n^{O(1)}$.

Next, we prove that \mathcal{A}' is correct. Assume that \mathcal{A}' fails to solve Π at a node u . Notice that this failure depends only on the $(t(n)+r)$ -hop neighborhood of u where the coloring constructed by \mathcal{A}' uses unique colors. In particular, we can extend this coloring of $B(u, t(n) + r)$ to a labeling of every node of the input graph with unique identifiers. The original algorithm \mathcal{A} fails to solve Π at u for these identifiers, a contradiction with \mathcal{A} being correct.

The proof of the second claim in the theorem is very similar and, in fact, easier, since $\Delta + 1$ -coloring has constant sequential local complexity and thus Linial’s color reductions are not needed. \square

As a helpful corollary of the second part of [Theorem 2.7](#), we can now see that any deterministic local algorithm \mathcal{A} solving some local problem can be converted into a deterministic sequential algorithm of the same asymptotic local complexity. This was actually not clear until now since our definition of deterministic sequential algorithms in [Definition 1.3](#) did not contain input unique identifiers.

³Here, $\log_n^*(s)$ returns how many times we need to take the logarithm of s , until its value drops below n , i.e., $\log_n^*(2^n) = 1$, $\log_n^*(2^{2^n}) = 2$ and so on.

2.1.2 Lower bound for coloring via round elimination

Finally, let us prove that the $\log^* n$ dependency for constructing maximal independent sets or colorings is necessary.

Theorem 2.8 (Naor [Nao91] and Linial [Lin92]). *The local complexity of computing $\Delta + 1$ coloring is $\Omega(\log^* n)$, even on graphs that are oriented paths.*

There are several known proofs of this theorem [Nao91; Lin92]; we will use the proof of Linial [Lin92] framed in the language of a powerful technique known as the *round elimination* (see [Suo13] for an introduction to it). In essence, given a local problem Π , round elimination is an automated technique that defines a problem Π' such that the round complexity of the fastest algorithm for Π' is exactly one round smaller than the round complexity of the fastest algorithm for Π (unless its complexity was already zero). This is very helpful for proving lower bounds: If we start with some problem Π and argue that even after t rounds of round elimination, the problem Π' that we ended up with is not solvable in zero rounds, we infer that the local complexity of Π is at least t .

Preparations for the lower bound: We will prove the lower bound in the extremely simple setup where we are promised that the input graph is an oriented path, i.e., the setup from Section 1.1. The lower bound will be for deterministic algorithms, so each node starts with a unique identifier. We will make a further restriction on the identifiers, we require that the input labeling with identifiers is increasing; that is, if for every oriented edge $e = uv$ going from u to v , we have $ID(u) < ID(v)$.

Next, let us discuss local algorithms. It will be helpful to think about them as functions in the spirit of Definition 1.2. A subtlety we need to be careful about is that in one step of round elimination, we don't want to directly convert a t -round local algorithm (that sees $2t + 1$ vertices) to a $t - 1$ round algorithm (that sees $2t - 1$ vertices). Instead, we want to convert an algorithm that sees t vertices to an algorithm that sees $t - 1$ vertices. To this end, we define an *edge-centered* $(t + 1/2)$ -round algorithm \mathcal{A} to be a local algorithm such that for an edge $e = uv$, the input to \mathcal{A} is the ball $B(uv, t - 1)$ defined as $B(u, t - 1) \cup B(v, t - 1)$. The output of \mathcal{A} is a label for the edge e . For example, $1/2$ -round local algorithm run on e has access to the two identifiers $ID(u)$ and $ID(v)$ and it maps the two identifiers to a label of e .

One half-round reduction: Here comes the heart of the proof; we will show that any given t -round local algorithm for coloring vertices with k colors can be converted to a $((t - 1/2))$ -round algorithm for coloring edges with 2^k colors.

Lemma 2.9. *Assume that for $t > 0$ we are given a node-centered t -round deterministic local algorithm \mathcal{A} that outputs a proper coloring with k colors on oriented paths with increasing unique identifiers. Then, there is an edge-centered $((t - 1/2))$ -round algorithm \mathcal{A}' that properly colors the edges with 2^k colors in the same setup. Similarly, edge-centered $(t + 1/2)$ -round algorithms can be converted into node-centered t -round algorithms.*

Proof. We will only discuss the conversion of a node-centered algorithm into an edge-centered one, the other case is very similar.

Given a t -round algorithm \mathcal{A} , we define a $(t - 1/2)$ -round edge-centered algorithm \mathcal{A}' as follows (see Figure 2.1): Given an edge $e = uv$ so that \mathcal{A}' has access to $B(uv, t - 1)$, we let \mathcal{A}' to consider all possible identifiers of the unique node $x \in B(v, t) \setminus B(uv, t - 1)$, i.e., the only node that \mathcal{A} sees but \mathcal{A}' doesn't. For each possible value of the identifier $ID(x)$ of x (we also consider the option that we are at the end of the path and the vertex x does not exist), the algorithm \mathcal{A}' simulates \mathcal{A} and records the color that \mathcal{A} outputs. The final output of \mathcal{A}' is a subset of $[k]$ that contains each color $c \in [k]$ whenever there is an identifier that makes \mathcal{A} output c . We can encode the set with k bits, i.e., we can view this set of original colors as a new color from the range $[2^k]$. This finishes the description of \mathcal{A}' .

Our task is to prove that \mathcal{A}' returns a proper coloring of edges. Here is what this question reduces to: Consider any increasing labeling of the oriented path with identifiers and consider any three consecutive nodes u, v, w (see Figure 2.2). We again let x be the unique node in $B(v, t) \setminus B(uv, t - 1)$ and y the subsequent neighbor of x . Consider the new colors $C_1, C_2 \in [2^k]$ that \mathcal{A}' outputs on the two edges uv and vw . We need to prove that the two colors C_1, C_2 are different. We do that by focusing on the original color $c \in [k]$ that \mathcal{A} outputs at v .

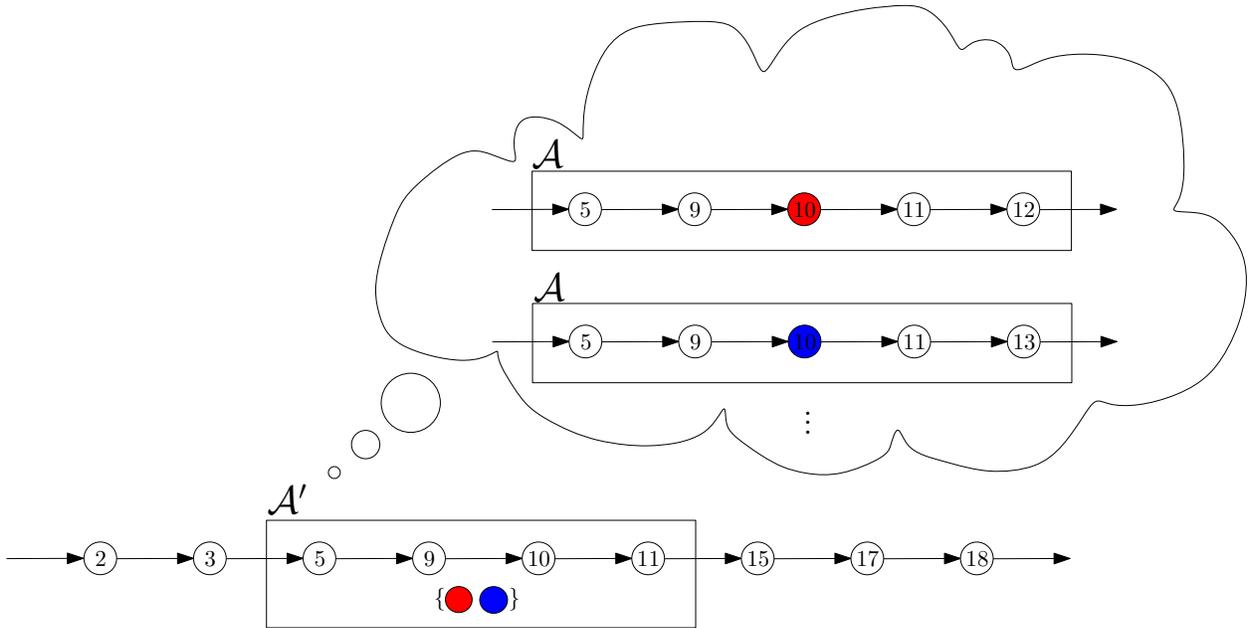


Figure 2.1: The definition of the 1.5-round algorithm \mathcal{A}' derived from a 2-round algorithm \mathcal{A} : The algorithm \mathcal{A}' considers all compatible one-node extensions of its neighborhood containing 4 nodes to 5-node-neighborhoods (in the picture, this corresponds to the following node having the identifier $12, 13, \dots, n^{O(1)}$, and the possibility that the node does not exist). The algorithm considers all possible answers that \mathcal{A} returns for those neighborhoods (the picture shows that the red and the blue color are two of those possible answers). The color that \mathcal{A}' uses to color the edge it is centered on is simply the set of all colors that \mathcal{A} returns for some extension (in the picture, it is the set containing the red and the blue color).

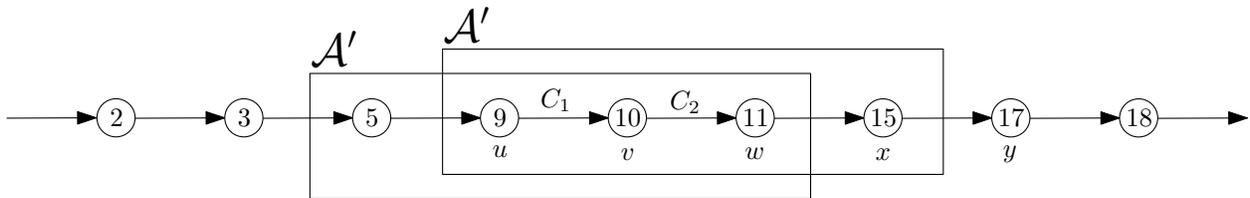


Figure 2.2: The situation from the proof of correctness of \mathcal{A}' : We need to argue that the two colors C_1, C_2 are different. To do so, we define c to be the color that \mathcal{A} outputs on the input $(5, 9, 10, 11, 15)$. By definition, we have $c \in C_1$. On the other hand, if $c \in C_2$, we get an existence of an identifier $ID'(y)$ such that \mathcal{A} returns c on the input $(9, 10, 11, 15, ID'(y))$. But then we consider the input $(5, 9, 10, 11, 15, ID'(y))$ and notice that \mathcal{A} colors two consecutive vertices with the same color c , a contradiction with its correctness.

On one hand, notice that $c \in C_1$ because the edge-centered algorithm \mathcal{A}' run at uv considers the actual identifier $ID(x)$ as a possibility and thus includes c in C_1 .

On the other hand, assume for contradiction that $c \in C_2$. That would imply that there is a certain identifier $ID'(y)$ that makes the output of \mathcal{A} at w to be the color c . But now consider changing the valid increasing labeling of identifiers that we started with by letting the identifier of y be equal to $ID'(y)$. Let's look at all the nodes in $B(vw, t)$. The identifiers on them are still a valid increasing identifier sequence.⁴ But after extending $B(vw, t)$ and its identifiers to a path on n vertices labeled with increasing identifiers, \mathcal{A} fails to solve k -coloring on that graphs since it outputs the same color c at both v and w , a contradiction. \square

We can now prove [Theorem 2.8](#) as follows: Assuming an $t = o(\log^* n)$ -round algorithm for 3-coloring of oriented paths, we apply [Lemma 2.9](#) $2t$ times until we end up with a 0-round algorithm \mathcal{A}_0 that colors oriented paths with less than n colors. But such \mathcal{A}_0 is simply a function mapping an input identifier from range $[n^{O(1)}]$ to a color in the smaller range $[n]$. Using the pigeonhole principle, we can find two identifiers that \mathcal{A}_0 maps to the same color and argue that if these two identifiers happen to be present at two neighboring nodes, \mathcal{A}_0 fails to output proper coloring.⁵

2.2 The Lovász Local Lemma Regime

We will next discuss Lovász local lemma, a very expressible local problem closely related to the third regime of problems from [Theorem 2.1](#).

Definition of Lovász local lemma: Lovász local lemma is the following very general local problem⁶. The problem is formally defined for bipartite graphs in which nodes of one part are labeled as *random variables* and the nodes of the other part are labeled as *bad events*. We denote the maximum degree of random variable nodes as $\Delta_{r.v.}$ and the maximum degree of bad event nodes as $\Delta_{b.e.}$. Each node u labeled as a random variable is additionally labeled with a probability space Ω_u . To simplify discussions, we will without loss of generality assume that each probability space is an infinite list of random bits (i.e., it is the uniform distribution on $[0, 1]$). Next, each node v labeled as a bad event that neighbors with nodes u_1, \dots, u_d with $d \leq \Delta_{b.e.}$ is additionally labeled with an event \mathcal{E}_v on the space $\prod_{i=1}^d \Omega_{u_i}$.

Often, it is useful to work only with a graph induced by bad-event nodes where two bad-event nodes are connected if they share a common random variable. In that case, we will talk about the *dependency-graph* formulation of Lovász local lemma and use Δ to denote its maximum degree. On the other hand, the setup with a bipartite graph with both bad-event nodes and random-variable nodes will be denoted as the *variable-event-graph*.

The crucial ingredient to the Lovász local lemma is a requirement on the bad events that, roughly speaking, says that we can use the union bound in the dependency-graph neighborhood of every bad event \mathcal{E}_v and conclude that with positive probability, neither \mathcal{E}_v , nor its neighboring bad events occur. Namely, in the *tight* version of Lovász local lemma, we are given a promise⁷ that each bad event \mathcal{E}_v has probability at most p where p is defined by the following *Lovász local lemma criterion*:

$$p \cdot (\Delta + 1) \leq 1/e. \tag{2.1}$$

A foundational result of Erdős and Lovász [[EL75](#)] is that even in the tight formulation, it is always possible to solve any instance of the local lemma problem, by which we mean that one can instantiate random variables so that no bad event occurs.

⁴This is the place in the argument where we need increasing and not just unique identifiers: With unique identifiers, the leftmost and rightmost node of $B(vw, t)$ could now have the same identifier.

⁵Round elimination may be used both to prove lower bounds and to construct algorithms. In particular, it can be used to derive that the round complexity of 3-coloring paths is $\frac{1}{2} \log^* n \pm O(1)$ [[RS15](#)].

⁶The Lovász-local-lemma problem does not quite satisfy the requirements for the local problem as we defined it in [Definition 1.1](#), but, morally speaking, it can be seen as such.

⁷Our definition of a local problem from [Definition 1.1](#) does not allow “promises” on input labels. To turn local lemma into a proper instance of a local problem, we can postulate that if the promise is not satisfied, the local solution around the node is not bound by any constraints.

A C -relaxed (or just polynomially-relaxed) version of Lovász local lemma, which is a bit more relevant to our applications, only requires that

$$p \cdot \Delta^C \leq 1 \tag{2.2}$$

This section will show both fast algorithms and lower bounds for the polynomially-relaxed Lovász local lemma.

Lovász local lemma is a very versatile and expressible local problem. Its deterministic round complexity on bounded degree graphs is $\Theta(\text{poly log } n)$, and its randomized complexity $\Theta(\text{poly log log } n)$.

2.2.1 Fast Algorithms for the Local Lemma

We first discuss how to solve any instance of the local lemma with a fast deterministic local algorithm. To this end, we will start with a randomized algorithm that we later derandomize by [Theorem 1.8](#). In a breakthrough result in the area of constructive algorithms for the local lemma, Moser and Tardos [[MT10](#)] presented an algorithm for it in the tight formulation. Moreover, they presented a parallel variant of their algorithm that can be interpreted as a local algorithm with round complexity $O(\log^2 n)$. This complexity was later improved by Chung, Pettie, and Su [[CPS17b](#)] to $O(\log n)$ rounds on bounded degree graphs. Plugging their result to [Theorem 1.8](#), we obtain the following theorem.

Theorem 2.10. *The deterministic round complexity of solving any instance of the tight version of Lovász local lemma on bounded degree graphs is $\tilde{O}(\log^4(n))$.*

Even faster randomized algorithm: Next, we will show that there is an even exponentially faster randomized algorithm for the relaxed version of the local lemma. We will discuss the algorithm of Fischer and Ghaheri [[FG17](#)] with round complexity $\text{poly log log } n$ which is based on the technique of shattering, a successful technique (cf. [Section 1.6](#)) that goes back to the work on algorithmic local lemma by Beck [[Bec91](#)].

The idea is as follows. Our algorithm will have two phases. In the first phase, we use a sequential local algorithm with constant complexity to do the following. Given an instance of the local lemma, the algorithm fixes *most* of the random variables in such a way that *most* bad events are satisfied (i.e., all random variables relevant for that bad event are fixed and it does not occur). The fraction of fixed random variables and satisfied events is in particular $1 - 1/\Delta^{O(1)}$. The price for this outcome is that remaining, unfixed, bad events have their probability slightly increased from $1/\Delta^C$ to $1/\Delta^{C'}$ for some $C' < C$.

Fortunately, we have an additional guarantee on the unfixed bad events. Due to the very small locality of the algorithm in the first phase, we can use an independence-like argument to show that the size of the connected components of unfixed bad events is $O(\Delta^{O(1)} \log n)$, with high probability (see [Figure 2.3](#)). We also say that the graph *shatters* into small components. We can thus run the best deterministic algorithm from [Theorem 2.10](#) as the second phase of the algorithm, to solve the remaining instance of the C' -relaxed local lemma. This second phase takes $\text{poly log } (\Delta \log n)$ rounds which is also the round complexity of the overall algorithm.⁸ We will now make the above discussion formal.

Theorem 2.11 (First phase of Fischer and Ghaheri [[FG17](#)]). *There is a randomized sequential local algorithm \mathcal{A} of constant local complexity that gets as input a variable-event graph of an instance of C -relaxed Lovász local lemma for large enough C . The algorithm fixes some bits of each random variable to concrete values so that conditioned on those fixed bits, we have the following two properties:*

1. *Each bad event has probability at most $1/(3\Delta)$.*
2. *Up to $1/\text{poly}(n)$ error probability, the residual dependency graph induced by bad events with non-zero probability has connected components of size $O(\Delta^3 \log n)$.*

⁸Our example algorithm from [Section 1.1](#) can be seen as a simple shattering algorithm.

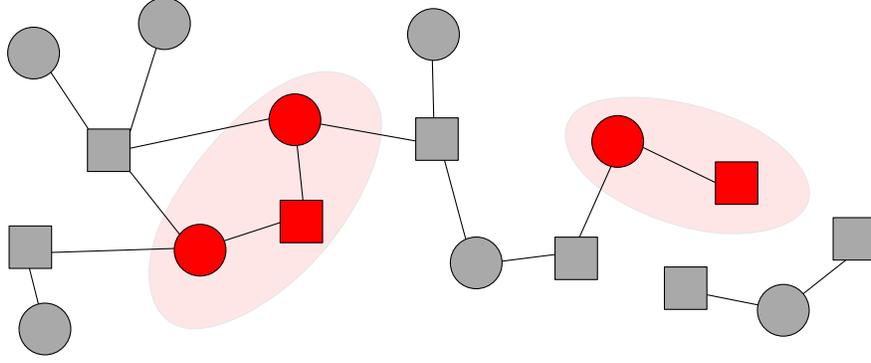


Figure 2.3: This picture shows an example variable-event graph corresponding to an instance of Lovász local lemma; the circles represent bad events and squares represent random variables. The picture shows the situation after the first phase of Fischer-Ghaffari algorithm (Theorem 2.11). Most random variables are set to a fixed value (grey squares). A small proportion of the random variables remains unset (red squares) and the bad events neighboring with an unset random variable (red circles) form connected components of diameter $O(\log n)$.

*Proof.*⁹ The algorithm \mathcal{A} is defined as follows. We iterate over the random variables and for each random-variable node u with neighboring bad-event nodes v_1, \dots, v_d , we perform the following process. We sample the random bits of Ω_u one by one. While we do that, we consider the probabilities of neighboring bad events $P(\mathcal{E}_i)$ for each $1 \leq i \leq d$. The first time it happens that for some i we have $P(\mathcal{E}_i) > 1/(6\Delta)$, i.e., the bad-event probability crosses the *dangerous threshold* of $1/(6\Delta)$, we stop sampling and do the following.

We label u , together with all other random variables neighboring v_i as *frozen* (if more v_i 's jumped over the dangerous threshold at the same time, we do this to all of them). We never sample bits of frozen random variables in the future. In particular, we stop sampling bits of u and continue with the next unfrozen random variable in the arbitrary order of our sequential local algorithm. This finishes the description of the algorithm.

To see that each bad event has probability at most $1/(3\Delta)$ during any point throughout the algorithm, we notice that if an event \mathcal{E} of probability $P(\mathcal{E}) = p$ depends on a single bit of randomness b , we have, by Markov's inequality, that $P(\mathcal{E}|b = 0), P(\mathcal{E}|b = 1) \leq 2p$. That is, the bad event probability in our process never increases multiplicatively by a factor larger than 2, which together with the definition of the dangerous threshold implies the desired bound.

To understand the size of connected components in the residual dependency graph (i.e., components of bad-event nodes of non-zero probability after we set the random bits), consider any fixed set S of bad-event nodes with the following two properties:

1. (*independence*) no two nodes $v_1, v_2 \in S$ are neighboring in the dependency graph G ,
2. (*connectivity*) S is connected in G^6 .

We will next prove that with high probability, no such set S of size $\Omega(\log n)$ survives to the residual graph.

First, we use the independence property to prove that the probability that all nodes in a fixed set S crossed the dangerous threshold during our process is exponentially small in the number of nodes $|S|$. Let us contemplate the behavior of the algorithm \mathcal{A} with respect to any bad-event node $v \in S$ and the random variables u_1, \dots, u_d neighboring v in the variable-event graph. We note that if bits b_1, \dots, b_{k-1} were already sampled and b_k is sampled next, we have $E_{b_k} [P(\mathcal{E}_v|b_1, \dots, b_k)] = P(\mathcal{E}_v|b_1, \dots, b_{k-1})$. Hence, viewing $P(\mathcal{E}_v)$ as a random variable that depends on bits sampled from $\Omega_u, u \in V(G)$, we conclude that its expectation is at most $1/\Delta^C$ and we can thus use Markov's inequality to conclude that the probability of $P(\mathcal{E}_v)$ crossing the dangerous threshold is at most $\frac{6\Delta}{\Delta^C} = 6/\Delta^{C-1}$. Moreover, we notice that if we first set the randomness of nodes in $V' = V(G) \setminus (u_1 \cup \dots \cup u_d)$ to whatever values, we can still make above argument and conclude

⁹The original algorithm and its analysis by Fischer and Ghaffari [FG17] contains subtle flaws. Here, we follow the proof by Maus [Mau].

that for all ways of setting $\Omega_u = \omega_u$ for all $u \in V'$ we have

$$P(\mathcal{E}_v | \forall u \in V' : \Omega_u = \omega_u) \leq 6/\Delta^{C-1}.$$

Since we assumed that S is an independent set of bad-event nodes in the dependency graph, we can thus inductively prove that the probability of all nodes in S crossing the dangerous threshold is at most $(\frac{6}{\Delta^{C-1}})^{|S|}$.

We next count the number of possible sets S of size t that satisfy the connectivity property: Each such S can be specified by fixing any node $u \in S$, and then specifying how one can walk for $2(|S| - 1)$ steps in G^6 so that the walk defines a spanning tree of $G^6[S]$. Hence, the number of sets S of size t is at most $n \cdot (2\Delta^6)^{2(t-1)}$. We can thus upper bound the existence of some connected surviving set S of size t as

$$n \cdot (2\Delta)^{12t} \cdot \left(\frac{6}{\Delta^{C-1}}\right)^t.$$

Choosing $C = O(1)$ and $t = O(\log n)$ large enough, we conclude that the size of this expression is at most $1/n^{O(1)}$.

Finally, for $t = O(\log n)$ from the above argument, consider any connected subset S of G of size at least $(2\Delta)^3 \cdot t$ and assume that S survived to the residual dependency graph. Then, we can find its subset $S' \subseteq S$ of size at least $|S'| \geq t$ where S' is independent in G^3 , yet S' is connected in G^4 . To see this, consider a greedy algorithm that starts with $S' = \{u\}$ for arbitrary $u \in S$, iterates through vertices of S and while we still can add at least one node $v \in S$ to S' and keep the independence property, we choose any v with distance 4 to S' and add it to S' . One can see that the final set S' has to be connected in G^4 . Also, $|S'| \geq t$ because adding a vertex to S' disqualifies at most $(2\Delta^3)$ vertices to be added to S' in the future.

If all nodes of S survived to the residual dependency graph, it means that every node in S' has to have a neighboring node that crossed the dangerous threshold. The set S' thus gives rise to a set S'' of size $|S''| = |S'| \geq t$ of nodes that all crossed the dangerous threshold. Moreover, by S' being independent in G^3 , we conclude that S'' is independent in G . Since S' was connected in G^4 , S'' is connected in G^6 . The set S'' thus satisfies the requirements of a set that, as we have proven, does not occur in the residual graph with high probability and we can thus conclude the same for the original connected set S . \square

Putting [Theorem 2.11](#) (simulated as a distributed algorithm by [Theorem 2.5](#)) together with [Theorem 2.10](#), we conclude that the following result holds.

Theorem 2.12 (Fischer and Ghahari [[FG17](#)]). *There exists a constant C such that the randomized round complexity of any instance of the C -relaxed Lovász local lemma on bounded-degree graphs is $\tilde{O}(\log^4 \log n)$.*

The dependency on Δ in the round complexity of Fischer-Ghahari algorithm is polynomial. It is unknown whether this polynomial dependency can be improved to logarithmic.

Problem 2.13. *Is there a randomized local algorithm for relaxed Lovász local lemma with round complexity $O(\text{poly log } \Delta + \text{poly log } \log n)$?*

Conjecture of Chang and Pettie [[CP19](#)]: We note that Chang and Pettie [[CP19](#), Conjecture 1] conjecture that any instance of the relaxed local lemma can be solved with randomized $O(\log \log n)$ complexity.

Problem 2.14 (Chang-Pettie Conjecture). *Is it true that the randomized round complexity of C -relaxed Lovász Local Lemma for some large enough C is $\Theta(\log \log n)$ on bounded degree graphs?*

We will see later in [Section 2.3.2](#) that the randomized round complexity $O(\log \log n)$ implies the deterministic round complexity of $O(\log n)$ via [Theorem 2.21](#).

Self-contained algorithm: A bit unfortunate property of [Theorem 2.12](#) is that it relies on the randomized entropy-compression algorithms from Moser and Tardos [[MT10](#)] or Chung, Pettie, and Su [[CPS17b](#)]. However, we can prove [Theorem 2.12](#) also in a self-contained way: First, we observe that after the first phase of the Fischer-Ghahari algorithm finished, each surviving bad event simply collects the whole connected component of the residual graph; then we solve the problem in each residual component by applying

the standard, existential, Lovász local lemma. This new randomized algorithm has complexity $O(\log n)$ on bounded degree graphs since this is the maximum diameter of residual components, with high probability. We can derandomize this algorithm using [Theorem 1.8](#) and use the resulting deterministic algorithm instead of the algorithm of Chung, Pettie, and Su [[CPS17b](#)] in the second phase of our shattering algorithm. This way, we get a simpler and more self-contained algorithm. Its downside is a worse value of C and a worse dependence of round complexity on Δ .

2.2.2 Lower Bound via Round Elimination

Next, we will show that the deterministic local complexity of solving a certain specific instance of Lovász local lemma is $\Omega(\log n)$. Later, we will prove in [Theorem 2.21](#) that this also implies a randomized lower bound of $\Omega(\log \log n)$, thus showing that the round complexity of Fischer-Ghaheri algorithm from [Theorem 2.12](#) is close to tight.

The problem we choose for the lower bound is *sinkless orientation*. We will define the problem only on trees of degree at most Δ where it is already hard. The task is to orient all the edges of the input tree so that no node is a *sink*, which is defined as a node such that all Δ neighboring edges point towards it (nodes of degree less than Δ are not constrained in any way).

Sinkless orientation can be seen as a specific instance of the local lemma: orienting each edge randomly corresponds to a random variable at each edge. Then, a vertex becoming a sink corresponds to a bad event of probability $2^{-\Delta}$. For large enough Δ , this is much smaller than the polynomial criterion $1/\Delta^C$ from [Equation \(2.2\)](#) which makes this problem a valid instance of the local lemma. We will next prove the following theorem.

Theorem 2.15 (Brandt, Fischer, Hirvonen, Keller, Lempiäinen, Rybicki, Suomela, and Uitto [[Bra+16](#)]). *For any constant Δ , the sinkless orientation problem has deterministic round complexity $\Omega(\log n)$ on the class of trees of degree at most Δ .*

Preparations: As in [Theorem 2.8](#), we will want to replace the uniqueness of identifiers with local constraints that imply that they are unique. In the proof of [Theorem 2.8](#), we worked with identifiers that were monotonically increasing (which implied they were unique), this time we will work with identifiers that are consistent with a so-called *ID graph* [[Kor+21](#); [Bal+23c](#); [Bra+22](#)].¹⁰

Definition 2.16 (ID graph). *Given a parameter Δ , an ID graph H is a graph on a set $[n]$ that we associate with unique identifiers. Every edge of the graph is colored with one of Δ many colors and we write H_i for the graph induced by the i -th color. We require that:*

1. *The girth of H , i.e., the length of the shortest cycle in H , is at least $\gamma \log_{\Delta} n$ for some fixed $\gamma > 0$.*
2. *Each independent set of each H_i has less than n/Δ vertices.*

Such a graph exists, which can be proven using the same argument as how one proves that high-girth high-chromatic graphs exist [[Bra+22](#); [Kor+21](#)].

We will fix any ID graph and work in a model *relative* to that ID graph. Here is what that means. First, we will assume that the input graph G is always a tree of degree at most Δ such that its edges are, moreover, properly Δ -colored on the input (i.e., each vertex is incident to edges of different colors). Additionally, we require that if two nodes u, v neighbor in G with an edge of color i , then their identifiers $ID(u), ID(v)$ neighbor in H_i .

We notice that this is a local constraint on input identifiers that does not imply they are unique. However, notice that whenever two vertices $u, v \in V(G)$ have the same identifier, we can consider the path between u and v in G and how it maps to a walk in H that starts in $ID(u)$ and finishes in $ID(v) = ID(u)$. The proper edge-coloring of G implies that the walk never goes from $x \in V(H)$ to $y \in V(H)$ and then back to x in the subsequent step. This implies that the walk contains a cycle, thus the distance of u and v is at least as large as the girth of H . That is, the identifiers are unique up to a large distance which is pretty much the same as them being unique (cf. [Theorem 2.7](#)).

¹⁰The usual usage of round elimination is for randomized algorithms instead of using the ID graph, but this would require a longer setup and more calculations.

Similarly to the lower bound of [Theorem 2.8](#), we will work with node-centered and edge-centered algorithms. For node-centered algorithms, solving sinkless orientation means that the algorithm outputs one of Δ many input edge colors at each node. Outputting a color i means that u decides that the edge of the color i goes outwards from u . The local constraint on this output vertex-coloring is that no two neighboring nodes should select the same edge going in between. We will call this variant of the problem *edge-grabbing*. On the other hand, edge-centered algorithms will solve the sinkless orientation as we described the problem: Each edge simply outputs how it is oriented and we have a local constraint at each vertex, requiring that it has at least one outgoing edge.

Round elimination: We proceed with performing the actual round elimination. Notice that the spirit of the following proof is very similar to the proof of [Lemma 2.9](#).

Lemma 2.17. *Assume that we are given a node-centered deterministic t -round local algorithm \mathcal{A} of round complexity at most $t \leq \log_{\Delta} n - 1$ that solves the edge-grabbing problem on trees of degree at most Δ relative to some fixed ID graph H . Then, there is a $t - 1/2$ -round edge-centered deterministic local algorithm \mathcal{A}' that solves sinkless orientation in the same setup.*

Similarly, $t + 1/2$ -round edge-centered algorithm for sinkless orientation implies a t -round node-centered algorithm for edge-grabbing.

Proof. We prove just the first part of the statement, the proof of the second part is similar, and doing it is a good exercise.

We start with any node-centered algorithm \mathcal{A} with round complexity t that solves the edge-grabbing problem. We define the edge-centered algorithm \mathcal{A}' of round complexity $t - 1/2$ as follows. For an edge $e = uv$, the algorithm first considers all the possible extensions of the (known) ball $B(uv, t - 1)$ to the ball $B(u, t)$ that is known to \mathcal{A} when it is run on u . By an extension, we mean first how the graph looks like (e.g., maybe some vertices on the boundary of $B(uv, t - 1)$ turn out to be leaves), and second, what the identifiers are (they have to be consistent with H). We consider all valid extensions and if at least one of them leads to \mathcal{A} grabbing the edge uv from u , then \mathcal{A}' orients the edge uv from u to v . After this is done, the algorithm makes an analogous reasoning for v ; again, whenever at least one extension of $B(uv, t - 1)$ to $B(v, t)$ decides to grab the edge uv , \mathcal{A}' orients it from v to u . If no vertex ever decides to grab uv , \mathcal{A}' decides to orient it arbitrarily. This finishes the description of \mathcal{A}' (see [Figure 2.4](#)).

To make sure that \mathcal{A}' is well-defined, we need to prove that it never happens that there is an extension of $B(uv, t - 1)$ to both $B(u, t)$ and $B(v, t)$ such that \mathcal{A} run on $B(u, t)$ decides to grab the edge uv and run on $B(v, t)$, it decides to grab the edge vu . To see this, notice that putting the two extensions $B(u, t), B(v, t)$ together, the graph $B(uv, t)$ has at most n nodes and its identifiers respect H .¹¹ We observe that in case \mathcal{A}' is not well-defined, \mathcal{A} fails to solve edge-grabbing on $B(uv, t)$, a contradiction with the assumed correctness of \mathcal{A} .

Moreover, the algorithm \mathcal{A}' solves sinkless orientation: For any node u of full degree Δ , the original algorithm \mathcal{A} decided to grab a certain edge uv . When we run \mathcal{A}' on uv , \mathcal{A}' will by definition orient this edge from u to v , thus u cannot be a sink. \square

Finishing the proof: Let us finish the proof of [Theorem 2.15](#).

Proof of [Theorem 2.15](#). Consider an input graph G which is a branching tree with $\varepsilon \log n$ layers and every non-leaf vertex has degree Δ . Note that for small enough $\varepsilon > 0$, the girth property of the ID graph implies that any labeling of G with identifiers from $[n]$ that respects a fixed ID graph H has unique identifiers. On the other hand, G has $O(\Delta^{\varepsilon \log n})$ nodes, so the range from which the identifiers are coming is $|V(G)|^{O(1)}$. Therefore, a deterministic local algorithm for sinkless orientation on G implies a local algorithm for that problem on G with identifiers consistent with H .

But [Lemma 2.17](#) shows that any $o(\log n)$ -round algorithm that works relative to H can be sped up to 0 round complexity. A 0-round algorithm \mathcal{A}_0 is simply a function that maps an input identifier to one of Δ many colors, i.e., which edge the vertex decides to grab. We can thus think of \mathcal{A}_0 as a coloring of H with Δ many colors. But notice that for any such vertex coloring, we can consider the largest color class i that has at least n/Δ colors. Then, we use the independence property of ID graphs from [Definition 2.16](#) to infer that

¹¹This part of the argument needs to work with the ID graph instead of unique identifiers.

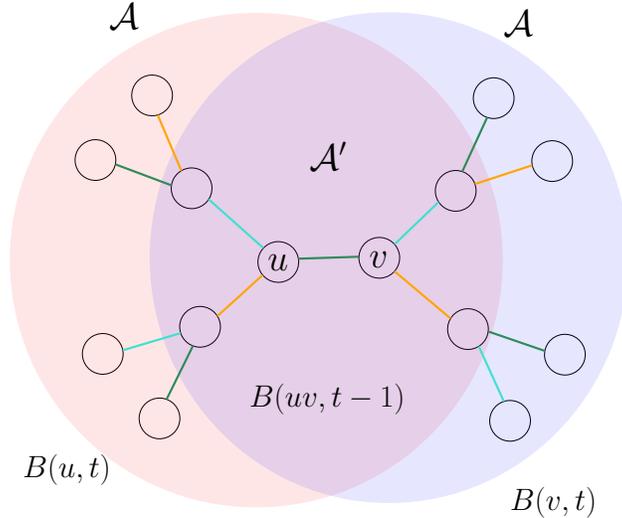


Figure 2.4: The picture shows the definition of the algorithm \mathcal{A}' for $t = 2$. Notice that our setup is a large Δ -regular tree with edges colored with Δ colors. The algorithm \mathcal{A}' has access only to the identifiers in $B(uv, t - 1)$ (the intersection of the red and the blue ball). To decide on the orientation of the edge uv , the algorithm considers all possible identifier extensions of $B(uv, t - 1)$ to $B(u, t)$ (the red ball), and if at least one such extension makes \mathcal{A} run at u grab the edge uv , \mathcal{A}' orients that edge towards v . We analogously orient this edge towards u if at least one extension of $B(uv, t - 1)$ to $B(v, t)$ (the blue ball) makes \mathcal{A} run at v grab the edge uv . We notice that it cannot happen that \mathcal{A}' wants to orient the edge uv in both directions: that would imply the existence of an identifier-labeling of $B(uv, t)$ on which \mathcal{A} is incorrect.

the set of vertices of color i cannot be independent in H_i , i.e., there is an edge uv in H where both u and v , as well as the edge uv are colored by the color i . We thus found two vertices that may be neighboring in the input graph G and the algorithm \mathcal{A}_0 decides to grab the edge connecting them from both endpoints of that edge, a contradiction with \mathcal{A}_0 being correct. \square

Sinkless orientation as the “simplest hard problem”: Notice that if we formulate sinkless orientation as an instance of Lovász local lemma, the bad event probability is equal to $2^{-\Delta}$, that is, the bad event probability is exponentially small compared to the polynomial guarantee in the relaxed criterion of Equation (2.2). It turns out that this is the threshold where an instance of the local lemma is still “hard”.

Theorem 2.18 ([Bra+16; BMU19; BGR20]). *On one hand, there is an instance of Lovász local lemma (namely sinkless orientation) with the criterion $p \cdot 2^\Delta \leq 1$ that has deterministic round complexity $\Omega(\log n)$. On the other hand, any instance of Lovász local lemma with the criterion $p \cdot 2^\Delta < 1$ has deterministic round complexity $O(\log^* n)$.*

A similar threshold phenomenon holds also if we work in the variable-event graph.

Theorem 2.19 ([FG17, Theorem 3.5], [Ber23b, Corollary 1.8]). *On one hand, there is an instance of Lovász local lemma¹² with the criterion $p \cdot \Delta_{r,v}^{\Delta_{b,e}} \leq 1$ has deterministic round complexity $\Omega(\log n)$. On the other hand, any instance of Lovász local lemma with the criterion $p \cdot \Delta_{r,v}^{\Delta_{b,e}} < 1$ has deterministic round complexity $O(\log^* n)$.*

2.3 Speedups and Slowdowns

The following section covers speedup and slowdown theorems which are at the heart of why we understand that there are sharp thresholds in the local complexities in Theorem 2.1.

¹²The instance is sinkless orientation on trees where one color class has degree $\Delta_{r,v}$ and the other has degree $\Delta_{b,e}$. One formulates it as an instance of the local lemma by letting each vertex of one color class grab a random outgoing edge. The proof that this variant of sinkless orientation is still hard seems to be missing in the literature.

The elegant idea behind speedups and slowdowns is that we can simply “lie” to algorithms about the size of the input graph, an idea closely related to the fooling argument we have already seen in the proof of [Theorem 2.7](#). To understand this technique, it may be useful to briefly recall that in [Definition 1.2](#), we defined a local algorithm as a function that takes two inputs. Firstly, it is n , the size of the input graph, and secondly, it is a $t(n)$ -hop neighborhood of a vertex that is additionally labeled by unique identifiers or random strings. We will use the notation \mathcal{A}_n to denote the algorithm \mathcal{A} when the first input is n , i.e., we will view \mathcal{A} as a sequence of functions $\mathcal{A}_1, \mathcal{A}_2, \dots$. In this section, we will contemplate what happens if we run \mathcal{A}_n on a graph of size $n' \neq n$. If $n' > n$, we are “speeding up” \mathcal{A} , while if $n' < n$, we are “slowing it down”.

2.3.1 Slowdowns

Let’s first see why slowdowns may be useful. As a first application, let us recall that we defined deterministic and randomized round complexities in [Definition 1.1](#) by requiring unique identifiers from the range $[n^{O(1)}]$ or error probability at most $1/n^{O(1)}$. We will next see that for algorithms with sufficiently small round complexity, we can replace $n^{O(1)}$ by n in the definition without changing its strength. For deterministic algorithms, this complements [Theorem 2.7](#) that shows how to replace very large identifiers with polynomially-sized ones. This follows by plugging in $f(n) = n^{O(1)}$ into the following slowdown theorem.

Theorem 2.20 (Chang, Kopelowitz, and Pettie [[CKP19](#)]). *Let f be any increasing function with $f(n) \geq n$, let Π be any local problem, and let us use $t_{f(n)}(n)$ to denote the deterministic (randomized) round complexity of solving Π if the input identifiers are from the range $[f(n)]$ (the error probability is required to be $1/f(n)$, respectively). Then,*

$$t_n(n) \leq t_{f(n)}(n) \leq t_n(f(n)).$$

The theorem holds for local complexities defined with respect to any subclass of graphs closed on adding isolated vertices. ¹³

Proof. We will prove just the deterministic version of the theorem. Notice that $t_n(n) \leq t_{f(n)}(n)$ follows directly from our assumption $f(n) \geq n$ and the definition: any algorithm expecting identifiers from the range $[f(n)]$ certainly works if they happen to be from the smaller range $[n]$.

Next, consider any deterministic algorithm \mathcal{A} that solves Π if the identifiers are from $[n]$ in round complexity $t(n)$. Our goal is to turn it into an algorithm \mathcal{A}' with round complexity $t'(n)$ that works if the identifiers are from $[f(n)]$. Think of \mathcal{A} as a sequence $\mathcal{A}_1, \mathcal{A}_2, \dots$ for each $n \in \mathbb{N}$. We define \mathcal{A}' by setting for each n that $\mathcal{A}'_n := \mathcal{A}_{f(n)}$. That is, we “lie” to the algorithm \mathcal{A} , telling it that the size of the graph is larger (namely $f(n)$) than what it actually is (namely n). Since the round complexity of \mathcal{A} on $f(n)$ -sized instances is $t(f(n))$, for the complexity of \mathcal{A}' on n -sized instances we have $t'(n) = t(f(n))$.

Moreover, since $\mathcal{A}_{f(n)}$ assumes that the unique identifiers are from $[f(n)]$, \mathcal{A}'_n also assumes that the unique identifiers from $[f(n)]$, making it a well-defined algorithm for the definition of local algorithm where identifiers are supposed to be from $[f(n)]$ on instances of size n .

Finally, we claim that \mathcal{A}' is a correct algorithm. To see this, suppose that \mathcal{A}'_n fails to solve Π on some graph G with n vertices labeled with unique identifiers from $[f(n)]$. Then, we go back to \mathcal{A} and run it on a graph G' defined as G together with $f(n) - n$ additional isolated vertices, all labeled with unique identifiers from $[f(n)]$. We notice that since Π is a local problem, a failure in G implies a failure in G' , and we get a contradiction with \mathcal{A} being correct. \square

As an example of a non-local problem where the range of identifiers matters, consider the leader-election problem¹⁴ where exactly one node of the input graph is to be selected. If the identifiers are from $[n]$, we can simply select the node with identifier 1. Otherwise, there is no local algorithm for it if the input graph is empty.

A similar argument to the proof of [Theorem 2.20](#) can be used to prove another sleep-well-at-night result: any local algorithm \mathcal{A} with round complexity $t(n)$ solving some local problem can be turned into an algorithm \mathcal{A}' of round complexity $t'(n) \leq t(n)$ which is a non-decreasing function of n .

¹³Looking at the proof, it is hard to come up with a reasonable class of graphs where the theorem does *not* apply.

¹⁴This is a fundamental problem in the broader area of distributed computing. The maximal independent set problem can be seen as a local variant of this problem.

The randomized version of [Theorem 2.20](#) turns out to be particularly interesting. Chang, Kopelowitz, and Pettie [[CKP19](#)] used it to show that *any* randomized algorithm can be derandomized, if we appropriately slow down its complexity.¹⁵

Theorem 2.21 (Chang, Kopelowitz, and Pettie [[CKP19](#)]). *If a local problem Π has randomized round complexity $t(n)$, its deterministic round complexity is $t(2^{O(n^2)})$.*

The theorem holds for any class of graphs closed on adding isolated vertices.

Proof. Let \mathcal{A} be a randomized local algorithm with error $1/n^{O(1)}$ and round complexity $t(n)$ solving Π . We start by applying [Theorem 2.20](#) with $f(n) = 2^{O(n^2)}$ to slow down \mathcal{A} to complexity $t(2^{O(n^2)})$ while pushing the error probability down to $2^{-\Omega(n^2)}$. We still call this new algorithm \mathcal{A} .

The deterministic algorithm \mathcal{A}' for Π will work as follows. Let C be such that the unique identifiers are coming from $[n^C]$. We will use a certain function h (a hash function that we soon specify) that maps each identifier from $[n^C]$ to an infinite bit string. We define \mathcal{A}' as follows: It first uses h to map each input identifier to an infinite bit string, and next it simulates \mathcal{A} using each bit string as the string of random bits.

To construct the function h that makes \mathcal{A}' correct, we simply choose a random one and notice that on any concrete graph, \mathcal{A}' fails to solve Π with probability at most $2^{-\Omega(n^2)}$ with the randomness over the choice of h . This is because on any concrete graph, \mathcal{A}' with random h is equivalent to running \mathcal{A} . But notice that the number of distinct graphs on n vertices labeled with polynomial identifiers is $2^{O(n^2)} \cdot (n^C)^n = 2^{O(n^2)}$. Hence, choosing the failure probability small enough, we can union-bound over all distinct labeled graphs and still get that random h works, with positive probability. We thus conclude that for some h , \mathcal{A}' is a valid deterministic algorithm for Π , as needed. \square

As a nice corollary, we are getting a very precise understanding of what is happening for the randomized round complexity of Lovász local lemma: On one hand, we have seen a randomized algorithm with round complexity $\text{poly log log}(n)$ ([Theorem 2.12](#)) that used a deterministic algorithm with round complexity $\text{poly log}(n)$ ([Theorem 2.10](#)) as a subroutine. But we can now see that *any* $\text{poly log log}(n)$ -round randomized algorithm would imply a $\text{poly log}(n)$ -round deterministic one via [Theorem 2.21](#). A similar observation applies not just for the local lemma, but for many other problems with randomized algorithms based on shattering (see [Section 1.6](#)). This explains why in local complexity deterministic algorithms are a big deal – not only they are used as subroutines of randomized algorithms, but we in fact understand that to make progress in the randomized world, progress in the deterministic world is often necessary.

2.3.2 Speedups

While slowdowns are very useful, the real fun starts when we try to speed up algorithms by lying to them that the size of the graph is *smaller* than it actually is. In fact, what can possibly go wrong if we tell an algorithm that the size of the graph is constant, i.e., we claim that $n = O(1)$? As it turns out, not much! One potential problem is that the algorithm may find out that its $t(n)$ -hop neighborhood contains more than n nodes and “crash”. The other problem is that the algorithm now requires very small identifiers (if we started with a deterministic algorithm) or that the failure probability of the algorithm increased to constant (if we started with a randomized one). This is where our story connects with our discussions of coloring and Lovász local lemma.

Theorem 2.22 (Chang and Pettie [[CP19](#)] and Chang, Kopelowitz, and Pettie [[CKP19](#)]). *Let \mathcal{A} be a local algorithm for a local problem Π with round complexity $t(n) = o(\log n)$. Then there is a local algorithm \mathcal{A}' solving Π such that*

1. *If \mathcal{A} was deterministic, so is \mathcal{A}' and its round complexity is $O(\log^* n)$.*
2. *If \mathcal{A} was randomized, so is \mathcal{A}' and its round complexity is the same as the randomized complexity of solving relaxed Lovász local lemma, i.e., it is $\tilde{O}(\log^4 \log(n))$.*

¹⁵Their technique is similar to Adelman’s theorem that proves that $\text{BPP} \subseteq \text{P/poly}$. See [[AB09](#), Theorem 7.14].

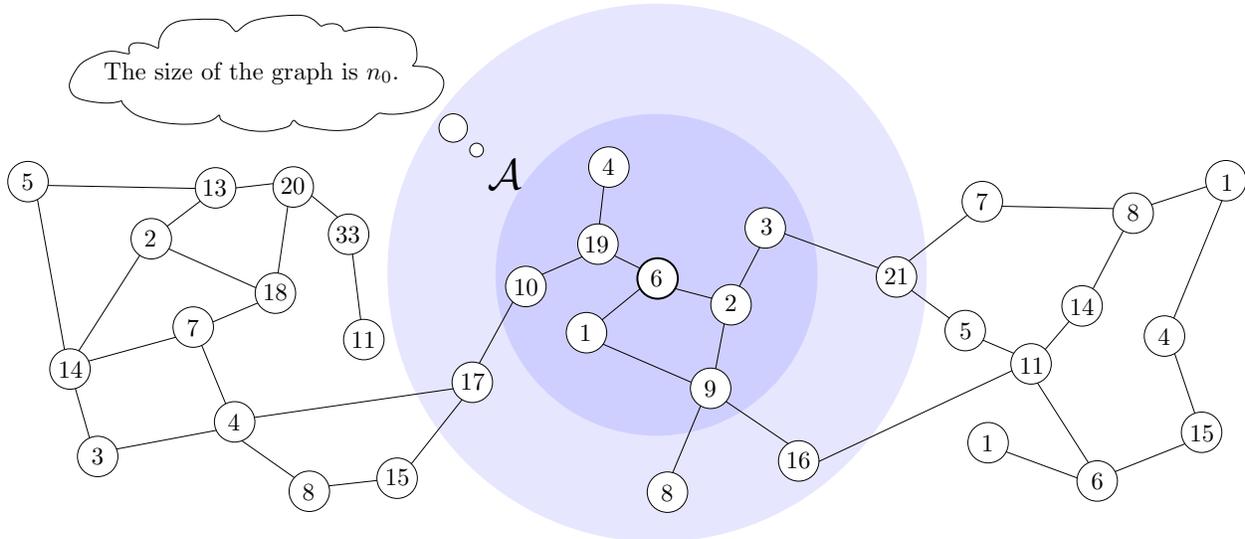


Figure 2.5: To speed up an algorithm \mathcal{A} , we tell it that the size of the graph is some constant n_0 , whatever the true size n is. In this case, the round complexity of the algorithm becomes 2 (see the smaller ball in the picture) and we assume that the checkability radius of the local problem solved by \mathcal{A} is 1.

The algorithm \mathcal{A} assumes that the identifiers are unique numbers of size up to $n_0^{O(1)}$. Since $n \gg n_0$, we cannot supply such identifiers. However, we notice that if \mathcal{A} is supplied with $n_0^{O(1)}$ -sized identifiers that are unique only in every ball of radius 3 (see the larger ball in the picture), the algorithm still has to work since. This is because a failure of \mathcal{A} around some node u implies a failure of \mathcal{A} on the graph $B(u, 3)$ which has size at most n_0 and is labeled with unique identifiers.

The theorem holds for the class of bounded degree graphs and any of its subclasses closed on taking subgraphs and adding isolated vertices. More generally, for round complexities defined relative to a subclass of bounded degree graphs of restricted growth, the complexity $t(n)$ needs to be such that we always have $|B(u, t(n))| = o(n)$ for any node u .

Proof. Let Π be a local problem with a checkability radius of r and \mathcal{A} be a deterministic (randomized) local algorithm for Π . We will assume that the identifiers are coming from a range $[n^C]$ (alternatively, the failure probability is $1/n^C$ for sufficiently large C). We will view \mathcal{A} as a sequence $\mathcal{A}_1, \mathcal{A}_2, \dots$, and we start by choosing n_0 to be a large enough constant so that

$$\Delta^0 + \Delta^1 + \dots + \Delta^{t(n_0)+r} \leq n_0. \quad (2.3)$$

We notice that such an n_0 exists by the requirement $t(n) = o(\log n)$. Our algorithm \mathcal{A}' will simulate \mathcal{A}_{n_0} for any input n (for $n < n_0$ we define $\mathcal{A}'_n = \mathcal{A}_n$).

Let us first handle the case of deterministic algorithms (see Figure 2.5). There, \mathcal{A}' needs to supply identifiers to its simulation of \mathcal{A}_{n_0} . This is done as follows. We compute a $(\Delta(G') + 1)$ -coloring of the power graph $G' = G^{2(t(n_0)+r)}$ such that the number of colors used is $\Delta(G') + 1 = \Delta^{O(t(n_0)+r)} \leq n_0^C$ where we used that C is large enough. We use those colors as identifiers for our simulation of \mathcal{A}_{n_0} . This finishes the description of \mathcal{A}' .

On one hand, the round complexity of the overall algorithm is $O(\log^* n)$ on bounded-degree graphs. On the other hand, suppose that \mathcal{A}' fails to solve the problem Π at some node u . We notice that Equation (2.3) yields that the $(t(n_0) + r)$ -hop neighborhood of u contains at most n_0 nodes. Moreover, this neighborhood is colored by \mathcal{A}' with unique colors from the range $[n_0^C]$. Recall that \mathcal{A}_{n_0} is defined on neighborhoods with at most n_0 nodes labeled with identifiers from $[n_0^C]$. Thus the algorithm \mathcal{A}' is well-defined. Additionally, a failure of \mathcal{A}' at u would imply a failure of \mathcal{A} on a neighborhood of size at most n_0 labeled with some unique identifiers from $[n_0^C]$. Adding isolated vertices to that neighborhood to make its size n_0 , we reach a contradiction with \mathcal{A}_{n_0} being correct on n_0 -sized instances.

We handle the case of randomized algorithms similarly: We again try to simulate \mathcal{A}_{n_0} for large enough

n_0 . We notice that we do not need to supply any identifiers in the simulation. However, if we simply simulate \mathcal{A} , we fail to solve Π at any fixed vertex u with probability $1/n_0^C$ as this is the failure probability of \mathcal{A}_{n_0} . Thus, we instead formulate the process of simulating \mathcal{A} as an instance of Lovász local lemma on a new variable-event graph G' as follows. Any original vertex u of the input graph G will be thought of as two new vertices in G' . One vertex, $u_{r.v.}$, is a random-variable vertex, and it corresponds to the random string at u . The other vertex, $u_{b.e.}$, corresponds to the event that if we run \mathcal{A}_{n_0} at u and use the random strings sampled from the random-variable vertices $v_{r.v.} \in B(u_{b.e.}, t(n_0))$, \mathcal{A}_{n_0} fails to solve Π at u .

We notice that the event at $u_{b.e.}$ depends only on the $(t(n_0) + r)$ -hop neighborhood of u in G . Thus the degree of this local-lemma instance (in the sense of its dependency-graph formulation) is at most $\Delta^0 + \Delta^1 + \dots + \Delta^{t(n_0)+r}$ which is in turn at most n_0 by Equation (2.3). On the other hand, the probability of each bad event is at most $1/n_0^C$. We can thus formulate the simulation of \mathcal{A}_{n_0} as an instance of a C -relaxed local lemma. For large enough C , we can then apply Theorem 2.12 to solve that instance.

Finally, we notice that both in the deterministic and the randomize case, we can generalize the requirement $t(n) = o(\log n)$ to requiring that t satisfies for each u that $|B(u, t(n))| = o(n)$. \square

We note that in the statement of Theorem 2.22, we do not literally require $t(n) = o(\log n)$ (or $|B(u, t(n))| = o(n)$), it should just be that for some large enough n_0 , we have $t(n_0) \leq \varepsilon \log n_0$, where ε is some small constant which is a function of Δ, r , and C .

Additional intuition yielded by the proof: The proof of Theorem 2.22 tells us a bit more: Basically, it allows us to think of local problems with deterministic round complexity $o(\log n)$ or, equivalently, $O(\log^* n)$, as “those problems that can be solved in a constant number of rounds, after we compute a distance coloring”. Moreover, the problems with randomized round complexity $o(\log n)$ or, equivalently poly $\log \log n$, can be thought of as “those problems that we can view as an instance of the local lemma”.

*Problems with deterministic round complexity $o(\log n)$ can be, in fact, solved in constantly many rounds, after a suitable coloring used as “fake identifiers” is computed.
Problems with randomized round complexity $o(\log n)$ can be formulated as instances of polynomially relaxed Lovász local lemma.*

Speedup below $\log^* n$: We will next sketch how one can speed up algorithms with round complexity $o(\log \log^* n)$ to complexity $O(1)$. Why the weird complexity $o(\log \log^* n)$? The following speedup argument relies more on the volume than on the radius¹⁶, so the importance of radius $o(\log \log^* n)$ is that the number of nodes the algorithm sees is $\Delta^{o(\log \log^* n)} = o(\log^* n)$. The $\log^* n$ volume is then the right threshold for which we can argue that any deterministic algorithm can be turned into a drastically simpler *order-invariant* algorithm.

An order-invariant algorithm is a deterministic local algorithm with an additional restriction: it does not have direct access to the identifiers written on the nodes; it only knows their *order*, meaning that it can only compare their relative size. For example, in the setup of Theorem 2.8 where we worked with an oriented path with increasing labels, an order-invariant algorithm would see the same input order at every vertex. This already shows that order-invariant algorithms struggle to solve interesting local problems on graphs that are paths.

Theorem 2.23 (Naor and Stockmeyer [NS95] and Chang, Kopelowitz, and Pettie [CKP19]). *Let \mathcal{A} be a deterministic local algorithm for a local problem Π with round complexity $o(\log \log^* n)$. Then, there is a deterministic local algorithm \mathcal{A}' solving Π with local complexity $O(1)$.*

The theorem holds for any subclass of bounded-degree graphs closed on taking subgraphs and adding isolated nodes.

Proof Sketch. We will sketch how any deterministic local algorithm \mathcal{A} with round complexity $t(n) = o(\log \log^* n)$ can be turned into an order-invariant algorithm \mathcal{A}' , using the hypergraph Ramsey theorem¹⁷. Consider the set $B = [n^{O(1)}]$ of identifiers and any of its subsets S of size $(2\Delta)^{t(n)} = o(\log^* n)$. Here, the function $(2\Delta)^{t(n)}$

¹⁶Compare with Theorem 3.2 classifying volume complexities where we have $o(\log^* n)$ speedup instead.

¹⁷Recall that in hypergraph Ramsey theorem we have a big set B , we color its small subsets S and we want to find a medium-sized set M such that all its small subsets $S \subseteq M$ have the same color.

upper bounds the maximum number of nodes in any $t(n)$ -hop neighborhood. We say that the *type* of S is the list of all possible $t(n)$ -hop neighborhoods, together with their labeling with identifiers from S , where for each labeled graph on this list we record the output of \mathcal{A} when being run on that labeled neighborhood.

We note that the number of types is relatively small, only roughly $2^{2^{O(\log^* n)^2}}$. For r the checkability radius of Π , we want to find a medium-sized set $M \subseteq B$ of size $(2\Delta)^{t(n)+r} = o(\log^* n)$ such that all subsets $S \subseteq M$ have the same type (i.e., the same color in the terminology of Ramsey's theorem). The known bounds for hypergraph Ramsey numbers [GRS91, §1, Theorem 2] allow us to conclude that B is large enough to always contain such a set M .

With such a set M at hand, we next define the order-invariant algorithm \mathcal{A}' for Π as follows: given an input order on vertices, the algorithm chooses an arbitrary set $S \subseteq M$ and outputs the answer that \mathcal{A} would give if the vertices were labeled with identifiers from S in the following way: The order induced by the identifiers from S has to be the same as the actual input order. It does not matter which $S \subseteq M$ we choose since they all have the same type. This finishes the description of \mathcal{A}' .

To see that \mathcal{A}' is valid, we assume for contradiction that it fails for some input order on vertices at some vertex u . Consider that input order on vertices in the $(t(n) + r)$ -hop neighborhood of u . Replace this input order with identifiers from M such that they induce the same order. The original algorithm \mathcal{A} behaves the same as the new algorithm \mathcal{A}' when it sees only the identifiers from M : hence, \mathcal{A} also fails. After adding isolated nodes to this neighborhood, we reach a contradiction.

Finally, let us contemplate the proof of [Theorem 2.22](#) again. The only reason why we could speedup $o(\log n)$ -round deterministic algorithm only to $O(\log^* n)$ instead of $O(1)$ was the necessity to compute a coloring that served as "fake" identifiers supplied to the simulation of an algorithm \mathcal{A}_{n_0} . But for order-invariant algorithms, this is not necessary. We can simply give \mathcal{A}_{n_0} exactly the same order as the order that appears on input; thus order-invariant algorithms can be sped up to complexity $O(1)$, as needed. \square

Tower-sized identifiers: We can wonder what would happen if we wanted to make the above proof of [Theorem 2.23](#) work for *all* deterministic local algorithms, not just those that encounter only $o(\log^* n)$ vertices. This would force us to change the small-set size from $o(\log^* n)$ to n in the proof. Since the hypergraph Ramsey numbers grow roughly as a tower function of the small-set size, this means that to make the proof work, we would have to start with unique identifiers from the range which contains numbers as large as the tower function of n . The rest of the proof then goes through and allows us to speed up $o(\log n)$ -round algorithms to $O(1)$ -round order-invariant algorithms.

This observation nicely complements [Theorem 2.7](#): While it is true that exponentially-, doubly-exponentially-, etc. sized identifiers have the same power as polynomially-sized ones, there is a transition somewhere around tower-function-sized identifiers. From then on, deterministic algorithms of $o(\log n)$ complexity with such huge identifiers can be turned into order-invariant ones. In other words, the identifiers are so large that no $o(\log n)$ -round algorithm can extract any meaningful information from them other than their relative order.

2.4 Classification of Local Problems

We can now put all the pieces together and prove [Theorem 2.1](#). Let's restate it here for convenience.

Theorem 2.1 (Classification of local problems with $o(\log n)$ complexity on bounded degree graphs). *Let us fix any Δ and the class of graphs of degree at most Δ . Then, any local problem with randomized round complexity $o(\log n)$ has one of the following three round complexities.*

1. Order-invariant regime: *The problem has $O(1)$ deterministic and randomized round complexity.*
2. Symmetry-breaking regime: *The deterministic and randomized round complexity of the problem lies between $\Omega(\log \log^* n)$ and $O(\log^* n)$ (both the deterministic and the randomized complexity is the same function).*
3. Lovász-local-lemma regime: *The problem has deterministic round complexity between $\Omega(\log n)$ and $\tilde{O}(\log^4 n)$. Its randomized round complexity is between $\Omega(\log \log n)$ and $\tilde{O}(\log^4 \log n)$.*

Proof. Let Π be any local problem of randomized local complexity $t(n) = o(\log n)$. We can use [Theorem 2.22](#) to formulate Π as an instance of Lovász local lemma. This instance is then solved either with the deterministic $\tilde{O}(\log^4 n)$ -round algorithm from [Theorem 2.10](#), or the randomized $\tilde{O}(\log^4 \log n)$ -round algorithm from [Theorem 2.12](#).

Next, let us assume that the randomized round complexity of Π is even $o(\log \log n)$. Then, we can use the derandomization of [Theorem 2.21](#) to conclude that the deterministic round complexity is $o(\log n)$. This allows us to use the speedup theorem of [Theorem 2.23](#) to conclude that the deterministic round complexity of Π is $O(\log^* n)$. Moreover, we notice that for any increasing function $t(n) = O(\log^* n)$, we have $t(2^{O(n^2)}) = O(t(n))$; this means that the derandomization from [Theorem 2.21](#) tells us that randomized and deterministic round complexities of Π are the same.

Finally, assume that the randomized local complexity of Π is even $o(\log \log^* n)$. Then, we can use [Theorem 2.23](#) to find an order-invariant constant-round local algorithm for Π via [Theorem 2.23](#). \square

Local problems with complexities below $o(\log n)$ are of three types:

1. *Those that can be solved in constant rounds by order-invariant algorithms,*
2. *those that are roughly as hard as the basic symmetry-breaking problems such as coloring or maximal independent set,*
3. *those that can be viewed as instances of Lovász local lemma.*

Improvements to [Theorem 2.1](#): Can we improve [Theorem 2.1](#)? The conjecture of Chang and Pettie [[CP19](#)] ([Problem 2.14](#)) suggests that in the local lemma regime, we may get rid of the polynomial gap between the lower and upper bounds. According to our current knowledge, there could be a local problem with deterministic and randomized round complexities $t_d(n), t_r(n)$ for any $t_d(n) = \Omega(\log n)$, $t_d(n) = \tilde{O}(\log^4 n)$, $t_r(n) = \Omega(\log \log n)$, $t_r(n) = \tilde{O}(\log^4 \log n)$, and $t_d(n) = O(t_r(2^n))$, where the last constraint follows from [Theorem 2.20](#).

On the other hand, it is known [[Bal+18](#)] that the (extremely narrow) regime between $\Omega(\log \log^* n)$ and $O(\log^* n)$ is densely populated by certain (artificial) local problems.

2.4.1 Sequential local complexities

Let us now discuss how sequential local complexities fit into the picture painted by [Theorem 2.1](#). We will first observe that the first two classes from our classification in [Theorem 2.1](#), i.e., $O(1)$ and $O(\log^* n)$ round complexities, are equal to the class of problems solvable with the sequential local complexity is $O(1)$.

Theorem 2.24. *On bounded degree graphs, local problems with round complexity $O(\log^* n)$ are exactly those whose sequential local complexity is $O(1)$.*

Proof sketch. On one hand, consider any local problem with constant sequential local complexity. We can use [Theorem 2.5](#) to simulate it with a distributed local algorithm in $O(\log^* n)$ rounds.

On the other hand, consider any local problem Π with local checkability r solvable in $O(\log^* n)$ round complexity. Recall that the proof of [Theorem 2.22](#) says that Π can in fact be solved with a local algorithm \mathcal{A} of round complexity $t = O(1)$ on any graph G if we are provided a proper coloring of a suitable power graph $G^{2(t+r)}$. Recall that $(\Delta + 1)$ -coloring has sequential local complexity $O(1)$, meaning that we can construct a sequential local algorithm of constant complexity for Π . \square

Next, we will show how our distributed speedup theorems imply speedups in the sequential world.

Theorem 2.25. *Let \mathcal{A} be a sequential local algorithm solving a local problem Π on bounded degree graphs. Assume that either*

1. *\mathcal{A} is deterministic and its sequential local complexity is $o(\log \log n)$,*
2. *or \mathcal{A} is randomized and its sequential local complexity is $o(\log \log \log n)$.*

Then, there is a distributed local algorithm \mathcal{A}' that solves Π in $O(\log^ n)$ round complexity.*

Proof. If \mathcal{A} is a deterministic sequential algorithm of local complexity $t(n) = o(\log \log n)$, we can simulate it with a distributed local algorithm via [Theorem 2.5](#). We get a deterministic algorithm with local complexity $\Delta^{O(t(n))} + O(t(n) \log^* n) = \log^{o(1)} n = o(\log n)$. Such an algorithm is then sped up to $O(\log^* n)$ complexity using [Theorem 2.22](#).

Similarly, if \mathcal{A} is randomized sequential algorithm of local complexity $o(\log \log \log n)$, [Theorem 2.5](#) implies that we get a randomized distributed local algorithm with $\log^{o(1)} \log n = o(\log \log n)$ round complexity. Such an algorithm can then be derandomized by [Theorem 2.21](#) into a deterministic algorithm of local complexity $o(\log n)$ which is in turn sped up to $O(\log^* n)$ using [Theorem 2.22](#). \square

Finally, we sketch how one can construct very fast sequential local algorithms for instances of the local lemma.

Theorem 2.26. *Let Π be a C -relaxed Lovász local lemma for sufficiently large C . Then, its deterministic sequential local complexity is $\tilde{O}(\log^4 \log n)$ and its randomized sequential local complexity is $\tilde{O}(\log^4 \log \log n)$ ¹⁸, on bounded degree graphs.*

Proof Sketch. To prove the first part of the theorem, we consider the randomized distributed local algorithm for local lemma from [Theorem 2.12](#). We turn it into a deterministic sequential local algorithm of the same complexity by the derandomization result of [Theorem 1.7](#).

To prove the second part of the theorem, we notice that our randomized distributed local algorithm from [Theorem 2.12](#) can certainly be implemented in the model of randomized sequential local algorithms. Recall that the algorithm reduces an instance of the local lemma to instances on graphs of size $O(\log n)$, as proven in [Theorem 2.11](#). Then, the best deterministic algorithm is applied to those instances to finish the job. While [Theorem 2.11](#) is stated as reducing from C -relaxed local lemma with some large C to local lemma with condition $p \leq 1/(3\Delta)$, we can generalize it to show that starting with a very large C' , we can reduce an instance of C' -relaxed local lemma to C -relaxed local lemma for C that is still arbitrarily large.

This allows us to run the deterministic sequential algorithm with complexity $\tilde{O}(\log^4 \log n)$ constructed in the first part of the proof instead of using the deterministic $O(\text{poly } \log n)$ -round algorithm from [Theorem 2.10](#). The final randomized sequential local complexity is thus $\tilde{O}(\log^4 \log \log n)$. \square

We note that for problems like sinkless orientation for which an $O(\log \log n)$ -round randomized algorithm is known [[GHK18](#)], the proofs of above [Theorems 2.25](#) and [2.26](#) imply tight bounds of $\Theta(\log \log n)$ and $\Theta(\log \log \log n)$ deterministic and randomized sequential local complexity on bounded degree graphs. In the case of sinkless orientation, the same bounds are also known even on unbounded degree graphs [[GHK18](#)]. If the conjecture of Chang and Pettie [[CP19](#)] ([Problem 2.14](#)) is true, then there is no gap between [Theorems 2.25](#) and [2.26](#).

We also note that [Theorems 2.25](#) and [2.26](#) show that while one can turn randomized distributed local algorithms into sequential deterministic ones via [Theorem 1.7](#), it is *not* possible to turn randomized sequential local algorithms into deterministic sequential local ones, without loss in their local complexity.

Putting everything together, we get the following refined classification theorem that also includes sequential local complexities.

Theorem 2.27 (Refined classification of local problems). *Let us fix any Δ and any class of graphs \mathcal{G} of degree at most Δ that is closed under taking subgraphs and adding isolated nodes. Let $t(n)$ be a function such that for any node u in some graph in \mathcal{G} we have $|B(u, t(n))| = o(n)$.*

Then, any local problem with randomized round complexity at most $t(n)$ has one of the following three complexities.

1. Order-invariant regime: *The problem has $O(1)$ deterministic (or randomized) round complexity and sequential local complexity.*
2. Symmetry-breaking regime: *The deterministic and randomized round complexity of the problem lies between $\Omega(\log \log^* n)$ and $O(\log^* n)$ (both the deterministic and the randomized complexity is the same function). Moreover, the sequential local complexity is $O(1)$.*

¹⁸It is an interesting exercise to open up the proof that the randomized sequential local complexity of local lemma is $\Theta(\text{poly } \log \log \log n)$ and try to think of theorems we have seen so far that do *not* go into that proof.

3. Lovász-local-lemma regime: *The problem has the following complexities:*

- (a) *deterministic round complexity is between $\Omega(\log n)$ and $\tilde{O}(\log^4 n)$,*
- (b) *randomized round complexity and deterministic sequential local complexity is between $\Omega(\log \log n)$ and $\tilde{O}(\log^4 \log n)$,*
- (c) *randomized sequential local complexity is between $\Omega(\log \log \log n)$ and $\tilde{O}(\log^4 \log \log n)$.*

Moreover, if any graph from the class satisfies that for any node u , we have $|B(u, r)| = 2^{O(r^{0.249})}$, then there are no local problems in the third class of problems.

The last part of the theorem follows by applying [Theorem 2.22](#): We know that every local problem in the local-lemma regime can be solved with a deterministic local algorithm of round complexity $\tilde{O}(\log^4 n)$; since $2^{(\tilde{O}(\log^4 n))^{0.249}} = o(n)$, the speedup to $O(\log^* n)$ complexity can be applied. If the conjecture of Chang and Pettie [[CP19](#)] ([Problem 2.14](#)) is true, we can get rid of the polynomial slack in the local-lemma regime complexities and the whole local-lemma class of problems is then empty for the class of subexponential growth graphs (see [Section 2.5](#)).

Problem 2.28. *Can local problem solvable with deterministic sequential local complexity $o(\log n)$ be sped up to $\text{poly} \log \log n$ deterministic sequential complexity (i.e., to the local-lemma regime)? What about problems with randomized sequential local complexity $o(\log n)$?*

2.5 Classification of Local Problems for Concrete Graph Classes

This section compiles results from a long line of work trying to extend the classification of local problems on bounded-degree graphs, i.e., [Theorem 2.1](#), to more specific graph classes. For some simple graph classes, we now have an almost complete understanding of the local complexity classes that go even beyond the $\log n$ local complexity threshold.

Locally checkable labeling problems: The theorems below are proven for the so-called *locally checkable labeling problems*. These are local problems that additionally allow one of the finitely many input colors on each node of the input graph. For example, a list coloring is an example of a locally checkable labeling problem.

Definition 2.29. *A locally checkable labeling problem (LCL) Π is formally a quintuple $(S_{in}, S_{out}, r, \Delta, \mathcal{P})$. Here, S_{in} is the finite set of input labels, S_{out} is the finite set of output labels, and \mathcal{P} is the set of allowed neighborhoods: Each neighborhood has maximum degree Δ , radius r , and each node is labeled with a label from S_{in} and S_{out} . Given an input graph of degree at most Δ and any input coloring of its nodes with S_{in} , the task is to find an output coloring by colors from S_{out} so that the r -hop neighborhood of each node is in \mathcal{P} .*

We did not have to distinguish between our definition of a local problem from [Definition 1.1](#) and the locally checkable labelings since, in the general graph classes like the class of all graphs or all bounded-degree graphs, it is straightforward to encode input labels as a part of the input graph. We thus chose the simpler definition that also allowed us to talk about local problems even outside of the setup of bounded-degree graphs. In the following theorems, especially for very restrictive classes like paths or grids, the difference between the two definitions matters, since the possibility of allowing inputs may substantially enlarge the number of possible local problems.

In the following theorems, we will always implicitly assume that they are for *solvable* locally checkable labelings, i.e., for problems where the solution is always guaranteed to exist. Otherwise, the problem does not have a well-defined local complexity class.

Computational aspect: We note that an additional dimension for all classification results below is the computational complexity of the classification problem, where the input is a description of locally checkable labeling, and the output is which class it belongs to. On the one hand, it is known that the classification problem is decidable on paths, albeit PSPACE-hard [[Bal+19a](#)]. On the other hand, the classification problem

is undecidable even for grids [Bra+17]. The decidability on trees is open; see [Problem 2.34](#). We will not discuss this dimension in the following theorem statements for brevity.

Additional remarks: We list the concrete classification theorems below. Unless stated otherwise, the deterministic and the randomized round complexity of the problem is always the same. We are always citing papers proving results specific to the given graph class, for example, general speedup and slowdown theorems of Chang and Pettie [CP19] and Chang, Kopelowitz, and Pettie [CKP19] are as a rule of thumb always an important part of the proof. Also, not all the papers we cite prove a part of the given theorem. Some of them only provide building blocks, analyze the computational complexity of the classification, etc.

List of known classifications: Let us now list known classification theorems on concrete graph classes.

Theorem 2.30 (Classification of local problems on paths [Bra+17; Bal+19a]). *Any solvable locally checkable labeling problem on oriented or unoriented paths with inputs has one of the following round complexities:*

1. $O(1)$,
2. $\Theta(\log^* n)$,
3. $\Theta(n)$.

Theorem 2.31 (Classification of local problems on grids [NS95; Bra+17; CKP19; GRB22]). *Any solvable locally checkable labeling problem on d -dimensional oriented grids¹⁹ has one of the following local complexities:*

1. $O(1)$,
2. $\Theta(\log^* n)$,
3. $\Theta(n^{1/d})$.

We note that on unoriented grids, the full classification is not known [GRB22].

Theorem 2.32 (Classification of local problems on bounded-degree rooted regular trees [Bal+21b; Bal+22a]). *Any solvable locally checkable labeling problem on bounded-degree rooted regular trees²⁰ has one of the following round complexities:*

1. $O(1)$,
2. $\Theta(\log^* n)$,
3. $\Theta(\log n)$,
4. $\Theta(n^{1/k})$ for some $k \in \mathbb{N}$.

Theorem 2.33 (Classification of local problems on bounded-degree trees [MR89; CP19; Bal+18; Cha20; Bal+20b; GRB22; Bal+22a; Bal+20a]). *Any solvable locally checkable labeling problem on trees has one of the following local complexities:*

1. $O(1)$,
2. $\Theta(\log^* n)$,
3. $\Theta(\log \log n)$ randomized, $\Theta(\log n)$ deterministic,
4. $\Theta(\log n)$ (both randomized and deterministic),
5. $\Theta(n^{1/k})$ for some $k \in \mathbb{N}$.

Problem 2.34. *Given an input locally checkable labeling problem on bounded-degree trees, is it decidable which class in [Theorem 2.33](#) it belongs to?*

¹⁹The edges of the input grid are consistently oriented.

²⁰Edges are oriented such that every vertex has at most one ingoing edge.

More Classifications: It is an intriguing question of how far these classifications can be extended, see e.g. the recent work extending the theory to minor-closed classes [Cha24] or extending the theory to unbounded-degree graphs [LPS23].

Another interesting class is the class of subexponential growth graphs, i.e., the class of bounded degree graphs where for every $\varepsilon > 0$ we can find r such that $|B(u, r)| \leq (1 + \varepsilon)^r$ ²¹. In this class of graphs, we can apply [Theorem 2.22](#) to speed up deterministic algorithms of complexity $O(\log n)$ (instead of $o(\log n)$) to $O(\log^* n)$. If the conjecture of [CP19] ([Problem 2.14](#)) is true, this has an interesting corollary: the local lemma regime from the classification theorem of [Theorem 2.1](#) would then not be present in this class of graphs and there would only two classes of local problems there (see [Theorem 2.27](#)). This leads to the following special case of [Problem 2.14](#):

Problem 2.35. *Is there a C such that all instances of the C -relaxed Lovász local lemma can be solved with round complexity $\Theta(\log^* n)$ on any class of graphs of subexponential growth?*

Some implications of a positive answer to this problem are known [Csó+22; BD23a].

²¹Formally, this is a family of classes, with one class for each possible dependency $r = r(\varepsilon)$.

Chapter 3

Applications

While the previous two chapters, [Chapters 1](#) and [2](#), discussed the general theory of local algorithms, the following chapter considers various popular models of distributed, parallel, or sublinear algorithms, as well as the field of descriptive combinatorics, and briefly discusses how techniques from local algorithms can help there, discussing a few cherry-picked examples per section. These sections are heavily influenced by the author’s particular interests; a different author would choose different applications.

3.1 Distributed Computing (CONGEST)

The CONGEST model [[Pel00](#)] is a model of distributed computing that is extremely tightly connected to local algorithms¹ Starting from the definition of a local algorithm as a message-passing protocol between nodes in a graph, the CONGEST model adds an additional requirement that each message is supposed to be *small*. In particular, it should fit into $O(\log n)$ bits. For example, a node can send its unique identifier in one round.

Many local algorithms in the literature are in fact stated as CONGEST algorithms since a number of local algorithms are readily implemented in the CONGEST model. However, not all of them. For example, the problem of counting for each node the number of triangles that contain it is a trivial problem for local algorithms but a very complex problem in the CONGEST model, where it requires polynomially many rounds [[CPZ19](#); [CS19](#)]. The general derandomization technique of Ghahramani, Harris, and Kuhn [[GHK18](#)] from [Theorem 1.8](#) is also not directly applicable to the CONGEST model, since computing conditional expectations in a neighborhood of a node requires collecting a lot of information from that neighborhood. However, the general approach is still helpful, and it was used to derandomize concrete problems like maximal independent set [[CPS17a](#)] or $\Delta + 1$ -coloring [[BKM20](#)].

It is an interesting question for which classes of graphs the definitions of local and CONGEST algorithms coincide for local problems. It is known that this is the case for local problems on trees [[Bal+21c](#)].

3.2 Local Computation Algorithms and the Volume Model

The model of local computation algorithms [[Rub+11](#); [Alo+12](#)], often denoted as LCA, is a model of sublinear algorithms closely related to the local model. We will first explain its variant, known as the *Volume* model by Rosenbaum and Suomela [[RS20](#)]. This model is very similar to local algorithms, but we measure the *volume* instead of the *radius*. A volume algorithm run at a node u starts with a set $S_0 = \{u\}$ and in the i -th step, the algorithm can pick a node $u_i \in S_i$ and an arbitrary index $1 \leq j \leq d$ where d is the degree of u_i . Then, it learns the identifier of the j -th neighbor v of u_i . We then set $S_{i+1} = S_i \cup \{v\}$. The volume complexity of the algorithm is the number of queries it makes until it decides to finish and output the label for u .

¹Local algorithms are often referred to as “distributed algorithms in the LOCAL model of distributed computing”.

In particular, we note that for graphs of maximum degree Δ , any $t(n)$ -round local algorithm can be turned into a volume algorithm of complexity $(2\Delta)^{t(n)}$ as the volume algorithm at u can simply query all the nodes in the $t(n)$ -hop neighborhood of u [PR07].

A local computational algorithm may have several definitions. The stateless local computation algorithm is defined as follows: we start with the definition of the volume model and add the additional requirement that the nodes start with unique identifiers from the range $[n]$, even for randomized algorithms. In each query, the algorithm can additionally ask for the node with the identifier j for any $1 \leq j \leq n$. The node u_i with identifier j is then added to the set of already seen nodes S_i . That is, the algorithm can use additional global queries that are, however, “blind”.

It is unclear whether global queries can help with solutions to local problems. We know that they do not help for very fast randomized local computation algorithms [Gö+16].

Problem 3.1. *Is there a local problem that distinguishes (either randomized or deterministic) Local computation algorithms and Volume models with respect to a natural class of graphs?*

As an example problem considered in the LCA model, there is a long line of work on local computation algorithms for maximal independent set and related problems [NO08; YY109; Ona+12; Beh22; Rub+11; Alo+12; LRY15; Gha16; RV16; GU19; Gha22] culminating with the randomized algorithm of Gha ari [Gha22] with volume complexity $\text{poly}(\Delta \log n)$.

For constant degree graphs, we can try to prove classification results similar to the classification of local problem complexities from Theorem 2.1. We have some understanding of the volume complexities of the three important local complexity classes from Theorem 2.1:

Theorem 3.2 ([PR07; EMR14; RS20; GRB22; BGR21]). *In the class of bounded-degree graphs, each local problem satisfies the following:*

1. *If its round complexity is $O(1)$ (order-invariant regime), then its randomized/deterministic volume complexity is $O(1)$.*
2. *If its randomized/deterministic round complexity is $\omega(1)$ and $O(\log^* n)$ (symmetry-breaking regime), its randomized/deterministic volume complexity is $\Theta(\log^* n)$.*
3. *If its randomized local complexity is $\omega(\log^* n)$ but $o(\log n)$ (Lovász local lemma regime), its randomized volume complexity is between $\Omega(\sqrt{\log n})$ and $O(\log n)$.*

We note that as far as we know, all local problems in the local-lemma regime have randomized volume complexity of $\Theta(\log n)$. In particular, sinkless orientation has this randomized volume complexity [BGR21].

Problem 3.3 (See Conjecture 1.3 in [BGR21]). *Is it true on bounded degree graphs that local problems from the local-lemma regime have randomized volume complexity $\Theta(\log n)$?*

We also note that in the model of deterministic volume complexity where the identifiers come from exponential, instead of polynomial range, we can prove that only the volume complexities $O(1)$, $\Theta(\log^* n)$, $\Theta(n)$ are possible [BGR21]. The same could be the case also for the standard, polynomial-range, identifiers.

Problem 3.4. *Are there any local problems such that their deterministic volume complexity (with polynomial-sized identifiers) on bounded degree graphs is $\omega(\log^* n)$ but $o(n)$?*

Equivalently, we could have asked whether there is a local problem whose deterministic volume complexity is different for polynomial-sized and exponential-sized identifiers.

3.3 PRAM

PRAM is a classical model of parallel algorithms studied extensively in the past 40 years [JáJ92]. It simplifies the complexity of practical parallel computing by assuming a simple model of a machine with multiple processors sharing a common memory. There are two complexity measures: *work*, i.e., the total number of instructions made by all processors together throughout the execution, and *depth*, i.e., the number of rounds

necessary to finish the computation in the case that the machine is equipped with as many processors as the algorithm requires.

Many techniques developed for distributed and local algorithms are useful in the design of parallel algorithms. As an example application for a local problem, Ghaheri and Haeupler [GH21] develop a randomized algorithm for maximal independent set on the so-called EREW variant of the PRAM model with the optimal depth $O(\log n)$ and $O(m \log^2 n)$ work, building on top of the local maximal independent set algorithm of Ghaheri [Gha16].

As another example application for a non-local problem, we note that one can generalize deterministic local algorithms for network decompositions discussed in Section 1.5 to get various clustering results for weighted undirected graphs like the following one.

Theorem 3.5 (Deterministic low-diameter clustering [Roz+22a]). *Let G be a graph and w its nonnegative weights. Then, there is a parallel algorithm with $\tilde{O}(m+n)$ work and $\tilde{O}(1)$ depth such that, given a parameter $R > 0$, it splits vertices of G into clusters such that*

1. *Each cluster has weighted diameter $\tilde{O}(R)$,*
2. *The total weight of edges that cross between different clusters is at most $\frac{1}{R} \sum_{e \in E(G)} w(e)$.*

Variants of this clustering result can be used for the design of deterministic parallel algorithms for the approximate shortest path problem or various metric embedding problems [Roz+22b; Roz+22a].

3.4 Massively Parallel Computing (MPC)

The final parallel model we discuss is the model of Massively parallel computing (MPC) [KSV10], the theoretical model behind the popular programming framework MapReduce and its variants. While the previous PRAM model focuses on the total work done by processors and ignores the cost of communication, the MPC model ignores the work done by processors and focuses on the communication cost.

Concretely, let us explain the so-called low-memory regime of MPC: At the beginning of any graph algorithm, the edges of the input graph are split into many machines, each capable of storing $O(n^\epsilon)$ bits in its memory. In one round, each machine can perform an arbitrary computation on its edges, and then it sends arbitrary information to any of its neighbors (the total information sent and received per machine is $O(n^\epsilon)$). We measure the number of rounds until the solution is computed.

Chang, Fischer, Ghaheri, Uitto, and Zheng [Cha+19a] constructed a randomized massively parallel algorithm for $\Delta + 1$ -coloring that works in $O(\log \log \log n)$ rounds by simulating the fastest randomized local algorithm for coloring of round complexity $\text{poly} \log \log n$ (see Chapter 3) with exponential speedup. This result is complemented by the work of Ghaheri, Kuhn, and Uitto [GKU19] that showed that this result cannot be improved, conditioned on a certain conjecture. Notice how understanding where the three logarithms are coming from requires an extensive understanding of coloring with local algorithms: The way we arrive at this complexity is that, first, starting with the trivial constant-round sequential local algorithm for coloring, we turn it into a deterministic distributed $\text{poly} \log n$ -round local algorithm using the general translation of Theorem 1.4. Next, this algorithm can be turned into an exponentially faster randomized algorithm using the shattering technique [CLP18]. Finally, this randomized algorithm is simulated, again with exponential speedup, in the massively parallel model.

There is also a line of work aiming to extend the classification of local problems to the massively parallel algorithms model, though the complexity landscape there is currently much less understood [Bal+22b].

3.5 Descriptive Combinatorics

Descriptive combinatorics [Mik90; MU17; LAO17; DF92; MU16; Gab00; ASS99; Mar16; CM17; CGP24; Csó+22; Ber23a] is an area at the intersection of combinatorics, measure theory, and set theory (see e.g. the surveys [KM16; Pik21; Ber22b]). It studies graphs that arise when one manipulates uncountably-infinitely large mathematical objects equipped with measure. The main object of interest is a *measurable graph*. Instead of a formal definition, let us discuss a particular example.

An example of a measurable graph: Consider a circle, i.e., a set of points of distance 1 from the origin in the plane \mathbb{R}^2 . A rotation of this circle by 1 radian counterclockwise induces a graph G_0 : We can draw an oriented edge from each point v on the circle to the point v' such that the rotation maps v to v' . This graph has uncountably many connected components, each of which is a doubly-infinite oriented path: this is because 1 is an irrational multiple of 2π and thus if we start “jumping” from v around the cycle with jumps of length 1, we never return to v . Moreover, the standard Lebesgue measure on the circle is telling us which subsets of vertices of G_0 we are “allowed to talk about”. The graph G_0 , together with the measure on top of it, is an example of a measurable graph.

Given a local problem Π , we can now ask whether it can be solved on a particular measurable graph. For example, consider the following problem: Color each vertex of G_0 with one of c colors such that the coloring is proper and each set of vertices of the same color is measurable. It is a folklore fact that this type of coloring is not possible with $c = 2$ colors, but it is possible with $c = 3$ colors (this is a special case of the following [Theorem 3.6](#)).

Formal connection between the two worlds: There is a close connection between local algorithms and descriptive combinatorics. Although this was to some extent understood earlier (see Elek and Lippner [\[EL10\]](#), Lovász [\[Lov12, Chapter 23.3\]](#)), it was only a breakthrough paper of Bernshteyn [\[Ber23a\]](#) that first realized the full power of the connection. Bernshteyn was able to prove general theorems of the type “if a local problem Π can be solved with a local algorithm of round complexity $O(\log^* n)$, then any measurable graph admits a Borel solution to Π ”.

As we did not define the terms “measurable graph” or “Borel solution” precisely, we will next only state perhaps the simplest concrete example manifesting the connection, and we point an interested reader to the recent survey of Bernshteyn [\[Ber22b\]](#) that covers this very exciting and recent topic in depth.

Theorem 3.6 ([\[GR21\]](#)). *Let Π be any local problem (without inputs). Then, the following two statements are equivalent:*

1. *There exists an $o(n)$ -round local algorithm solving Π on sufficiently large oriented cycles.*
2. *One can label each vertex of G_0 with a label of Π such that all vertices satisfy the constraints of Π , and for each label ℓ , the set of vertices with that label is Lebesgue-measurable².*

By now, we know of several more results of this type, both for general bounded-degree graphs [\[Ber23a; Ber23b\]](#) and for special cases like trees [\[Bra+21\]](#) and grids [\[Gao+18; GR23\]](#).

Transfer of techniques: Many recent results in the area of descriptive combinatorics are adopting techniques from local algorithms that are, by now, familiar to readers of this text: network decomposition [\[BY23\]](#), Lovász local lemma [\[Ber23a; Ber23b; Csó+22; BW23; BY23\]](#), derandomizations [\[Ber23b; GR23\]](#), ID graphs [\[Bra+22; Bra+21\]](#), and others.

This connection also led to some progress in the area of local algorithms. For example, Bernshteyn [\[Ber22a\]](#) used ideas from a measurable Vizing theorem of Grebík and Pikhurko [\[GP20\]](#) to construct a fast local algorithm for $(\Delta+1)$ -edge coloring problem. Brandt, Chang, Grebík, Grunau, Rozhoř, and Vidnyánszky [\[Bra+22\]](#) adapted a lower bound of Marks [\[Mar16\]](#) to give a different (and similarly simple) lower bound for sinkless orientation from [Theorem 2.15](#).

3.6 Other Models

There are many more models of local/distributed/parallel/sublinear computation that are, in one way or another, related to local algorithms. In many of these models, it not only makes sense to try to solve concrete problems but often, one can also hope that the theory of local algorithms, such as the classification of the local problems on bounded-degree graphs, can be extended similarly to, e.g., [Theorem 3.2](#).

Uniform algorithms: One favorite model of the author is the model of *uniform* local complexity. In this model, a randomized local algorithm does not know the size of the input graph, n . It simply looks at larger and larger neighborhoods of a given vertex, until it decides that it has seen enough to compute the

²One can replace “Lebesgue measurable” with “a union of finitely many intervals” and the theorem still holds.

output label at the node. This model was independently studied by local algorithms community [KSV12] and community of probabilists [HSW17; Hol17; Spi20] where it is known as *finitary factors of i.i.d.*

In the $o(\log n)$ -local-complexity regime, uniform algorithms can be seen as a more powerful version of classical local algorithms. This is because there are certain local problems such that their solution requires a few nodes to see the whole graph, while most of the nodes can output their solution after seeing their $O(1)$ -radius neighborhood [GR23]. On the other hand, local problems solvable by uniform local algorithms still often admit measurable solutions defined in descriptive combinatorics [GR21; GR23; Bra+22]. Thus, uniform local algorithms can be seen as interpolating between the extremely clean picture painted by the classification of $o(\log n)$ -round algorithms (Theorem 2.1), and the much less well-understood complexity classes coming from descriptive combinatorics.

The connection between classical local algorithms and uniform ones would be much cleaner if the following problem was resolved (see [Bra+22]):

Problem 3.7. *Is there a C and a uniform randomized local algorithm on bounded degree graphs for C -relaxed Lovász local lemma such that after $\text{poly} \log \log(1/\varepsilon)$ rounds, at least $1 - \varepsilon$ fraction of nodes know the solution?*

Other models: There are many more models of locality studied in the literature. Let us list a sample of them. The list includes the averaged local complexity [Feu17; BT19; CGP20; Bal+23b] (closely connected to uniform algorithms), online local model [Akb+23; Cha+23; Akb+24], dynamic local model [Akb+23; Akb+24], local mending model [MSV23], supported local model [SS13; Kor+21], quantum local model [GKM09; LNR19; Coi+22; Akb+24], awake (energy) complexity [Cha+19b; Cha+20; CGP20; GP22; GP23], local certification [GS11; FHK12; Fra+13; Feu21], finitely dependent colorings [HL+15; Hol23], or computable combinatorics [QW22].

Final Remarks

This text aims to present the area of local algorithms in a way that highlights what the author sees as its key aspect: Unlike typical subareas of computer science and discrete mathematics that are usually unified by a set of useful techniques and important results, the area of local algorithms extends beyond this norm. We have a clean theory of the model of local algorithms that leads to a clear complexity-theoretical picture. We also understand that local algorithms have applications to a number of other models studied by the broader algorithmic community. In this sense, local complexity is similar to much more established fields like communication complexity. The author believes that the theory has the potential for diverse extensions and applications and encourages the reader to identify and explore the next one.

Bibliography

- [Akb+24] Amirreza Akbari, Xavier Coiteux-Roy, Francesco d’Amore, François Le Gall, Henrik Lievonen, Darya Melnyk, Augusto Modanese, Shreyas Pai, Marc-Olivier Renou, Václav Rozho , et al. “Online Locality Meets Distributed Quantum Computing”. In: *arXiv preprint arXiv:2403.01903* (2024).
- [Akb+23] Amirreza Akbari, Navid Eslami, Henrik Lievonen, Darya Melnyk, Joonas Särkijärvi, and Jukka Suomela. “Locality in Online, Dynamic, Sequential, and Distributed Graph Algorithms”. In: *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*. Ed. by Kousha Etessami, Uriel Feige, and Gabriele Puppis. Vol. 261. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 10:1–10:20. ISBN: 978-3-95977-278-5. DOI: [10.4230/LIPIcs.ICALP.2023.10](https://doi.org/10.4230/LIPIcs.ICALP.2023.10). URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2023.10>.
- [ASS99] Alexander S. Kechris, Sławomir Solecki, and Stevo Todorcević. “Borel chromatic numbers.” In: *Adv. Math.* 141.1 (1999), pp. 1–44.
- [ABI86] Noga Alon, Laszlo Babai, and Alon Itai. “A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem”. In: *Journal of Algorithms* 7.4 (1986), pp. 567–583.
- [Alo+12] Noga Alon, Ronitt Rubinfeld, Shai Vardi, and Ning Xie. “Space-efficient local computation algorithms”. In: *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*. SIAM. 2012, pp. 1132–1139.
- [Alt+93] Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. “On sparse spanners of weighted graphs”. In: *Discrete & Computational Geometry* 9.1 (1993), pp. 81–100.
- [AB09] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- [Åst+10] Matti Åstrand, Valentin Polishchuk, Joel Rybicki, Jukka Suomela, and Jara Uitto. “Local algorithms in (weakly) coloured graphs”. In: *arXiv preprint arXiv:1002.0125* (2010).
- [Awe+89] Baruch Awerbuch, Andrew V. Goldberg, Michael Luby, and Serge A. Plotkin. “Network Decomposition and Locality in Distributed Computation”. In: *Proc. 30th IEEE Symp. on Foundations of Computer Science (FOCS)*. 1989, pp. 364–369.
- [Bal+19a] Alkida Balliu, Sebastian Brandt, Yi-Jun Chang, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. “The Distributed Complexity of Locally Checkable Problems on Paths is Decidable”. In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. PODC ’19. Toronto ON, Canada: Association for Computing Machinery, 2019, pp. 262–271. ISBN: 9781450362177.
- [Bal+22a] Alkida Balliu, Sebastian Brandt, Yi-Jun Chang, Dennis Olivetti, Jan Studený, and Jukka Suomela. “Efficient Classification of Locally Checkable Problems in Regular Trees”. In: *36th International Symposium on Distributed Computing*. 2022.
- [Bal+20a] Alkida Balliu, Sebastian Brandt, Yuval Efron, Juho Hirvonen, Yannic Maus, Dennis Olivetti, and Jukka Suomela. “Classification of Distributed Binary Labeling Problems”. In: *34th International Symposium on Distributed Computing (DISC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. 2020.

- [Bal+22b] Alkida Balliu, Sebastian Brandt, Manuela Fischer, Rustam Latypov, Yannic Maus, Dennis Olivetti, and Jara Uitto. "Exponential Speedup over Locality in MPC with Optimal Memory." In: *36th International Symposium on Distributed Computing: DISC 2022*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. 2022, pp. 9–1.
- [Bal+19b] Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. "Lower bounds for maximal matchings and maximal independent sets". In: *Proc. Foundations of Computer Science (FOCS)*. 2019.
- [Bal+22c] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. "Distributed Δ -coloring plays hide-and-seek". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 464–477.
- [Bal+22d] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. "Distributed edge coloring in time polylogarithmic in Δ ". In: *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*. 2022, pp. 15–25.
- [Bal+23a] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. "Distributed maximal matching and maximal independent set on hypergraphs". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 2632–2676.
- [Bal+21a] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. "Improved distributed lower bounds for MIS and bounded (out-) degree dominating sets in trees". In: *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*. 2021, pp. 283–293.
- [Bal+23b] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, Dennis Olivetti, and Gustav Schmid. "On the Node-Averaged Complexity of Locally Checkable Problems on Trees". In: *37th International Symposium on Distributed Computing*. 2023.
- [BBO20] Alkida Balliu, Sebastian Brandt, and Dennis Olivetti. "Distributed Lower Bounds for Ruling Sets". In: *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*. 2020, pp. 365–376. DOI: [10.1109/FOCS46700.2020.00042](https://doi.org/10.1109/FOCS46700.2020.00042). URL: <https://doi.org/10.1109/FOCS46700.2020.00042>.
- [Bal+21b] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, Jan Studený, Jukka Suomela, and Aleksandr Tereshchenko. "Locally checkable problems in rooted trees". In: *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*. 2021, pp. 263–272.
- [Bal+20b] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, and Jukka Suomela. "Almost global problems in the LOCAL model". In: *Distributed Computing (2020)*, pp. 1–23.
- [Bal+20c] Alkida Balliu, Sebastian Brandt, Dennis Olivetti, and Jukka Suomela. "How Much Does Randomness Help with Locally Checkable Problems?" In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. PODC '20. Virtual Event, Italy: Association for Computing Machinery, 2020, pp. 299–308. ISBN: 9781450375825.
- [Bal+21c] Alkida Balliu, Keren Censor-Hillel, Yannic Maus, Dennis Olivetti, and Jukka Suomela. "Locally Checkable Labelings with Small Messages". In: *35th International Symposium on Distributed Computing*. 2021.
- [Bal+18] Alkida Balliu, Juho Hirvonen, Janne H Korhonen, Tuomo Lempiäinen, Dennis Olivetti, and Jukka Suomela. "New classes of distributed time complexity". In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 2018, pp. 1307–1318.
- [Bal+19c] Alkida Balliu, Juho Hirvonen, Christoph Lenzen, Dennis Olivetti, and Jukka Suomela. "Locality of not-so-weak coloring". In: *International Colloquium on Structural Information and Communication Complexity*. Springer. 2019, pp. 37–51.
- [Bal+19d] Alkida Balliu, Juho Hirvonen, Dennis Olivetti, and Jukka Suomela. "Hardness of minimal symmetry breaking in distributed computing". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. 2019, pp. 369–378.
- [Bal+23c] Alkida Balliu, Janne H Korhonen, Fabian Kuhn, Henrik Lievonen, Dennis Olivetti, Shreyas Pai, Ami Paz, Joel Rybicki, Stefan Schmid, Jan Studený, et al. "Sinkless Orientation Made Simple". In: *Symposium on Simplicity in Algorithms (SOSA)*. SIAM. 2023, pp. 175–191.

- [BKO20] Alkida Balliu, Fabian Kuhn, and Dennis Olivetti. "Distributed edge coloring in time quasi-polylogarithmic in Δ ". In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. 2020, pp. 289–298.
- [BKO22] Alkida Balliu, Fabian Kuhn, and Dennis Olivetti. "Improved Distributed Fractional Coloring Algorithms". In: *25th International Conference on Principles of Distributed Systems*. 2022, p. 1.
- [BKM20] Philipp Bamberger, Fabian Kuhn, and Yannic Maus. "Efficient Deterministic Distributed Coloring with Small Bandwidth". In: *Proc. Principles of Distributed Computing (PODC)*. 2020, to appear, arXiv:1912.02814.
- [Bar15] Leonid Barenboim. "Deterministic $(\Delta + 1)$ -Coloring in Sublinear (in Δ) Time in Static, Dynamic and Faulty Networks". In: *Proc. 34th ACM Symp. on Principles of Distributed Computing (PODC)*. 2015, pp. 345–354.
- [BE10] Leonid Barenboim and Michael Elkin. "Deterministic Distributed Vertex Coloring in Polylogarithmic Time". In: *Proc. 29th Symp. on Principles of Distributed Computing (PODC)*. 2010, pp. 410–419.
- [BE11] Leonid Barenboim and Michael Elkin. "Deterministic distributed vertex coloring in polylogarithmic time". In: *Journal of the ACM (JACM)* 58.5 (2011), pp. 1–25.
- [BE13] Leonid Barenboim and Michael Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments*. Morgan & Claypool Publishers, 2013.
- [BE08] Leonid Barenboim and Michael Elkin. "Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition". In: *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*. 2008, pp. 25–34.
- [BEG18] Leonid Barenboim, Michael Elkin, and Uri Goldenberg. "Locally-Iterative Distributed $(\Delta + 1)$ -Coloring below Szegedy-Vishwanathan Barrier, and Applications to Self-Stabilization and to Restricted-Bandwidth Models". In: *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*. PODC '18. Egham, United Kingdom: Association for Computing Machinery, 2018, pp. 437–446. ISBN: 9781450357951.
- [BEK15] Leonid Barenboim, Michael Elkin, and Fabian Kuhn. "Distributed $(\Delta + 1)$ -Coloring in Linear (in Δ) Time". In: *SIAM Journal on Computing* 43.1 (2015), pp. 72–95.
- [Bar+16] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. "The Locality of Distributed Symmetry Breaking". In: *Journal of the ACM* 63 (3 2016), 20:1–20:45.
- [BT19] Leonid Barenboim and Yaniv Tzur. "Distributed symmetry-breaking with improved vertex-averaged complexity". In: *20th International Conference on Distributed Computing and Networking, ICDCN 2019*. Association for Computing Machinery. 2019, pp. 31–40.
- [BS07] Surender Baswana and Sandeep Sen. "A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs". In: *Random Structures & Algorithms* 30.4 (2007), pp. 532–563.
- [Bec91] József Beck. "An algorithmic approach to the Lovász local lemma. I". In: *Random Structures & Algorithms* 2.4 (1991), pp. 343–365.
- [Beh22] Soheil Behnezhad. "Time-optimal sublinear algorithms for matching and vertex cover". In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 873–884.
- [Ber22a] Anton Bernshteyn. "A fast distributed algorithm for $(\Delta + 1)$ -edge-coloring". In: *Journal of Combinatorial Theory, Series B* 152 (2022), pp. 319–352.
- [Ber22b] Anton Bernshteyn. "Descriptive combinatorics and distributed algorithms". In: *NOTICES OF THE AMERICAN MATHEMATICAL SOCIETY* 69.9 (2022).
- [Ber23a] Anton Bernshteyn. "Distributed algorithms, the Lovász Local Lemma, and descriptive combinatorics". In: *Inventiones mathematicae* 233.2 (2023), pp. 495–542.
- [Ber23b] Anton Bernshteyn. "Probabilistic constructions in continuous combinatorics and a bridge to distributed algorithms". In: *Advances in Mathematics* 415 (2023), p. 108895.

- [BD23a] Anton Bernshteyn and Abhishek Dhawan. "Borel Vizing's Theorem for Graphs of Subexponential Growth". In: *arXiv e-prints* (2023), arXiv:2307.
- [BD23b] Anton Bernshteyn and Abhishek Dhawan. "Fast algorithms for Vizing's theorem on bounded degree graphs". In: *arXiv preprint arXiv:2303.05408* (2023).
- [BW23] Anton Bernshteyn and Felix Weilacher. "Borel versions of the Local Lemma and LOCAL algorithms for graphs of finite asymptotic separation index". In: *arXiv preprint arXiv:2308.14941* (2023).
- [BY23] Anton Bernshteyn and Jing Yu. "Coarse embeddings into grids and asymptotic dimension for Borel graphs of polynomial growth". In: *arXiv preprint arXiv:2302.04727* (2023).
- [Bez+22] Marcel Bezdrighin, Michael Elkin, Mohsen Ghafari, Christoph Grunau, Bernhard Haeupler, Saeed Ilchi, and Václav Rozhoň. "Deterministic distributed sparse and ultra-sparse spanners and connectivity certificates". In: *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures*. 2022, pp. 1–10.
- [Bra+16] S. Brandt, O. Fischer, J. Hirvonen, B. Keller, T. Lempiäinen, J. Rybicki, J. Suomela, and J. Uitto. "A lower bound for the distributed Lovász local lemma". In: *Proc. 48th ACM Symp. on Theory of Computing (STOC)*. 2016, pp. 479–488.
- [Bra19] Sebastian Brandt. "An automatic speedup theorem for distributed problems". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. 2019, pp. 379–388.
- [Bra+22] Sebastian Brandt, Yi-Jun Chang, Jan Grebík, Christoph Grunau, Václav Rozhoň, and Zoltán Vidnyánszky. "Local Problems on Trees from the Perspectives of Distributed Algorithms, Finitary Factors, and Descriptive Combinatorics". In: *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. 2022.
- [Bra+21] Sebastian Brandt, Yi-Jun Chang, Jan Grebík, Christoph Grunau, Václav Rozhoň, and Zoltán Vidnyánszky. "On homomorphism graphs". In: *arXiv preprint arXiv:2111.03683* (2021).
- [BGR20] Sebastian Brandt, Christoph Grunau, and Václav Rozhoň. "Generalizing the Sharp Threshold Phenomenon for the Distributed Complexity of the Lovász Local Lemma". In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. 2020, pp. 329–338.
- [BGR21] Sebastian Brandt, Christoph Grunau, and Václav Rozhoň. "The randomized local computation complexity of the Lovász local lemma". In: *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*. 2021, pp. 307–317.
- [Bra+17] Sebastian Brandt, Juho Hirvonen, Janne H. Korhonen, Tuomo Lempiäinen, Patric R.J. Östergård, Christopher Purcell, Joel Rybicki, Jukka Suomela, and Przemysław Uznański. "LCL Problems on Grids". In: *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC '17*. Washington, DC, USA: Association for Computing Machinery, 2017, pp. 101–110. ISBN: 9781450349925.
- [BMU19] Sebastian Brandt, Yannic Maus, and Jara Uitto. "A sharp threshold phenomenon for the distributed complexity of the Lovász local lemma". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. 2019, pp. 389–398.
- [BO20] Sebastian Brandt and Dennis Olivetti. "Truly tight-in- δ bounds for bipartite maximal matching and variants". In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. 2020, pp. 69–78.
- [CPS17a] Keren Censor-Hillel, Merav Parter, and Gregory Schwartzman. "Derandomizing Local Distributed Algorithms under Bandwidth Restrictions". In: *31st International Symposium on Distributed Computing (DISC 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2017.
- [Cha20] Yi-Jun Chang. "The Complexity Landscape of Distributed Locally Checkable Problems on Trees". In: *34th International Symposium on Distributed Computing (DISC 2020)*. Ed. by Hagit Attiya. Vol. 179. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 18:1–18:17. ISBN: 978-3-95977-168-9. DOI: [10.4230/LIPIcs.DISC.2020.18](https://doi.org/10.4230/LIPIcs.DISC.2020.18). URL: <https://drops.dagstuhl.de/opus/vol1texte/2020/13096>.

- [Cha24] Yi-Jun Chang. "The Distributed Complexity of Locally Checkable Labeling Problems Beyond Paths and Trees". In: *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik. 2024.
- [Cha+20] Yi-Jun Chang, Varsha Dani, Thomas P Hayes, and Seth Pettie. "The energy complexity of BFS in radio networks". In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. 2020, pp. 273–282.
- [Cha+19a] Yi-Jun Chang, Manuela Fischer, Mohsen Ghafari, Jara Uitto, and Yufan Zheng. "The Complexity of $(\Delta + 1)$ Coloring in Congested Clique, Massively Parallel Computation, and Centralized Local Computation". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. ACM. 2019, pp. 471–480.
- [CG21] Yi-Jun Chang and Mohsen Ghafari. "Strong-Diameter Network Decomposition". In: *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*. PODC'21. Virtual Event, Italy: Association for Computing Machinery, 2021, pp. 273–281. ISBN: 9781450385480.
- [Cha+18] Yi-Jun Chang, Qizheng He, Wenzheng Li, Seth Pettie, and Jara Uitto. "The complexity of distributed edge coloring with small palettes". In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2018, pp. 2633–2652.
- [CKP19] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. "An Exponential Separation between Randomized and Deterministic Complexity in the LOCAL Model". In: *SIAM Journal on Computing* 48.1 (2019), pp. 122–143.
- [Cha+19b] Yi-Jun Chang, Tsvi Kopelowitz, Seth Pettie, Ruosong Wang, and Wei Zhan. "Exponential separations in the energy complexity of leader election". In: *ACM Transactions on Algorithms (TALG)* 15.4 (2019), pp. 1–31.
- [CLP18] Yi-Jun Chang, Wenzheng Li, and Seth Pettie. "An Optimal Distributed $(\Delta + 1)$ -Coloring Algorithm?" In: *Proc. 50th ACM Symp. on Theory of Computing (STOC)*. 2018.
- [Cha+23] Yi-Jun Chang, Gopinath Mishra, Hung Tuan Nguyen, Mingyang Yang, and Yu-Cheng Yeh. "A Tight Lower Bound for 3-Coloring Grids in the Online-LOCAL Model". In: *arXiv preprint arXiv:2312.01384* (2023).
- [CP19] Yi-Jun Chang and Seth Pettie. "A Time Hierarchy Theorem for the LOCAL Model". In: *SIAM Journal on Computing* 48.1 (2019), pp. 33–69.
- [CPZ19] Yi-Jun Chang, Seth Pettie, and Hengjie Zhang. "Distributed triangle detection via expander decomposition". In: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2019, pp. 821–840.
- [CS19] Yi-Jun Chang and Thatchaphol Saranurak. "Improved distributed expander decomposition and nearly optimal triangle enumeration". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. 2019, pp. 66–73.
- [CGP20] Soumyottam Chatterjee, Robert Gmyr, and Gopal Pandurangan. "Sleeping is efficient: Mis in $o(1)$ -rounds node-averaged awake complexity". In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. 2020, pp. 99–108.
- [CPS17b] Kai-Min Chung, Seth Pettie, and Hsin-Hao Su. "Distributed algorithms for the Lovász local lemma and graph coloring". In: *Distributed Computing* 30.4 (2017), pp. 261–280.
- [Coi+22] Xavier Coiteux-Roy, Francesco d'Amore, Rishikesh Gajjala, Fabian Kuhn, François Le Gall, Henrik Lievonen, Augusto Modanese, Marc-Olivier Renou, Gustav Schmid, and Jukka Suomela. "No distributed quantum advantage for approximate graph coloring." In: *Quantum Information and Computation* 22.15&16 (2022), pp. 1261–1279.
- [CV86] Richard Cole and Uzi Vishkin. "Deterministic Coin Tossing with Applications to Optimal Parallel List Ranking". In: *Information and Control* 70.1 (1986), pp. 32–53.
- [CGP24] Clinton T Conley, Jan Grebik, and Oleg Pikhurko. "Divisibility of spheres with measurable pieces". In: *L'Enseignement Mathématique* (2024).

- [CM17] Clinton T. Conley and Benjamin D. Miller. "Measure reducibility of countable Borel equivalence relations". In: *Ann. of Math. (2)* 185.2 (2017), pp. 347–402. ISSN: 0003-486X.
- [Csó+22] Endre Csóka, Łukasz Grabowski, András Máthé, Oleg Pikhurko, and Konstantinos Tyros. "Moser-Tardos Algorithm with small number of random bits". In: *arXiv preprint arXiv:2203.05888* (2022).
- [CP23] Endre Csóka and András Pongrácz. *Local algorithms for the maximum flow and minimum cut in bounded-degree networks*. 2023. arXiv: [1005.0513](https://arxiv.org/abs/1005.0513) [cs. DS].
- [DS24] Sameep Dahal and Jukka Suomela. "Distributed half-integral matching and beyond". In: *Theoretical Computer Science* 982 (2024), p. 114278.
- [Der+08] Bilel Derbel, Cyril Gavoille, David Peleg, and Laurent Viennot. "On the Locality of Distributed Sparse Spanner Construction". In: *Proceedings of the Twenty-Seventh ACM Symposium on Principles of Distributed Computing*. PODC '08. Toronto, Canada: Association for Computing Machinery, 2008, pp. 273–282. ISBN: 9781595939890.
- [DKM19] Janosch Deurer, Fabian Kuhn, and Yannic Maus. "Deterministic Distributed Dominating Set Approximation in the CONGEST Model". In: *Proc. Principles of Distributed Computing (PODC)*. 2019, pp. 94–103.
- [DF92] Randall Dougherty and Matthew Foreman. "Banach-Tarski paradox using pieces with the property of Baire". In: *Proc. Nat. Acad. Sci. U.S.A.* 89.22 (1992), pp. 10726–10728. ISSN: 0027-8424.
- [EL10] Gábor Elek and Gábor Lippner. "Borel oracles. An analytical approach to constant-time algorithms". In: *Proceedings of the American Mathematical Society* 138.8 (2010), pp. 2939–2947.
- [EN17] Michael Elkin and Ofer Neiman. "Efficient algorithms for constructing very sparse spanners and emulators". In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2017, pp. 652–669.
- [EFF82] Paul Erdős, Peter Frankl, and Zoltán Füredi. "Families of finite sets in which no set is covered by the union of two others". In: *Journal of Combinatorial Theory, Series A* 33.2 (1982), pp. 158–166.
- [EL75] Paul Erdős and László Lovász. "Problems and results on 3-chromatic hypergraphs and some related questions". In: *Infinite and finite sets* 10.2 (1975), pp. 609–627.
- [EMR14] Guy Even, Moti Medina, and Dana Ron. "Deterministic stateless centralized local algorithms for bounded degree graphs". In: *European Symposium on Algorithms*. Springer. 2014, pp. 394–405.
- [Fao+23] Salwa Faour, Mohsen Ghafari, Christoph Grunau, Fabian Kuhn, and Václav Rozhoň. "Local distributed rounding: Generalized to MIS, matching, set cover, and beyond". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 4409–4447.
- [Feu17] Laurent Feuilloley. "How long it takes for an ordinary node with an ordinary ID to output?" In: *International Colloquium on Structural Information and Communication Complexity*. Springer. 2017, pp. 263–282.
- [Feu21] Laurent Feuilloley. "Introduction to local certification". In: *Discrete Mathematics & Theoretical Computer Science* 23. Distributed Computing and Networking (2021).
- [FG17] Manuela Fischer and Mohsen Ghafari. "Sublogarithmic Distributed Algorithms for Lovász Local Lemma, and the Complexity Hierarchy". In: *31 International Symposium on Distributed Computing*. 2017.
- [FGK17] Manuela Fischer, Mohsen Ghafari, and Fabian Kuhn. "Deterministic Distributed Edge-Coloring via Hypergraph Maximal Matching". In: *Proc. 58th IEEE Symp. on Foundations of Computer Science (FOCS)*. 2017.
- [FMU22] Manuela Fischer, Slobodan Mitrović, and Jara Uitto. "Deterministic $(1 + \epsilon)$ -approximate maximum matching with poly $(1/\epsilon)$ passes in the semi-streaming model and beyond". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 248–260.

- [FHK16] P. Fraigniaud, M. Heinrich, and A. Kosowski. "Local Conflict Coloring". In: *Proc. 57th IEEE Symp. on Foundations of Computer Science (FOCS)*. 2016, pp. 625–634.
- [Fra+13] Pierre Fraigniaud, Mika Göös, Amos Korman, and Jukka Suomela. "What Can Be Decided Locally without Identifiers?" In: *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*. PODC '13. Montréal, Québec, Canada: Association for Computing Machinery, 2013, pp. 157–165. ISBN: 9781450320658.
- [FHK12] Pierre Fraigniaud, Magnús M Halldórsson, and Amos Korman. "On the impact of identifiers on local decision". In: *International Conference On Principles Of Distributed Systems*. Springer. 2012, pp. 224–238.
- [FK23] Marc Fuchs and Fabian Kuhn. "List Defective Colorings: Distributed Algorithms and Applications". In: *37th International Symposium on Distributed Computing (DISC 2023)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik. 2023.
- [Gab00] Damien Gaboriau. "Coût des relations d'équivalence et des groupes". In: *Invent. Math.* 139.1 (2000), pp. 41–98. ISSN: 0020-9910.
- [Gao+18] Su Gao, Steve Jackson, Edward Krohne, and Brandon Seward. "Continuous combinatorics of abelian group actions". In: *arXiv preprint arXiv:1803.03872* (2018).
- [Gav+09] Cyril Gavoille, Ralf Klasing, Adrian Kosowski, Łukasz Kuszner, and Alfredo Navarra. "On the complexity of distributed graph coloring with local minimality constraints". In: *Networks: An International Journal* 54.1 (2009), pp. 12–19.
- [GKM09] Cyril Gavoille, Adrian Kosowski, and Marcin Markiewicz. "What Can Be Observed Locally? Round-Based Models for Quantum Distributed Computing". In: *Proceedings of the 23rd International Conference on Distributed Computing*. DISC'09. Elche, Spain: Springer-Verlag, 2009, pp. 243–257. ISBN: 3642043542.
- [Gha16] Mohsen Ghafari. "An improved distributed algorithm for maximal independent set". In: *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 2016, pp. 270–277.
- [Gha22] Mohsen Ghafari. "Local computation of maximal independent set". In: *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 438–449.
- [GG23] Mohsen Ghafari and Christoph Grunau. "Faster Deterministic Distributed MIS and Approximate Matching". In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 1777–1790.
- [Gha+23] Mohsen Ghafari, Christoph Grunau, Bernhard Haeupler, Saeed Ilchi, and Václav Rozhořel. "Improved distributed network decomposition, hitting sets, and spanners, via derandomization". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 2532–2566.
- [GGR21] Mohsen Ghafari, Christoph Grunau, and Václav Rozhořel. "Improved Deterministic Network Decomposition". In: *Proc. of the 32nd ACM-SIAM Symp. on Discrete Algorithms (SODA)*. USA: Society for Industrial and Applied Mathematics, 2021, pp. 2904–2923. ISBN: 9781611976465.
- [GH21] Mohsen Ghafari and Bernhard Haeupler. "A time-optimal randomized parallel algorithm for mis". In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2021, pp. 2892–2903.
- [GHK18] Mohsen Ghafari, David Harris, and Fabian Kuhn. "On derandomizing local distributed algorithms". In: *Proc. Foundations of Computer Science (FOCS)*. 2018, pp. 662–673.
- [Gha+18a] Mohsen Ghafari, Juho Hirvonen, Fabian Kuhn, and Yannic Maus. "Improved distributed delta-coloring". In: *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*. 2018, pp. 427–436.
- [GK22] Mohsen Ghafari and Fabian Kuhn. "Deterministic distributed vertex coloring: Simpler, faster, and without network decomposition". In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 1009–1020.

- [GKM17] Mohsen Ghafari, Fabian Kuhn, and Yannic Maus. "On the Complexity of Local Distributed Graph Problems". In: *Proc. 49th ACM Symp. on Theory of Computing (STOC)*. 2017, pp. 784–797.
- [Gha+18b] Mohsen Ghafari, Fabian Kuhn, Yannic Maus, and Jara Uitto. "Deterministic Distributed Edge-Coloring with Fewer Colors". In: *Proc. 50th ACM Symp. on Theory of Computing (STOC)*. 2018.
- [GKU19] Mohsen Ghafari, Fabian Kuhn, and Jara Uitto. "Conditional Hardness Results for Massively Parallel Computation from Distributed Lower Bounds". In: *Proc. Foundations of Computer Science (FOCS)*. 2019.
- [GP22] Mohsen Ghafari and Julian Portmann. "Average Awake Complexity of MIS and Matching". In: *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures*. SPAA '22. Philadelphia, PA, USA: Association for Computing Machinery, 2022, pp. 45–55. ISBN: 9781450391467.
- [GP23] Mohsen Ghafari and Julian Portmann. "Distributed MIS with Low Energy and Time Complexities". In: *Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing*. PODC '23. Orlando, FL, USA: Association for Computing Machinery, 2023, pp. 146–156.
- [GS17] Mohsen Ghafari and Hsin-Hao Su. "Distributed Degree Splitting, Edge Coloring, and Orientations". In: *Proc. 28th ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 2017, pp. 2505–2523.
- [GU19] Mohsen Ghafari and Jara Uitto. "Sparsifying distributed algorithms with ramifications in massively parallel computation and centralized local computation". In: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2019, pp. 1636–1653.
- [Gha20] Ghafari, Mohsen. *Distributed Graph Algorithms*. Lecture notes. 2020. URL: <https://disco.ethz.ch/courses/fs20/podc/lecturenotes/DGA.pdf>.
- [GPS88] Andrew V. Goldberg, Sergey A. Plotkin, and Gregory E. Shannon. "Parallel Symmetry-Breaking in Sparse Graphs". In: *SIAM Journal on Discrete Mathematics* 1.4 (1988), pp. 434–446.
- [Göo+16] Mika Göös, Juho Hirvonen, Reut Levi, Moti Medina, and Jukka Suomela. "Non-local probes do not help with many graph problems". In: *Distributed Computing: 30th International Symposium, DISC 2016, Paris, France, September 27–29, 2016. Proceedings 30*. Springer. 2016, pp. 201–214.
- [GHS14] Mika Göös, Juho Hirvonen, and Jukka Suomela. "Linear-in-Delta Lower Bounds in the LOCAL Model". In: *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*. PODC '14. Paris, France: Association for Computing Machinery, 2014, pp. 86–95. ISBN: 9781450329446.
- [GHS13] Mika Göös, Juho Hirvonen, and Jukka Suomela. "Lower Bounds for Local Approximation". In: *J. ACM* 60.5 (Oct. 2013). ISSN: 0004-5411.
- [GS11] Mika Göös and Jukka Suomela. "Locally Checkable Proofs". In: *Proc. 30th Symp. on Principles of Distributed Computing (PODC)*. 2011, pp. 159–168.
- [GRS91] Ronald L Graham, Bruce L Rothschild, and Joel H Spencer. *Ramsey theory*. Vol. 20. John Wiley & Sons, 1991.
- [Gre23] Jan Grebík. "Measurable Vizing's theorem". In: *arXiv preprint arXiv:2303.16440* (2023).
- [GP20] Jan Grebík and Oleg Pikhurko. "Measurable versions of Vizing's theorem". In: *Advances in Mathematics* 374 (2020), p. 107378.
- [GR21] Jan Grebík and Václav Rozhoř. "Classification of local problems on paths from the perspective of descriptive combinatorics". In: *Extended Abstracts EuroComb 2021: European Conference on Combinatorics, Graph Theory and Applications*. Springer. 2021, pp. 553–559.
- [GR23] Jan Grebík and Václav Rozhoř. "Local problems on grids from the perspective of distributed algorithms, finitary factors, and descriptive combinatorics". In: *Advances in Mathematics* 431 (2023), p. 109241.

- [GRB22] Christoph Grunau, Václav Rozhoř, and Sebastian Brandt. "The landscape of distributed complexities on trees and beyond". In: *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*. 2022, pp. 37–47.
- [HKP01] M. Håkoniemi, M. Karoński, and A. Panconesi. "On the Distributed Complexity of Computing Maximal Matchings". In: *SIAM Journal on Discrete Math.* 15.1 (2001), pp. 41–57.
- [Har19] David G Harris. "Distributed local approximation algorithms for maximum matching in graphs and hypergraphs". In: *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2019, pp. 700–724.
- [HSS16] David G. Harris, Johannes Schneider, and Hsin-Hao Su. "Distributed $(\Delta+1)$ -Coloring in Sublogarithmic Rounds". In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '16. Cambridge, MA, USA: Association for Computing Machinery, 2016, pp. 465–478. ISBN: 9781450341325.
- [HS20] Juho Hirvonen and Jukka Suomela. "Distributed Algorithms". In: *Textbook, Aalto University* (2020).
- [HS12] Juho Hirvonen and Jukka Suomela. "Distributed Maximal Matching: Greedy is Optimal". In: *Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing*. PODC '12. Madeira, Portugal: Association for Computing Machinery, 2012, pp. 165–174. ISBN: 9781450314503.
- [HL+15] Alexander Holroyd, Thomas Liggett, et al. "Symmetric 1-dependent colorings of the integers". In: *Electronic Communications in Probability* 20 (2015).
- [Hol23] Alexander E Holroyd. "Symmetrization for finitely dependent colouring". In: *arXiv preprint arXiv:2305.13980* (2023).
- [Hol17] Alexander E. Holroyd. "One-dependent coloring by finitary factors". In: *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* 53.2 (2017), pp. 753–765.
- [HSW17] Alexander E. Holroyd, Oded Schramm, and David B. Wilson. "Finitary coloring". In: *Ann. Probab.* 45.5 (Sept. 2017), pp. 2867–2898. URL: <https://doi.org/10.1214/16-AOP1127>.
- [JáJ92] Joseph JáJá. *An introduction to parallel algorithms*. Addison Wesley Longman Publishing Co., Inc., 1992.
- [Joh99] Ojvind Johansson. "Simple distributed $\Delta+1$ -coloring of graphs". In: *Information Processing Letters* 70.5 (1999), pp. 229–232. ISSN: 0020-0190.
- [KSV10] Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii. "A model of computation for MapReduce". In: *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*. SIAM. 2010, pp. 938–948.
- [KM16] Alexander S Kechris and Andrew S Marks. "Descriptive graph combinatorics". In: *preprint* 2020.9 (2016).
- [KSS81] Daniel J. Kleitman, J Shearer, and Dean Sturtevant. "Intersections of k-element sets". In: *Combinatorica* 1.4 (1981), pp. 381–384.
- [Kor+21] Janne H. Korhonen, Ami Paz, Joel Rybicki, Stefan Schmid, and Jukka Suomela. "Brief Announcement: Sinkless Orientation Is Hard Also in the Supported LOCAL Model". In: *35th International Symposium on Distributed Computing (DISC 2021)*. Ed. by Seth Gilbert. Vol. 209. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, 58:1–58:4. ISBN: 978-3-95977-210-5. DOI: [10.4230/LIPIcs.DISC.2021.58](https://doi.org/10.4230/LIPIcs.DISC.2021.58). URL: <https://drops.dagstuhl.de/opus/vol1texte/2021/14860>.
- [KSV12] Amos Korman, Jean-Sébastien Sereni, and Laurent Viennot. "Toward more localized local algorithms: removing assumptions concerning global knowledge". In: *Distributed Computing* 26.5-6 (Sept. 2012), pp. 289–308. ISSN: 1432-0452. DOI: [10.1007/s00446-012-0174-8](https://doi.org/10.1007/s00446-012-0174-8). URL: <http://dx.doi.org/10.1007/s00446-012-0174-8>.
- [Kuh20] Fabian Kuhn. "Faster deterministic distributed coloring through recursive list coloring". In: *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2020, pp. 1244–1259.

- [Kuh09] Fabian Kuhn. "Local Weak Coloring Algorithms and Implications on Deterministic Symmetry Breaking". In: *Proc. 21st ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*. 2009.
- [KMW16] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. "Local Computation: Lower and Upper Bounds". In: *Journal of the ACM* 63.2 (2016).
- [KW06] Fabian Kuhn and Roger Wattenhofer. "On the Complexity of Distributed Graph Coloring". In: *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing*. PODC '06. Denver, Colorado, USA: Association for Computing Machinery, 2006, pp. 7–15. ISBN: 1595933840.
- [LNR19] François Le Gall, Harumichi Nishimura, and Ansis Rosmanis. "Quantum Advantage for the LOCAL Model in Distributed Computing". In: *36th International Symposium on Theoretical Aspects of Computer Science*. 2019.
- [LRY15] Reut Levi, Ronitt Rubinfeld, and Anak Yodpinyanee. "Local computation algorithms for graphs of non-constant degrees". In: *Proceedings of the 27th ACM symposium on Parallelism in Algorithms and Architectures*. 2015, pp. 59–61.
- [LPS23] Henrik Lievonen, Timothé Picavet, and Jukka Suomela. "Distributed Binary Labeling Problems in High-Degree Graphs". In: *arXiv preprint arXiv:2312.12243* (2023).
- [Lin92] Nati Linial. "Locality in Distributed Graph Algorithms". In: *SIAM Journal on Computing* 21.1 (1992), pp. 193–201.
- [LS93] Nati Linial and Michael Saks. "Low Diameter Graph Decompositions". In: *Combinatorica* 13.4 (1993), pp. 441–454.
- [Lov12] László Lovász. *Large networks and graph limits*. Vol. 60. American Mathematical Soc., 2012.
- [Lub86] Michael Luby. "A Simple Parallel Algorithm for the Maximal Independent Set Problem". In: *SIAM Journal on Computing* 15 (1986), pp. 1036–1053.
- [LAO17] Lukasz Grabowski, András Máthé, and Oleg Pikhurko. "Measurable circle squaring." In: *Ann. of Math.* 185.2 (2017), pp. 671–710.
- [MU16] Andrew Marks and Spencer Unger. "Baire measurable paradoxical decompositions via matchings". In: *Advances in Mathematics* 289 (2016), pp. 397–410.
- [Mar16] Andrew S Marks. "A determinacy approach to Borel combinatorics". In: *Journal of the American Mathematical Society* 29.2 (2016), pp. 579–600.
- [MU17] Andrew S Marks and Spencer T Unger. "Borel circle squaring". In: *Annals of Mathematics* 186.2 (2017), pp. 581–605.
- [Mau] Y. Maus. personal communication.
- [Mau21] Yannic Maus. "Distributed graph coloring made easy". In: *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures*. 2021, pp. 362–372.
- [MT20] Yannic Maus and Tigran Tonoyan. "Local Conflict Coloring Revisited: Linial for Lists". In: *34th International Symposium on Distributed Computing (DISC 2020)*. Ed. by Hagit Attiya. Vol. 179. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020, 16:1–16:18. ISBN: 978-3-95977-168-9. DOI: [10.4230/LIPIcs.DISC.2020.16](https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.DISC.2020.16). URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.DISC.2020.16>.
- [MSV23] Darya Melnyk, Jukka Suomela, and Neven Villani. "Mending Partial Solutions with Few Changes". In: *26th International Conference on Principles of Distributed Systems*. 2023.
- [Mik90] Miklós Laczkovich. "Equidecomposability and discrepancy; a solution of Tarski's circle-squaring problem." In: *J. Reine Angew. Math.* 404 (1990), pp. 77–117.
- [MPX13] Gary L Miller, Richard Peng, and Shen Chen Xu. "Parallel graph decompositions using random shifts". In: *Proceedings of the twenty-fifth annual ACM symposium on Parallelism in algorithms and architectures*. 2013, pp. 196–203.

- [MR89] Gary L. Miller and John H. Reif. "Parallel Tree Contraction—Part I: Fundamentals". In: *Advances in Computing Research* 5 (1989), pp. 47–72.
- [MT10] Robin A Moser and Gábor Tardos. "A constructive proof of the general Lovász local lemma". In: *Journal of the ACM (JACM)* 57.2 (2010), pp. 1–15.
- [Nao91] Moni Naor. "A lower bound on probabilistic algorithms for distributive ring coloring". In: *SIAM Journal on Discrete Mathematics* 4.3 (1991), pp. 409–412.
- [NS95] Moni Naor and Larry Stockmeyer. "What Can Be Computed Locally?" In: *SIAM Journal on Computing* 24.6 (1995), pp. 1259–1277.
- [NO08] Huy N. Nguyen and Krzysztof Onak. "Constant-Time Approximation Algorithms via Local Improvements". In: *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. 2008, pp. 327–336. doi: [10.1109/FOCS.2008.81](https://doi.org/10.1109/FOCS.2008.81).
- [Ona+12] Krzysztof Onak, Dana Ron, Michal Rosen, and Ronitt Rubinfeld. "A near-optimal sublinear-time algorithm for approximating the minimum vertex cover size". In: *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*. SIAM. 2012, pp. 1123–1131.
- [PR01] Alessandro Panconesi and Romeo Rizzi. "Some simple distributed algorithms for sparse networks". In: *Distributed computing* 14.2 (2001), pp. 97–100.
- [PS92] Alessandro Panconesi and Aravind Srinivasan. "Improved distributed algorithms for coloring and network decomposition problems". In: *Proc. 24th ACM Symp. on Theory of Computing (STOC)*. 1992, pp. 581–592.
- [PR07] Michal Parnas and Dana Ron. "Approximating the minimum vertex cover in sublinear time and a connection to distributed algorithms". In: *Theoretical Computer Science* 381.1-3 (2007), pp. 183–196.
- [Pel00] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [Pik21] Oleg Pikhurko. *Borel Combinatorics of Locally Finite Graphs*. 2021. arXiv: [2009.09113 \[math.CO\]](https://arxiv.org/abs/2009.09113).
- [QW22] Long Qian and Felix Weilacher. "Descriptive Combinatorics, Computable Combinatorics, and ASI Algorithms". In: *arXiv preprint arXiv:2206.08426* (2022).
- [RV16] Omer Reingold and Shai Vardi. "New techniques and tighter bounds for local computation algorithms". In: *Journal of Computer and System Sciences* 82.7 (2016), pp. 1180–1200.
- [RS20] Will Rosenbaum and Jukka Suomela. "Seeing Far vs. Seeing Wide: Volume Complexity of Local Graph Problems". In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. PODC '20. Virtual Event, Italy: Association for Computing Machinery, 2020, pp. 89–98. ISBN: 9781450375825.
- [Roz+22a] Václav Rozhoř, Michael Elkin, Christoph Grunau, and Bernhard Haeupler. "Deterministic low-diameter decompositions for weighted graphs and distributed and parallel applications". In: *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 1114–1121.
- [RG20] Václav Rozhoř and Mohsen Ghafari. "Polylogarithmic-time deterministic network decomposition and distributed derandomization". In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 2020, pp. 350–363.
- [Roz+22b] Václav Rozhoř, Christoph Grunau, Bernhard Haeupler, Goran Zuzic, and Jason Li. "Undirected $(1+\epsilon)$ -Shortest Paths via Minor-Aggregates: Near-Optimal Deterministic Parallel and Distributed Algorithms". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: Association for Computing Machinery, 2022, pp. 478–487. ISBN: 9781450392648.
- [RHG23] Václav Rozhoř, Bernhard Haeupler, and Christoph Grunau. "A Simple Deterministic Distributed Low-Diameter Clustering". In: *Symposium on Simplicity in Algorithms (SOSA)*. SIAM. 2023, pp. 166–174.
- [Rub+11] Ronitt Rubinfeld, Gil Tamir, Shai Vardi, and Ning Xie. "Fast local computation algorithms". In: *arXiv preprint arXiv:1104.1377* (2011).

- [RS15] Joel Rybicki and Jukka Suomela. "Exact bounds for distributed graph colouring". In: *International Colloquium on Structural Information and Communication Complexity*. Springer. 2015, pp. 46–60.
- [SS13] Stefan Schmid and Jukka Suomela. "Exploiting Locality in Distributed SDN Control". In: *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*. HotSDN '13. Hong Kong, China: Association for Computing Machinery, 2013, pp. 121–126. ISBN: 9781450321785.
- [SW10] Johannes Schneider and Roger Wattenhofer. "A New Technique for Distributed Symmetry Breaking". In: *Proceedings of the 29th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*. PODC '10. Zurich, Switzerland: Association for Computing Machinery, 2010, pp. 257–266. ISBN: 9781605588889.
- [Spi20] Yinon Spinka. "Finitely dependent processes are finitary". In: *Ann. Probab.* 48.4 (2020), pp. 2088–2117.
- [Suo10] Jukka Suomela. "Distributed Algorithms for Edge Dominating Sets". In: *Proceedings of the 29th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*. PODC '10. Zurich, Switzerland: Association for Computing Machinery, 2010, pp. 365–374. ISBN: 9781605588889.
- [Suo13] Jukka Suomela. "Survey of local algorithms". In: *ACM Computing Surveys* 45.2 (2013).
- [Suo20] Jukka Suomela. "Using Round Elimination to Understand Locality". In: *SIGACT News* 51.3 (Sept. 2020), pp. 63–81. ISSN: 0163-5700.
- [Suo23] Suomela, Jukka. *Open problems related to locality in distributed graph algorithms*. 2023. URL: <https://jukkasuomela.fi/open/>.
- [SV93] M. Szegedy and S. Vishwanathan. "Locality Based Graph Coloring". In: *Proc. 25th ACM Symp. on Theory of Computing (STOC)*. 1993, pp. 201–207.
- [YYI09] Yuichi Yoshida, Masaki Yamamoto, and Hiro Ito. "An improved constant-time approximation algorithm for maximum matchings". In: *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 2009, pp. 225–234.