

A Time Hierarchy Theorem for the LOCAL Model*

Yi-Jun Chang
University of Michigan
cyjjun@umich.edu

Seth Pettie
University of Michigan
pettie@umich.edu

Abstract

The celebrated *Time Hierarchy Theorem* for Turing machines states, informally, that more problems can be solved given more time. The extent to which a time hierarchy-type theorem holds in the classic distributed LOCAL model has been open for many years. In particular, it is consistent with previous results that all natural problems in the LOCAL model can be classified according to a small *constant* number of complexities, such as $O(1)$, $O(\log^* n)$, $O(\log n)$, $2^{O(\sqrt{\log n})}$, etc.

In this paper we establish the first time hierarchy theorem for the LOCAL model and prove that several *gaps* exist in the LOCAL time hierarchy. Our main results are as follows:

- We define an infinite set of simple coloring problems called *Hierarchical $2\frac{1}{2}$ -Coloring*. A correctly colored graph can be confirmed by simply checking the neighborhood of each vertex, so this problem fits into the class of *locally checkable labeling* (LCL) problems. However, the complexity of the k -level Hierarchical $2\frac{1}{2}$ -Coloring problem is $\Theta(n^{1/k})$, for $k \in \mathbb{Z}^+$. The upper and lower bounds hold for both general graphs and trees, and for both randomized and deterministic algorithms.
- Consider any LCL problem on *bounded degree trees*. We prove an automatic-speedup theorem that states that any *randomized* $n^{o(1)}$ -time algorithm solving the LCL can be transformed into a *deterministic* $O(\log n)$ -time algorithm. Together with a previous result [6], this establishes that on trees, there are no natural deterministic complexities in the ranges $\omega(\log^* n) - o(\log n)$ or $\omega(\log n) - n^{o(1)}$.
- We expose a gap in the *randomized* time hierarchy on general graphs. Roughly speaking, any randomized algorithm that solves an LCL problem in sublogarithmic time can be sped up to run in $O(T_{LLL})$ time, which is the complexity of the distributed Lovász local lemma problem, currently known to be $\Omega(\log \log n)$ and $O(\log n)$.

Finally, we revisit Naor and Stockmeyer's characterization of $O(1)$ -time LOCAL algorithms for LCL problems (as *order-invariant w.r.t. vertex IDs*) and calculate the complexity gaps that are directly implied by their proof. For n -rings we see a $\omega(1) - o(\log^* n)$ complexity gap, for $(\sqrt{n} \times \sqrt{n})$ -tori an $\omega(1) - o(\sqrt{\log^* n})$ gap, and for bounded degree trees and general graphs, an $\omega(1) - o(\log(\log^* n))$ complexity gap.

*Supported by NSF Grants CCF-1514383 and CCF-1637546.

1 Introduction

The goal of this paper is to understand the spectrum of natural problem complexities that can exist in the LOCAL model [29, 35] of distributed computation, and to quantify the value of randomness in this model. Whereas the time hierarchy of Turing machines is known¹ to be very “smooth”, recent work [6, 5] has exhibited strange “gaps” in the LOCAL complexity hierarchy. Indeed, prior to this work it was not even known if the LOCAL model could support more than a small *constant* number of problem complexities. Before surveying prior work in this area, let us formally define the deterministic and randomized variants of the LOCAL model, and the class of *locally checkable labeling* (LCL) problems, which are intuitively those graph problems that can be computed locally in *nondeterministic constant time*.

In both the DetLOCAL and RandLOCAL models the input graph $G = (V, E)$ and communications network are identical. Each vertex hosts a processor and all vertices run the same algorithm. Each edge supports communication in both directions. The computation proceeds in synchronized *rounds*. In a round, each processor performs some computation and sends a message along each incident edge, which is delivered before the beginning of the next round. Each vertex v is initially aware of its degree $\deg(v)$, a port numbering mapping its incident edges to $\{1, \dots, \deg(v)\}$, certain global parameters such as $n \stackrel{\text{def}}{=} |V|$, $\Delta \stackrel{\text{def}}{=} \max_{v \in V} \deg(v)$, and possibly other information.² The only measure of efficiency is the number of rounds. All local computation is free and the size of messages is unbounded. *Henceforth “time” refers to the number of rounds*. The differences between DetLOCAL and RandLOCAL are as follows.

DetLOCAL: In order to avoid trivial impossibilities, all vertices are assumed to hold unique $\Theta(\log n)$ -bit IDs. Except for the information about $\deg(v)$, $\text{ID}(v)$, and the port numbering, the initial state of v is identical to every other vertex. The algorithm executed at each vertex is deterministic.

RandLOCAL: In this model each vertex may locally generate an unbounded number of independent truly random bits. (There are no globally shared random bits.) Except for the information about $\deg(v)$ and its port numbering, the initial state of v is identical to every other vertex. Algorithms in this model operate for a specified number of rounds and have some probability of *failure*, the definition of which is problem specific. We set the maximum tolerable global probability of failure to be $1/n$.

Clearly RandLOCAL algorithms can generate distinct IDs (w.h.p.) if desired. Observe that the role of “ n ” is different in the two LOCAL models: in DetLOCAL it affects the ID length whereas in RandLOCAL it affects the failure probability.

LCL Problems. Naor and Stockmeyer [33] introduced *locally checkable labelings* to formalize a large class of natural graph problems. Fix a class \mathcal{G} of possible input graphs and let Δ be the maximum degree in any such graph. Formally, an LCL problem \mathcal{P} for \mathcal{G} has a radius r , constant size input and output alphabets $\Sigma_{\text{in}}, \Sigma_{\text{out}}$, and a set \mathcal{C} of acceptable configurations.³ Each $C \in \mathcal{C}$ is a graph centered at a specific vertex, in which each vertex has a degree, a port numbering, and two labels from Σ_{in} and Σ_{out} . Given the input graph $G(V, E, \phi_{\text{in}})$ where $\phi_{\text{in}} : V(G) \rightarrow \Sigma_{\text{in}}$, an acceptable output is any function $\phi_{\text{out}} : V(G) \rightarrow \Sigma_{\text{out}}$ such that for each $v \in V(G)$, the subgraph induced by $N^r(v)$ (denoting the r -neighborhood of v together with information stored there: vertex degrees, port numberings, input labels, and output labels) is isomorphic to a member of \mathcal{C} . An LCL can be described explicitly by enumerating a finite number of acceptable configurations. LCLs can be generalized to graph classes with unbounded degrees.

Many natural symmetry breaking problems can be expressed as LCLs, such as MIS, maximal matching, (α, β) -ruling sets, $(\Delta + 1)$ -vertex coloring, and sinkless orientation.

1.1 The Complexity Landscape of LOCAL

The complexity landscape for LCL problems is defined by “natural” complexities (sharp lower and upper bounds for specific LCL problems) and provably empty *gaps* in the complexity spectrum. We now have

¹For *any* time-constructible function $T(n)$, there is a problem solvable in $O(T(n))$ but not $o(T(n))$ time [21, 11].

²The assumption that global parameters are common knowledge can sometimes be removed; see Korman, Sereni, and Viennot [26].

³All of these parameters may depend on Δ .

an almost perfect understanding of the complexity landscape for two simple topologies: n -rings [8, 29, 32, 33, 6] and $(\sqrt{n} \times \sqrt{n})$ -tori [33, 6, 5]. See Figure 1, Top and Middle. On the n -ring, the only possible problem complexities are $O(1)$, $\Theta(\log^* n)$ (e.g., 3-coloring), and $\Theta(n)$ (e.g., 2-coloring, if bipartite). The gaps between these three complexities are obtained by *automatic speedup theorems*. Naor and Stockmeyer’s [33] characterization of $O(1)$ -time LCL algorithms actually implies that any $o(\log^* n)$ -time algorithm on the n -ring can be transformed to run in $O(1)$ time; see Appendix A. Chang, Kopelowitz, and Pettie [6] showed that any $o(n)$ -time RandLOCAL algorithm can be made to run in $O(\log^* n)$ time in DetLOCAL.

The situation with $(\sqrt{n} \times \sqrt{n})$ -tori is almost identical [5]: every known LCL has complexity $O(1)$, $\Theta(\log^* n)$ (e.g., 4-coloring), or $\Theta(\sqrt{n})$ (e.g., 3-coloring). Whereas the gap implied by [33] is $\omega(1) - o(\log^* n)$ on the n -ring, it is $\omega(1) - o(\sqrt{\log^* n})$ on the $(\sqrt{n} \times \sqrt{n})$ -torus; see Appendix A.⁴

Whereas randomness is known not to help in n -rings [33, 6], it is an open question on tori [5]. Whereas the classification question is decidable on n -rings (whether an LCL is $O(\log^* n)$ or $\Omega(n)$, for example) this question is *undecidable* on $(\sqrt{n} \times \sqrt{n})$ -tori [33, 5].

The gap theorems of Chang et al. [6] show that no LCL problem on general graphs has DetLOCAL complexity in the range $\omega(\log^* n) - o(\log_\Delta n)$, nor RandLOCAL complexity in the range $\omega(\log^* n) - o(\log_\Delta \log n)$. Some problems exhibit an exponential separation ($O(\log_\Delta \log n)$ vs. $\Omega(\log_\Delta n)$) between their RandLOCAL and DetLOCAL complexities, such as Δ -coloring degree- Δ trees [4, 6] and sinkless orientation [4, 14]. More generally, Chang et al. [6] proved that the RandLOCAL complexity of *any* LCL problem on graphs of size n is, holding Δ fixed, at least its deterministic complexity on instances of size $\sqrt{\log n}$. Thus, on the class of degree $\Delta = O(1)$ graphs there were only five known natural complexities: $O(1)$, $\Theta(\log^* n)$, randomized $\Theta(\log \log n)$, $\Theta(\log n)$, and $\Theta(n)$. For non-constant Δ , the RandLOCAL lower bounds of Kuhn, Moscibroda, and Wattenhofer [27] imply $\Omega(\min\{\frac{\log \Delta}{\log \log \Delta}, \sqrt{\frac{\log n}{\log \log n}}\})$ lower bounds on $O(1)$ -approximate vertex cover, MIS, and maximal matching. This $\Omega(\log \Delta / \log \log \Delta)$ lower bound is only known to be tight for $O(1)$ -approximate vertex cover [2]; the best maximal matching [3] and MIS [12] algorithms’ dependence on Δ is $\Omega(\log \Delta)$. The $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ lower bound is not known to be tight for any problem, but is almost tight for maximal matching on bounded arboricity graphs [3], e.g. trees or planar graphs.

New Results. In this paper we study the LOCAL complexity landscape on more general topologies: bounded degree trees and general graphs; see Figure 1, Bottom. We establish a new complexity gap for trees, a *potential* complexity gap for general graphs (which depends on the distributed complexity of the constructive Lovász local lemma), and a new infinite hierarchy of coloring problems with polynomial time complexities. In more detail,

- We prove that on the class of degree bounded trees, no LCL has complexity in the range $\omega(\log n) - n^{o(1)}$. Specifically, any $n^{o(1)}$ -time RandLOCAL algorithm can be converted to an $O(\log n)$ -time DetLOCAL algorithm. Moreover, given a description of an LCL problem \mathcal{P} , it is *decidable* whether the RandLOCAL complexity of \mathcal{P} is $n^{\Omega(1)}$ or the DetLOCAL complexity of \mathcal{P} is $O(\log n)$.

It turns out that this gap is maximal: we cannot extend it lower than $\omega(\log n)$ [6], nor higher than $n^{o(1)}$, as we show below.

- We define an infinite class of LCL problems called *Hierarchical $2\frac{1}{2}$ -Coloring*. We prove that k -level Hierarchical $2\frac{1}{2}$ -Coloring has complexity $\Theta(n^{1/k})$. The upper bound holds in DetLOCAL on general graphs, and the lower bound holds even on degree-3 trees in RandLOCAL. Thus, in contrast to rings and tori, trees and general graphs support an *infinite* number of natural problem complexities.
- Suppose we have a RandLOCAL algorithm for general graphs running in $C(\Delta) + o(\log_\Delta n)$ time. We can transform this algorithm to run in $O(C(\Delta) \cdot T_{LLL})$ time, where T_{LLL} is the complexity of a weak (i.e., “easy”) version of the constructive Lovász local lemma. At present, T_{LLL} is known to be $\Omega(\log \log n)$ and $O(\log n)$ on bounded degree graphs. If it turns out that T_{LLL} is sublogarithmic, this establishes a new RandLOCAL complexity gap between $\omega(T_{LLL})$ and $o(\log n)$ on bounded degree graphs.

⁴J. Suomela (personal communication, 2017) has a proof that there is an $\omega(1) - o(\log^* n)$ complexity gap for tori, at least for LCLs that do not use port numberings or input labels. The issues that arise with port numbering and input labels can be very subtle.




Figure 1: **Top:** the complexity landscape for LCL problems on the n -ring. **Middle:** the complexity landscape for LCL problems on the $(\sqrt{n} \times \sqrt{n})$ -torus. Refer to [33, 6, 5] and Appendix A for proofs of the complexity gaps (‘X’) on rings and tori. **Bottom:** the complexity landscape for LCL problems on bounded degree trees. The $\omega(\log^* n) - o(\log n)$ DetLOCAL gap and $\omega(\log^* n) - o(\log \log n)$ RandLOCAL gap are due to [6]. The $\omega(\text{LLL}) - o(\log n)$ and $\omega(\log n) - n^{o(1)}$ gaps are new. Refer to Appendix A for the $\omega(1) - o(\log(\log^* n))$ gap. It is unknown whether there are $\omega(n^{1/(k+1)}) - o(n^{1/k})$ gaps. With the exception of the $\omega(\log n) - n^{o(1)}$ gap, this is exactly the known complexity landscape for general bounded degree graphs as well.

Finally, it seems to be folklore that Naor and Stockmeyer’s work [33] implies *some kind of* complexity gap, which has been cited as $\omega(1) - o(\log^* n)$ [5, p. 2]. However, to our knowledge, no proof of this complexity gap has been published. We show how Naor and Stockmeyer’s approach implies complexity gaps that depend on the graph topology:

- $\omega(1) - o(\log^* n)$ on rings.
- $\omega(1) - o(\sqrt{\log^* n})$ on tori.⁵
- $\omega(1) - o(\log(\log^* n))$ on bounded degree trees and general graphs.

These gaps apply to the general class of LCL problems defined in this paper, in which vertices initially hold an input label and possible port numbering.⁶ J. Suomela (personal communication) gave a proof that the $\omega(1) - o(\log^* n)$ gap exists in tori as well, for LCL problems without input labels or port numbering.

Commentary. Our $\omega(\log n) - n^{o(1)}$ complexity gap for trees is interesting from both a technical and greater philosophical perspective, due to the fact that many natural problems have been “stuck” at $n^{o(1)}$ complexities for decades. Any DetLOCAL algorithm that relies on network decompositions [34] currently takes $2^{O(\sqrt{\log n})}$ time. If our automatic speedup theorem could be extended to the class of all graphs, this would immediately yield $O(\log n)$ -time algorithms for MIS, $(\Delta + 1)$ -coloring, and many other LCLs.

All the existing automatic speedup theorems are quite different in terms of proof techniques. Naor and Stockmeyer’s approach is based on Ramsey theory. The speedup theorems of [6, 5] use the fact that $o(\log_\Delta n)$ algorithms on general graphs⁷ cannot “see” the whole graph, and can therefore be efficiently tricked into

⁵Under certain requirements, this gap can be widened to $\omega(1) - o(\log^* n)$. See Appendix A.

⁶Port numberings are needed to represent ‘edge labeling’ problems (like maximal matching, edge coloring, and sinkless orientation) unambiguously as vertex labelings. They are not needed for native ‘vertex labeling’ problems like $(\Delta + 1)$ -coloring or MIS.

⁷(and $o(n)$ algorithms on n -rings and $o(\sqrt{n})$ algorithms on $(\sqrt{n} \times \sqrt{n})$ -tori)

thinking the graph has constant size. Our $n^{o(1)} \rightarrow O(\log n)$ speedup theorem introduces an entirely new set of techniques based on classic automata theory. We show that any LCL problem gives rise to a regular language that represents partial labelings of the tree that can be consistently extended to total labelings. By applying the pumping lemma for regular languages, we can “pump” the input tree into a much larger tree that behaves similar to the original tree. The advantage of creating a larger imaginary tree is that each vertex can (mentally) simulate the behavior of an $n^{o(1)}$ -time algorithm on the imaginary tree, merely by inspecting its $O(\log n)$ -neighborhood in the *actual* tree. Moreover, because the pumping operation preserves properties of the original tree, a labeling of the imaginary tree can be efficiently converted to a labeling of the original tree.

1.2 Related Results

There are several LOCAL lower bounds for natural problems that do not quite fit in the LCL framework. Göös, Hirvonen, and Suomela [15] proved a sharp $\Omega(\Delta)$ lower bound for fractional maximal matching and Göös and Suomela proved $\Omega(\log n)$ lower bounds on $(1 + \delta)$ -approximating the minimum vertex cover, $\delta > 0$, even on degree-3 graphs. See [28, 23] for lower bounds on coloring problems that apply to constrained algorithms or a constrained version of the LOCAL model.

In recent years there have been efforts to develop a ‘complexity theory’ of locality. The gap theorems of [33, 6, 5] have already been discussed. Suomela surveys [36] the class of problems that can be computed with $O(1)$ time. Fraigniaud et al. [10] defined a distributed model for locally deciding graph properties; see [9] for a survey of variants of the local distributed decision model. Göös and Suomela [16] considered the *proof* complexity (measured in terms of bits-per-vertex label) of locally verifying graph properties. Very recently, Ghaffari, Kuhn, and Maus [13] defined the SLOCAL model (sequential LOCAL) and exhibited several complete problems for this model, inasmuch as a $\text{polylog}(n)$ -time DetLOCAL algorithm for any complete problem implies a $\text{polylog}(n)$ DetLOCAL algorithm for every $\text{polylog}(n)$ -time problem in SLOCAL.⁸

1.3 Organization

In Section 2 we introduce Hierarchical $2\frac{1}{2}$ -Coloring and prove that the k -level variant of this problem has complexity $\Theta(n^{1/k})$. In Section 3 we prove the $n^{o(1)} \rightarrow O(\log n)$ speedup theorem for bounded degree trees. In Section 4 we discuss the constructive Lovász local lemma and prove the $o(\log_{\Delta} n) \rightarrow T_{LLL}$ randomized speedup theorem. In Section 5 we discuss open problems and outstanding conjectures. Appendix A reviews Naor and Stockmeyer’s characterization of $O(1)$ -time LCL algorithms, using Ramsey theory, and explains how it implies gaps in the complexity hierarchy that depend on graph topology.

2 An Infinitude of Complexities: Hierarchical $2\frac{1}{2}$ -Coloring

In this section we give an infinite sequence $(\mathcal{P}_k)_{k \in \mathbb{Z}^+}$ of LCL problems, where the complexity of \mathcal{P}_k is precisely $\Theta(n^{1/k})$.⁹ The upper bound holds on general graphs in DetLOCAL and the lower bound holds in RandLOCAL, even on degree-3 trees. Informally, the task of \mathcal{P}_k is to 2-color (with $\{\varnothing, \mathfrak{c}\}$) certain specific subgraphs of the input graph. Some vertices are *exempt* from being colored (in which case they are labeled \mathfrak{v}), and in addition, it is possible to decline to 2-color certain subgraphs, by labeling them \mathfrak{q} .

There are no input labels. The output label set is $\Sigma_{\text{out}} = \{\varnothing, \mathfrak{c}, \mathfrak{q}, \mathfrak{v}\}$.¹⁰ The problem \mathcal{P}_k is an LCL defined by the following rules.

Levels Subsequent rules depend on the *levels* of vertices. Let V_i , $i \in \{1, \dots, k + 1\}$, be the set of vertices on

⁸The class of $O(1)$ -time SLOCAL algorithms is, roughly speaking, those graph labelings that can be computed sequentially and greedily, by a truly local algorithm. This class is a *strict* subset of LCLs.

⁹Brandt et al. [5, Appendix A.3] described an LCL that has complexity $\Theta(\sqrt{n})$ on general graphs, but not trees. It may be possible to generalize their LCL to any complexity of the form $\Theta(n^{1/k})$.

¹⁰Venus, Mars, Mercury, Saturn.

level i , defined as follows.

$$\begin{aligned}
G_1 &= G \\
G_i &= G_{i-1} - V_{i-1}, && \text{for } i \in [2, k+1] \\
V_i &= \{v \in V(G_i) \mid \deg_{G_i}(v) \leq 2\}, && \text{for } i \in [1, k] \\
V_{k+1} &= V(G_{k+1}) && \text{(the remaining vertices)}
\end{aligned}$$

Remember that vertices know their degrees, so a vertex in V_1 deduces this with 0 rounds of communication. In general the level of v can be calculated from information in $N^k(v)$.

Exemption A vertex labeled \mathfrak{v} is called *exempt*. No V_1 vertex is labeled \mathfrak{v} ; all V_{k+1} vertices are labeled \mathfrak{v} . Any V_i vertex is labeled \mathfrak{v} iff it is adjacent to a lower level vertex labeled $\mathfrak{q}, \mathfrak{o}$, or \mathfrak{v} . Define $D_i \subseteq V_i$ to be the set of level i exempt vertices.

Two-Coloring Vertices not covered by the exemption rule are labeled one of $\mathfrak{q}, \mathfrak{o}, \mathfrak{q}$.

- Any vertex in V_i , $i \in [1, k]$, labeled \mathfrak{q} has no neighbor in V_i labeled \mathfrak{q} or \mathfrak{q} .
- Any vertex in V_i , $i \in [1, k]$, labeled \mathfrak{o} has no neighbor in V_i labeled \mathfrak{o} or \mathfrak{q} .
- Any vertex in $V_k - D_k$ with exactly 0 or 1 neighbors in $V_k - D_k$ must be labeled \mathfrak{q} or \mathfrak{o} .

Commentary. The Level rule states that the graph induced by V_i consists of paths and cycles. The Two-Coloring rule implies that each component of non-exempt vertices in the graph induced by $V_i - D_i$ must either (a) be labeled uniformly by \mathfrak{q} or (b) be properly 2-colored by $\{\mathfrak{q}, \mathfrak{o}\}$. Every *path* in $V_k - D_k$ must be properly 2-colored, but *cycles* in $V_k - D_k$ are allowed to be labeled uniformly by \mathfrak{q} . This last provision is necessary to ensure that *every* graph can be labeled according to \mathcal{P}_k since there is no guarantee that cycles in $V_k - D_k$ are bipartite.

Remark 1. As stated \mathcal{P}_k is an LCL with an alphabet size of 4 and a radius k , since the coloring rules refer to levels, which can be deduced by looking up to radius k . On the other hand, we can also represent \mathcal{P}_k as an LCL with radius 1 and alphabet size $4k$ by including a vertex's level in its output label. A correct level assignment can be verified within radius 1. For example, level 1 vertices are those with degree at most 2, and a vertex is labeled $i \in [2, k]$ iff all but at most 2 neighbors have levels less than i .

Theorem 1. The DetLOCAL complexity of \mathcal{P}_k on general graphs is $O(n^{1/k})$.

Proof. The algorithm fixes the labeling of V_1, \dots, V_k, V_{k+1} in order, according to the following steps. Assume that all vertices in V_1, \dots, V_{i-1} have already been labeled.

- Compute D_i according to the Exemption rule. (E.g., $D_1 = \emptyset$, $D_{k+1} = V_{k+1}$.)
- Each path in the subgraph induced by $V_i - D_i$ calculates its length. If it contains at most $\lceil 2n^{1/k} \rceil$ vertices, it properly 2-colors itself with $\{\mathfrak{q}, \mathfrak{o}\}$; longer paths and cycles in $V_i - D_i$ label themselves uniformly by \mathfrak{q} .

This algorithm correctly solves \mathcal{P}_k provided that it never labels a path in $V_k - D_k$ with \mathfrak{q} . Let U_i be the subgraph induced by those vertices in $V_1 \cup \dots \cup V_i$ labeled \mathfrak{q} . Consider a connected component C in U_i whose V_i -vertices are arranged in a path (not a cycle). We argue by induction that C has at least $2n^{i/k}$ vertices. This is clearly true in the base case $i = 1$: if a path component of U_1 were colored \mathfrak{q} , it must have more than $\lceil 2n^{1/k} \rceil$ vertices. Now assume the claim is true for $i - 1$ and consider a component C of U_i . If the V_i -vertices in C form a path, it must have length greater than $2n^{1/k}$. Each vertex in that path must be adjacent to an endpoint of a V_{i-1} path. Since V_{i-1} paths have two endpoints, the V_i path is adjacent to at least $\lceil 2n^{1/k} \rceil / 2 \geq n^{1/k}$ components in U_{i-1} , each of which has size at least $2n^{(i-1)/k}$, by the inductive hypothesis. Thus, the size of C is at least $n^{1/k} \cdot 2n^{(i-1)/k} + 2n^{1/k} > 2n^{i/k}$. Because there are at most n vertices in the graph, it is impossible for V_k vertices arranged in a path to be colored \mathfrak{q} . \square

Theorem 2. The RandLOCAL complexity of \mathcal{P}_k on trees with maximum degree $\Delta = 3$ is $\Omega(n^{1/k})$.

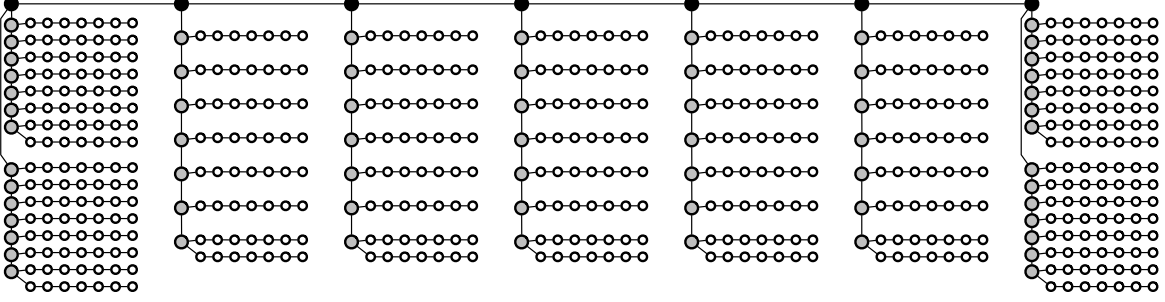


Figure 2: The graph H_k with parameters $k = 3, x = 7$. White vertices are in V_1 , gray in V_2 , and black in V_3 . $V_4 = V_{k+1}$ is empty.

Proof. Fix an integer parameter x and define a sequence of graphs $(H_i)_{1 \leq i \leq k}$ as follows. Each H_i has a *head* and a *tail*.

- H_1 is a path (or *backbone*) of length x . One end of the path is the head and the other end the tail.
- To construct H_i , $i \in [2, k-1]$, begin with a *backbone* path (v_1, v_2, \dots, v_x) , with head v_1 and tail v_x . Form $x+1$ copies $(H_{i-1}^{(j)})_{1 \leq j \leq x+1}$ of H_{i-1} , where $v^{(j)}$ is the head of $H_{i-1}^{(j)}$. Connect $v^{(j)}$ to v_j by an edge, for $j \in [1, x]$, and also connect $v^{(x+1)}$ to v_x by an edge.
- H_k is constructed exactly as above, except that we generate $x+2$ copies of H_{k-1} and connect the heads of two copies of H_{k-1} to both v_1 and v_x . See Figure 2 for an example with $k = 3$.

Let us make several observations about the construction of H_k . First, it is a tree with maximum degree 3. Second, when decomposing $V(H_k)$ into levels $(V_1, \dots, V_k, V_{k+1})$, V_i is precisely the union of the backbones in all copies of H_i , and $V_{k+1} = \emptyset$. Third, the number of vertices in H_k is $\Theta(x^k)$, so a $o(n^{1/k})$ algorithm for \mathcal{P}_k must run in $o(x)$ time on H_k .

Consider a RandLOCAL algorithm \mathcal{A} solving \mathcal{P}_k on H_k within $t < x/5 - O(1)$ time, that fails with probability p_{fail} . If \mathcal{A} is a good algorithm then $p_{\text{fail}} \leq 1/|V(H_k)|$. However, we will now show that p_{fail} is constant, independent of $|V(H_k)|$.

Define \mathcal{E}_i to be the event that $D_i \neq \emptyset$ and $p_i = \Pr(\mathcal{E}_i)$. By an induction from $i = 2$ to k , we prove that $p_i \leq 2(i-1) \cdot p_{\text{fail}}$.

Base case. We first prove that

$$\Pr(H_k \text{ is not correctly colored according to } \mathcal{P}_k \mid \mathcal{E}_2) \geq 1/2.$$

Conditioning on \mathcal{E}_2 means that $D_2 \neq \emptyset$. Fix any $v \in D_2$ and let P be a copy of H_1 (a path) adjacent to v . In order for $v \in D_2$, it must be that P is properly 2-colored with $\{\varphi, \sigma\}$. Since $t < x/5 - O(1)$, there exist two vertices u and u' in P such that

1. $N^t(u)$, $N^t(u')$, and $N^t(v)$ are disjoint sets,
2. the subgraphs induced by $N^t(u)$ and $N^t(u')$ are isomorphic, and
3. the distance between u and u' is odd.

Let p_φ and p_σ be the probabilities that u/u' is labeled φ and σ , respectively. A proper 2-coloring of P assigns u and u' different colors, and that occurs with probability $2p_\varphi p_\sigma \leq 2p_\varphi(1-p_\varphi) \leq 1/2$. Moreover, this holds independent of the random bits generated by vertices in $N^t(v)$. The algorithm fails unless u, u' have different colors, thus $p_{\text{fail}} \geq p_2/2$, and hence $p_2 \leq 2 \cdot p_{\text{fail}}$.

Inductive Step. Let $3 \leq i \leq k$. The inductive hypothesis states that $p_{i-1} \leq 2(i-2) \cdot p_{\text{fail}}$. By a proof similar to the base case, we have that:

$$\Pr(H_k \text{ is not correctly colored according to } \mathcal{P}_k \mid \mathcal{E}_i \setminus \mathcal{E}_{i-1}) \geq 1/2.$$

We are conditioning on $\mathcal{E}_i \setminus \mathcal{E}_{i-1}$. If this event is empty, then $p_i \leq p_{i-1} \leq 2(i-2) \cdot p_{\text{fail}}$ and the induction is complete. On the other hand, if $\mathcal{E}_i \setminus \mathcal{E}_{i-1}$ holds then there is some $v \in D_i$ adjacent to a copy of H_{i-1} with backbone path P , where $P \cap D_{i-1} = \emptyset$. In other words, if H_k is colored according to \mathcal{P}_k then P must be properly 2-colored with $\{\varphi, \mathfrak{O}\}$. The argument above shows this occurs with probability at least $1/2$. Thus,

$$p_{\text{fail}} = \Pr(H_k \text{ is incorrectly colored}) \geq \Pr(\mathcal{E}_i \setminus \mathcal{E}_{i-1})/2 \geq (p_i - p_{i-1})/2,$$

or $p_i \leq 2p_{\text{fail}} + p_{i-1} \leq 2(i-1)p_{\text{fail}}$, completing the induction.

Finally, let P be the path induced by vertices in V_k . The probability that \mathcal{E}_k holds ($P \cap D_k \neq \emptyset$) is $p_k \leq 2(k-1) \cdot p_{\text{fail}}$. On the other hand, $\Pr(H_k \text{ not colored correctly} \mid \overline{\mathcal{E}_k}) \geq 1/2$ by the argument above, hence $p_{\text{fail}} \geq (1 - p_k)/2$, or $p_k \geq 1 - 2p_{\text{fail}}$. Combining the upper and lower bounds on p_k we conclude that $p_{\text{fail}} \geq (2k)^{-1}$ is constant, independent of $|V(H_k)|$. Thus, algorithm \mathcal{A} cannot succeed with high probability. \square


3 A Complexity Gap on Bounded Degree Trees

In this section we prove an $n^{o(1)} \rightarrow O(\log n)$ speedup theorem for LCL problems on bounded degree trees. The progression of definitions and lemmas in Sections 3.2–3.13 is logical, but does not reveal the high level structure of the proof. Section 3.1 gives an informal tour of the proof and its key ideas. Throughout, \mathcal{P} is a radius- r LCL and \mathcal{A} is an $n^{o(1)}$ -time algorithm for \mathcal{P} on bounded degree trees.

3.1 A Tour of the Proof

Consider this simple way to decompose a tree in $O(\log n)$ time, inspired by Miller and Reif [30]. Iteratively remove paths of degree-2 vertices (*compress*) and vertices with degree 0 or 1 (*rake*). Vertices removed in iteration i are at *level* i . If $O(\log n)$ *rakes* alone suffice to decompose a tree then it has $O(\log n)$ diameter and any LCL can be solved in $O(\log n)$ time on such a graph. Thus, we mainly have to worry about the situation where *compress* removes very long ($\omega(1)$ -length) paths.


The first observation is that it is easy to split up long degree-2 paths of level- i vertices into constant length paths, by artificially promoting a well-spaced subset of level- i vertices to level $i+1$. Thus, we have a situation that looks like this: level- i vertices are arranged in an $O(1)$ -length path, each the root of a (colored) subtree of level- $(< i)$ vertices that were removed in previous rake/compress steps, and bookended by level- $(> i)$ (black) vertices. Call the subgraph between the bookends H .




In our approach it is the level- $(i+1)$ vertices that are in charge of coordinating the labeling of level- $(\leq i)$ vertices in their purview. In this diagram, H is in the purview of both black bookends. We only have one tool available for computing a labeling of this subgraph: an $n^{o(1)}$ -time RandLOCAL algorithm \mathcal{A} that works w.h.p. What would happen if we *simulated* \mathcal{A} on H ? The simulation would fail catastrophically of course, since it needs to look up to an $n^{o(1)}$ radius, to parts of the graph far outside of H .

Note that the colored subtrees are unbounded in terms of size and depth. Nonetheless, they fall into a *constant* number of equivalence classes in the following sense. The *class* of a rooted tree is the set of

all labelings of the r -neighborhood of its root that can be extended to total labelings of the tree that are consistent with \mathcal{P} .



In other words, the large and complex graph H can be succinctly encoded as a simple class vector (c_1, c_2, \dots, c_x) , where c_j is the class of the j th colored tree. Consider the set of all labelings of H that are consistent with \mathcal{P} . This set can also be succinctly represented by listing the labelings of the r -neighborhoods of the bookends that can be extended to all of H , while respecting \mathcal{P} . The set of these partial labelings defines the *type* of H . We show that H 's type can be computed by a finite automaton that reads the class vector (c_1, \dots, c_x) one character at a time. By the pigeonhole principle, if x is sufficiently large then the automaton loops, meaning that (c_1, \dots, c_x) can be written as $x \circ y \circ z$, which has the same type as every $x \circ y^j \circ z$, for all $j \geq 1$. This *pumping lemma for trees* lets us dramatically expand the size of H without affecting its type, i.e., how it interacts with the outside world beyond the bookends.



This diagram illustrates the pumping lemma with a substring of $|y| = 3$ trees (rooted at gray vertices) repeated $j = 3$ times. Now let us reconsider the simulation of \mathcal{A} . If we first pump H to be long enough, and then simulate \mathcal{A} on the middle section of pumped- H , \mathcal{A} must, according to its $n^{o(1)}$ time bound, compute a labeling *without needing any information outside of pumped- H* , i.e., beyond the bookends. Thus, we can use \mathcal{A} to *pre-commit* to a labeling of a small (radius- r) subgraph of pumped- H . Given this pre-commitment, the left and right bookends no longer need to coordinate their activities: everything left (right) of the pre-committed zone is now in the purview of the left (right) bookend. Interestingly, these manipulations (tree surgery and pre-commitments) can be repeated for each i , yielding a hierarchy of *imaginary* trees such that a proper labeling at one level of the hierarchy implies a proper labeling at the previous level.

Roadmap. This short proof sketch has been simplified to the point that it is riddled with small inaccuracies. Nonetheless, it does accurately capture the difficulties, ideas, and techniques used in the actual proof. In Section 3.2 we formally define the notion of a partially labeled graph, i.e., one with certain vertices pre-committed to their output labels. Section 3.3 defines a surgical “cut-and-paste” operation on graphs. Section 3.4 defines a partition of the vertices of a subgraph H , which differentiates between vertices that “see” the outside graph, and those that see only H . Section 3.5 defines an equivalence relation on graphs that, intuitively, justifies surgically replacing a subgraph with an equivalent graph. Sections 3.6 and 3.7 explore properties of the equivalence relation. Section 3.8 introduces the pumping lemma for trees, and Section 3.9 defines a specialized Rake/Compress-style graph decomposition. Section 3.10 presents the operations **Extend** (which pumps a subtree) and **Label** (which labels a small subset of the vertices of a tree) in terms of a black-box *labeling function* f . Section 3.11 defines the set of all (partially labeled) trees that can be encountered, by considering the interplay between the graph decomposition, **Extend**, and **Label**. It is important that for each tree encountered, its partial labeling can be extended to a complete labeling consistent with \mathcal{P} ; whether this actually holds depends on the choice of black-box f . Section 3.13 shows how a *feasible labeling function* f can be extracted from any $n^{o(1)}$ -time algorithm \mathcal{A} and Section 3.12 shows that \mathcal{P} can be solved in $O(\log n)$.

time, given a feasible labeling function.

3.2 Partial Labeled Graphs

A *partially labeled graph* $\mathcal{G} = (G, \mathcal{L})$ is a graph G together with a function $\mathcal{L} : V(G) \rightarrow \Sigma_{\text{out}} \cup \{\perp\}$. The vertices in $\mathcal{L}^{-1}(\perp)$ are *unlabeled*. A *complete labeling* $\mathcal{L}' : V(G) \rightarrow \Sigma_{\text{out}}$ for \mathcal{G} is one that labels all vertices and is consistent with \mathcal{G} 's partial labeling, i.e., $\mathcal{L}'(v) = \mathcal{L}(v)$ whenever $\mathcal{L}(v) \neq \perp$. A *legal* labeling is a complete labeling that is *locally consistent* for all $v \in V(G)$, i.e., the labeled subgraph induced by $N^r(v)$ is consistent with the LCL \mathcal{P} . Here $N^r(v)$ is the set of all vertices within distance r of v .


All graph operations can be extended naturally to partially labeled graphs. For instance, a subgraph of a partially labeled graph $\mathcal{G} = (G, \mathcal{L})$ is a pair $\mathcal{H} = (H, \mathcal{L}')$ such that H is a subgraph of G , and \mathcal{L}' is \mathcal{L} restricted to the domain $V(H)$. With slight abuse of notation, we usually write $\mathcal{H} = (H, \mathcal{L})$.

3.3 Graph Surgery

Let $\mathcal{G} = (G, \mathcal{L})$ be a partially labeled graph, and let $\mathcal{H} = (H, \mathcal{L})$ be a subgraph of \mathcal{G} . The *poles* of \mathcal{H} are those vertices in $V(H)$ that are adjacent to some vertex in the outside graph $V(G) - V(H)$. We define an operation **Replace** that surgically removes \mathcal{H} and replaces it with some \mathcal{H}' .

Replace Let $S = (v_1, \dots, v_p)$ be a list of the poles of \mathcal{H} and let $S' = (v'_1, \dots, v'_p)$ be a designated set of poles in some partially labeled graph \mathcal{H}' . The partially labeled graph $\mathcal{G}' = \text{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$ is constructed as follows. Beginning with \mathcal{G} , replace \mathcal{H} with \mathcal{H}' , and replace any edge $\{u, v_i\}$, $u \in V(G) - V(H)$, with $\{u, v'_i\}$. If the poles S, S' are clear from context, we may also simply write $\mathcal{G}' = \text{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$. Writing $\mathcal{G}' = (G', \mathcal{L}')$ and $\mathcal{H}' = (H', \mathcal{L}')$, there is a natural 1-1 correspondence between the vertices in $V(G) - V(H)$ and $V(G') - V(H')$.

In the proof of our $n^{o(1)} \rightarrow O(\log n)$ speedup theorem we only consider unipolar and bipolar graphs ($p \in \{1, 2\}$) but for maximum generality we define everything w.r.t. graphs having $p \geq 1$ poles.



Given a legal labeling \mathcal{L}_\diamond of \mathcal{G} , we would like to know whether there is a legal labeling \mathcal{L}'_\diamond of \mathcal{G}' that agrees with \mathcal{L}_\diamond , i.e., $\mathcal{L}_\diamond(v) = \mathcal{L}'_\diamond(v')$ for each $v \in V(G) - V(H)$ and the corresponding $v' \in V(G') - V(H')$. Our goal is to define an equivalence relation $\tilde{\sim}$ on partially labeled graphs (with designated poles) so that the following is true: if $(\mathcal{H}, S) \tilde{\sim} (\mathcal{H}', S')$, then such a legal labeling \mathcal{L}'_\diamond must exist, regardless of the choice of \mathcal{G} and \mathcal{L}_\diamond . Observe that since \mathcal{P} has radius r , the interface between $V(H)$ (or $V(H')$) and the rest of the graph only occurs around the $O(r)$ -neighborhoods of the poles of \mathcal{H} (or \mathcal{H}'). This motivates us to define a certain partition of \mathcal{H} 's vertices that depends on its poles and r .

3.4 A Tripartition of the Vertices

Let $\mathcal{H} = (H, \mathcal{L})$ be a partially labeled graph with poles $S = (v_1, \dots, v_p)$. Define $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ to be a tripartition of $V(H)$, where $D_1 = \bigcup_{v \in S} N^{r-1}(v)$, $D_2 = \bigcup_{v \in D_1} N^r(v) - D_1$, and $D_3 = V(H) - (D_1 \cup D_2)$. See Figure 3 for an illustration.

Consider the partition $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ of a partially labeled graph $\mathcal{H} = (H, \mathcal{L})$. Let $\mathcal{L}_* : D_1 \cup D_2 \rightarrow \Sigma_{\text{out}}$ assign output labels to $D_1 \cup D_2$. We say that \mathcal{L}_* is *extendible* (to all of $V(H)$) if there




Figure 3: A partially labeled subgraph \mathcal{H} with poles $S = (s, t)$, embedded in a larger graph \mathcal{G} . In the partition $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$, D_1 is the set of vertices in $V(H)$ within radius $r - 1$ of S , D_2 are those within radius $2r - 1$ of S , excluding D_1 , and D_3 is the rest of $V(H)$. When \mathcal{H} is embedded in some larger graph \mathcal{G} , D_0 denotes the remaining vertices in $V(\mathcal{G}) - V(H)$.

exists a complete labeling \mathcal{L}_\diamond of H such that \mathcal{L}_\diamond agrees with \mathcal{L} where it is defined, agrees with \mathcal{L}_* on $D_1 \cup D_2$, and is locally consistent with \mathcal{P} on all vertices in $D_2 \cup D_3$.¹¹

3.5 An Equivalence Relation on Graphs

Consider two partially labeled graphs \mathcal{H} and \mathcal{H}' with poles $S = (v_1, \dots, v_p)$ and $S' = (v'_1, \dots, v'_p)$, respectively. Let $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ and $\xi(\mathcal{H}', S') = (D'_1, D'_2, D'_3)$. Define $\mathcal{Q} = (Q, \mathcal{L})$ and $\mathcal{Q}' = (Q', \mathcal{L}')$ as the subgraphs of \mathcal{H} and \mathcal{H}' induced by the vertices in $D_1 \cup D_2$ and $D'_1 \cup D'_2$, respectively.

The relation $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$ holds if and only if there is a 1-1 correspondence $\phi : (D_1 \cup D_2) \rightarrow (D'_1 \cup D'_2)$ meeting the following conditions.

Isomorphism. The two graphs Q and Q' are isomorphic under ϕ . Moreover, for each $v \in D_1 \cup D_2$ and its corresponding vertex $v' = \phi(v) \in D'_1 \cup D'_2$, (i) $\mathcal{L}(v) = \mathcal{L}'(v')$, (ii) if the underlying LCL problem has input labels, then the input labels of v and v' are the same, and (iii) v is the i th pole in S iff v' is the i th pole in S' .

Extendibility. Let \mathcal{L}_* be *any* assignment of output labels to vertices in $D_1 \cup D_2$ and let \mathcal{L}'_* be the corresponding labeling of $D'_1 \cup D'_2$ under ϕ . Then \mathcal{L}_* is extendible (to $V(H)$) if and only if \mathcal{L}'_* is extendible (to $V(H')$).

Notice that there could be many 1-1 correspondences between $D_1 \cup D_2$ and $D'_1 \cup D'_2$ that satisfy the isomorphism requirement, though only some subset may satisfy the extendibility requirement due to differences in the topology and partial labeling of D_3 and D'_3 . Any ϕ meeting both requirements is a *witness* of the relation $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$.

3.6 Properties of the Equivalence Relation

Let $\mathcal{G}' = \text{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$. Consider the partitions $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ and $\xi(\mathcal{H}', S') = (D'_1, D'_2, D'_3)$ and let $D_0 = V(\mathcal{G}) - V(H)$ and $D'_0 = V(\mathcal{G}') - V(H')$ be the remaining vertices in \mathcal{G} and \mathcal{G}' , respectively.

If $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$ then there exists a 1-1 correspondence $\phi : (D_0 \cup D_1 \cup D_2) \rightarrow (D'_0 \cup D'_1 \cup D'_2)$ such that (i) ϕ restricted to D_0 is the natural 1-1 correspondence between D_0 and D'_0 and (ii) ϕ restricted to $D_1 \cup D_2$ witnesses the relation $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$. Such a 1-1 correspondence ϕ is called *good*. We have the following lemma.

¹¹We are not concerned whether \mathcal{L}_\diamond is consistent with \mathcal{P} for vertices in D_1 . Ultimately, \mathcal{H} will be a subgraph of a larger graph \mathcal{G} . Since the r -neighborhoods of vertices in D_1 will intersect $V(\mathcal{G}) - V(H)$, the labeling of H does not provide enough information to tell if these vertices' r -neighborhoods will be consistent with \mathcal{P} . See Figure 3.

Lemma 1. Let $\mathcal{G}' = \text{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$. Consider the partitions $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ and $\xi(\mathcal{H}', S') = (D'_1, D'_2, D'_3)$ and let $D_0 = V(G) - V(H)$ and $D'_0 = V(G') - V(H')$. Suppose that $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$, so there is a good 1-1 correspondence $\phi : (D_0 \cup D_1 \cup D_2) \rightarrow (D'_0 \cup D'_1 \cup D'_2)$. Let \mathcal{L}_\diamond be a complete labeling of \mathcal{G} that is locally consistent for all vertices in $D_2 \cup D_3$. Then there exists a complete labeling \mathcal{L}'_\diamond of \mathcal{G}' such that the following conditions are met.

Condition 1. $\mathcal{L}_\diamond(v) = \mathcal{L}'_\diamond(v')$ for each $v \in D_0 \cup D_1 \cup D_2$ and its corresponding vertex $v' = \phi(v) \in D'_0 \cup D'_1 \cup D'_2$. Moreover, if \mathcal{L}_\diamond is locally consistent for v , then \mathcal{L}'_\diamond is locally consistent for v' .

Condition 2. \mathcal{L}'_\diamond is locally consistent for all vertices in $D'_2 \cup D'_3$.

Proof. We construct \mathcal{L}'_\diamond as follows. First of all, for each $v \in D_0 \cup D_1 \cup D_2$, fix $\mathcal{L}'_\diamond(\phi(v)) = \mathcal{L}_\diamond(v)$. It remains to show how to assign output labels to vertices in D'_3 to meet Conditions 1 and 2.

Let \mathcal{L}_* be \mathcal{L}_\diamond restricted to the domain $D_1 \cup D_2$. Similarly, let \mathcal{L}'_* be \mathcal{L}'_\diamond restricted to $D'_1 \cup D'_2$. Due to the fact that \mathcal{L}_\diamond is locally consistent for all vertices in $D_2 \cup D_3$, the labeling \mathcal{L}_* is extendible for \mathcal{H} . Since $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$, the labeling \mathcal{L}'_* is also extendible for \mathcal{H}' . Thus, we can set $\mathcal{L}'_\diamond(v')$ for all $v' \in D'_3$ in such a way that \mathcal{L}'_\diamond is locally consistent for all vertices in $D'_2 \cup D'_3$. Therefore, Condition 2 is met.

To see that (the second part of) Condition 1 is also met, observe that for $v \in D_0 \cup D_1$, $N^r(v) \subseteq D_0 \cup D_1 \cup D_2$. Therefore, if \mathcal{L}_\diamond is locally consistent for $v \in D_0 \cup D_1$, then \mathcal{L}'_\diamond is locally consistent for $\phi(v)$ since they have the same radius- r neighborhood view. Condition 2 already guarantees that \mathcal{L}'_\diamond is locally consistent for all $v' \in D'_2$.¹² \square

Theorem 3 provides a user-friendly corollary of Lemma 1, which does not mention the tripartition ξ .

Theorem 3. Let $\mathcal{G} = (G, \mathcal{L})$ and $\mathcal{H} = (H, \mathcal{L})$ be a subgraph \mathcal{G} . Suppose \mathcal{H}' is a graph for which $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$ and let $\mathcal{G}' = \text{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$. We write $\mathcal{G}' = (G', \mathcal{L}')$ and $\mathcal{H}' = (H', \mathcal{L}')$. Let \mathcal{L}_\diamond be a complete labeling of \mathcal{G} that is locally consistent for all vertices in H . Then there exists a complete labeling \mathcal{L}'_\diamond of \mathcal{G}' such that the following conditions are met.

- For each $v \in V(G) - V(H)$ and its corresponding $v' \in V(G') - V(H')$, we have $\mathcal{L}_\diamond(v) = \mathcal{L}'_\diamond(v')$. Moreover, if \mathcal{L}_\diamond is locally consistent for v , then \mathcal{L}'_\diamond is locally consistent for v' .
- \mathcal{L}'_\diamond is locally consistent for all vertices in H' .

Theorem 3 has several useful consequences. If \mathcal{L}_\diamond is a legal labeling of \mathcal{G} , then the output labeling \mathcal{L}'_\diamond of \mathcal{G}' guaranteed by Theorem 3 is also legal. Observe that setting $\mathcal{G} = \mathcal{H}$ in Theorem 3 implies $\mathcal{G}' = \mathcal{H}'$. Suppose that \mathcal{H} admits a legal labeling. For any (\mathcal{H}', S') such that $(\mathcal{H}', S') \stackrel{*}{\sim} (\mathcal{H}, S)$, the partially labeled graph \mathcal{H}' also admits a legal labeling. Thus, whether \mathcal{H} admits a legal labeling is determined by the equivalence class of (\mathcal{H}, S) (for any choice of S).

Roughly speaking, Theorem 4 shows that the equivalence class of (\mathcal{G}, X) is preserved after replacing a subgraph \mathcal{H} of \mathcal{G} by another partially labeled graph \mathcal{H}' such that $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$.

Theorem 4. Let $\mathcal{G} = (G, \mathcal{L})$ and let $\mathcal{H} = (H, \mathcal{L})$ be a subgraph \mathcal{G} . Suppose \mathcal{H}' is such that $(\mathcal{H}, S) \stackrel{*}{\sim} (\mathcal{H}', S')$ for some pole lists S, S' . Let $\mathcal{G}' = \text{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$ be a partially labeled graph. Designate a set $X \subseteq (V(G) - V(H)) \cup S$ as the poles of \mathcal{G} , listed in some order, and let X' be the corresponding list of vertices in \mathcal{G}' . It follows that $(\mathcal{G}, X) \stackrel{*}{\sim} (\mathcal{G}', X')$.

Proof. Consider the four partitions $\xi(\mathcal{H}, S) = (B_1, B_2, B_3)$, $\xi(\mathcal{H}', S') = (B'_1, B'_2, B'_3)$, $\xi(\mathcal{G}, X) = (D_1, D_2, D_3)$, and $\xi(\mathcal{G}', X') = (D'_1, D'_2, D'_3)$. We write $B_0 = V(G) - V(H)$ and $B'_0 = V(G') - V(H')$. Let ϕ be any good 1-1 correspondence from $B_0 \cup B_1 \cup B_2$ to $B'_0 \cup B'_1 \cup B'_2$. Because $X \subseteq B_0 \cup S$, we have $D_1 \cup D_2 \subseteq B_0 \cup B_1 \cup B_2$ and $D'_1 \cup D'_2 \subseteq B'_0 \cup B'_1 \cup B'_2$. To show that $(\mathcal{G}, X) \stackrel{*}{\sim} (\mathcal{G}', X')$, it suffices to prove that ϕ (restricted to the domain $D_1 \cup D_2$) is a witness to the relation $(\mathcal{G}, X) \stackrel{*}{\sim} (\mathcal{G}', X')$.

Let $\mathcal{L}_* : (D_1 \cup D_2) \rightarrow \Sigma_{\text{out}}$ and \mathcal{L}'_* be the corresponding labeling of $D'_1 \cup D'_2$. All we need to do is show that \mathcal{L}_* is extendible to $V(G)$ if and only if \mathcal{L}'_* is extendible for $V(G')$. Since we can also write

¹²It is this lemma that motivates our definition of the tripartition $\xi(\mathcal{H}, S)$. It is not clear how an analogue of Lemma 1 could be proved using the seemingly more natural bipartition, i.e., by collapsing D_1, D_2 into one set.

$\mathcal{G} = \text{Replace}(\mathcal{G}', (\mathcal{H}', S'), (\mathcal{H}, S))$, it suffices to show just one direction, i.e., if \mathcal{L}_* is extendible then \mathcal{L}'_* is extendible.

Suppose that \mathcal{L}_* is extendible. Then there exists an output labeling \mathcal{L}_\diamond of \mathcal{G} such that (i) for each $v \in D_1 \cup D_2$, we have $\mathcal{L}_*(v) = \mathcal{L}_\diamond(v)$, and (ii) \mathcal{L}_\diamond is locally consistent for all vertices in $D_2 \cup D_3$. Observe that $D_2 \cup D_3 \supseteq B_2 \cup B_3$. By Lemma 1, there exists a complete labeling \mathcal{L}'_\diamond of \mathcal{G}' such that the two conditions in Lemma 1 are met. We show that this implies that \mathcal{L}'_* is extendible.

Lemma 1 guarantees that $\mathcal{L}_\diamond(v) = \mathcal{L}'_\diamond(\phi(v))$ for each $v \in B_0 \cup B_1 \cup B_2$ and its corresponding vertex $\phi(v) \in B'_0 \cup B'_1 \cup B'_2$. Since $D'_1 \cup D'_2 \subseteq B'_0 \cup B'_1 \cup B'_2$, we have $\mathcal{L}'_*(v') = \mathcal{L}'_\diamond(v')$ for each $v' \in D'_1 \cup D'_2$.

Since \mathcal{L}_\diamond is locally consistent for all vertices in $D_2 \cup D_3$, Lemma 1 guarantees that \mathcal{L}'_\diamond is locally consistent for all vertices in $D'_2 \cup D'_3$. More precisely, due to Condition 1, \mathcal{L}'_\diamond is locally consistent for all vertices in $(D'_2 \cup D'_3) - B'_3$; due to Condition 2, \mathcal{L}'_\diamond is locally consistent for all vertices in $B'_2 \cup B'_3$.

Thus, \mathcal{L}'_* is extendible, as the complete labeling \mathcal{L}'_\diamond of \mathcal{G}' satisfies: (i) for each $v' \in D'_1 \cup D'_2$, we have $\mathcal{L}'_*(v') = \mathcal{L}'_\diamond(v')$, and (ii) \mathcal{L}'_\diamond is locally consistent for all vertices in $D'_2 \cup D'_3$. \square

3.7 The Number of Equivalence Classes

An important feature of $\tilde{*}$ is that it has a *constant* number of equivalence classes, for any fixed number p of poles. Which constant is not important, but we shall work out an upper bound nonetheless.¹³

Consider a partially labeled graph \mathcal{H} with poles $S = (v_1, \dots, v_p)$. Let $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ and define $\mathcal{Q} = (Q, \mathcal{L})$ to be the subgraph of \mathcal{H} induced by $D_1 \cup D_2$. Observe that the equivalence class of (\mathcal{H}, S) is determined by (i) the topology of Q (including its input labels from Σ_{in} , if \mathcal{P} has input labels), (ii) the locations of the poles $S \subseteq V(Q)$ in Q , and (iii) the subset of all output labelings of $V(Q) = D_1 \cup D_2$ that are extendible.

The number of vertices in $D_1 \cup D_2$ is at most $p\Delta^{2r}$. The total number of distinct graphs of at most $p\Delta^{2r}$ vertices (with input labels from Σ_{in} and a set of p designated poles) is at most $2^{\binom{p\Delta^{2r}}{2}} |\Sigma_{\text{in}}|^{p\Delta^{2r}}$. The total number of output labelings of $D_1 \cup D_2$ is at most $|\Sigma_{\text{out}}|^{p\Delta^{2r}}$. Therefore, the total number of equivalence classes of graphs with p poles is at most $2^{\binom{p\Delta^{2r}}{2}} |\Sigma_{\text{in}}|^{p\Delta^{2r}} 2^{|\Sigma_{\text{out}}|^{p\Delta^{2r}}}$, which is constant whenever $\Delta, r, |\Sigma_{\text{in}}|, |\Sigma_{\text{out}}|$, and p are.

3.8 A Pumping Lemma for Trees

In this section we consider partially labeled trees with one and two poles; they are called *unipolar* (or *rooted*) and *bipolar*, respectively. Let $\mathcal{T} = (T, \mathcal{L})$ be a unipolar tree with pole $S = (z)$, $z \in V(T)$. Define $\text{Class}(\mathcal{T})$ to be the equivalence class of (\mathcal{T}, S) w.r.t. $\tilde{*}$. Notice that whether a partially labeled rooted tree \mathcal{T} admits a legal labeling is determined by $\text{Class}(\mathcal{T})$ (Theorem 3). We say that a class c is *good* if each partially labeled rooted tree in the class c admits a legal labeling; otherwise the class c is *bad*. We write \mathcal{C} to denote the set of all classes. Notice that $|\mathcal{C}|$ is constant. The following lemma is a specialization of Theorem 4.

Lemma 2. *Let \mathcal{T} be a partially labeled rooted (unipolar) tree, and let \mathcal{T}' be a rooted subtree of \mathcal{T} , whose leaves are also leaves of \mathcal{T} . Let \mathcal{T}'' be another partially labeled rooted tree such that $\text{Class}(\mathcal{T}') = \text{Class}(\mathcal{T}'')$. Then replacing \mathcal{T}' with \mathcal{T}'' does not alter the class of \mathcal{T} .*

Let $\mathcal{H} = (H, \mathcal{L})$ be a bipolar tree with poles $S = (s, t)$. The unique oriented path in H from s to t is called the *core path* of \mathcal{H} . It is more convenient to express a bipolar tree as a *sequence of rooted/unipolar trees*, as follows. The partially labeled bipolar tree $\mathcal{H} = (\mathcal{T}_i)_{i \in [k]}$ is formed by arranging the roots of unipolar trees (\mathcal{T}_i) into a path $P = (v_1, \dots, v_k)$, where v_i is the root/pole of \mathcal{T}_i . The two poles of \mathcal{H} are $s = v_1$ and $t = v_k$, so P is the core path of \mathcal{H} . Define $\text{Type}(\mathcal{H})$ as the equivalence class of $(\mathcal{H}, S = (s, t))$ w.r.t. $\tilde{*}$. The following lemma follows from Theorem 4.

Lemma 3. *Let \mathcal{H} be a partially labeled bipolar tree with poles (s, t) . Let \mathcal{T} be \mathcal{H} , but regarded as a unipolar tree rooted at s . Then $\text{Class}(\mathcal{T})$ is determined by $\text{Type}(\mathcal{H})$. If we write $\mathcal{H} = (\mathcal{T}_i)_{i \in [k]}$, then $\text{Type}(\mathcal{H})$ is determined by $\text{Class}(\mathcal{T}_1), \dots, \text{Class}(\mathcal{T}_k)$.*

¹³For the sake of simplicity, in the calculation we assume that the underlying LCL problem does not refer to port-numbering. It is straightforward to see that even if port-numbering is taken into consideration, the number of equivalence classes (for any fixed number p) is still a constant.

Let $\mathcal{G} = (G, \mathcal{L})$ be a partially labeled graph, and let $\mathcal{H} = (H, \mathcal{L})$ be a bipolar subtree of \mathcal{G} with poles (s, t) . Let \mathcal{H}' be another partially labeled bipolar tree. Recall that $\mathcal{G}' = \text{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$ is defined as the partially labeled graph resulting from replacing the subgraph \mathcal{H} with \mathcal{H}' in \mathcal{G} . We write $\mathcal{G}' = (G', \mathcal{L}')$ and $\mathcal{H}' = (H', \mathcal{L}')$. The following lemmas follow from Theorems 3 and 4.

Lemma 4. *Consider $\mathcal{G}' = \text{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$. If $\text{Type}(\mathcal{H}') = \text{Type}(\mathcal{H})$ and \mathcal{G} admits a legal labeling \mathcal{L}_\diamond , then \mathcal{G}' admits a legal labeling \mathcal{L}'_\diamond such that $\mathcal{L}_\diamond(v) = \mathcal{L}'_\diamond(v')$ for each vertex $v \in V(G) - V(H)$ and its corresponding $v' \in V(G') - V(H')$.*

Lemma 5. *Suppose that $\mathcal{G} = (\mathcal{T}_i)_{i \in [k]}$ is a partially labeled bipolar tree, $\mathcal{H} = (\mathcal{T}_i, \dots, \mathcal{T}_j)$ is a bipolar subtree of \mathcal{G} , and \mathcal{H}' is some other partially labeled bipolar tree with $\text{Type}(\mathcal{H}') = \text{Type}(\mathcal{H})$. Then $\mathcal{G}' = \text{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$ is a partially labeled bipolar tree and $\text{Type}(\mathcal{G}') = \text{Type}(\mathcal{G})$.*

Lemma 6. *Let $\mathcal{H} = (\mathcal{T}_i)_{i \in [k]}$ and $\mathcal{H}' = (\mathcal{T}_i)_{i \in [k+1]}$ be identical to \mathcal{H} in its first k trees. Then $\text{Type}(\mathcal{H}')$ is a function of $\text{Type}(\mathcal{H})$ and $\text{Class}(\mathcal{T}_{k+1})$.*

Lemma 6 is what allows us to bring classical automata theory into play. Suppose that we somehow computed and stored $c_i = \text{Class}(\mathcal{T}_i)$ at the root of \mathcal{T}_i . Lemma 6 implies that a finite automaton walking along the core path of $\mathcal{H}' = (\mathcal{T}_i)_{i \in [k+1]}$, can compute $\text{Type}(\mathcal{H}')$, by reading the vector (c_1, \dots, c_{k+1}) one character at a time. The number of states in the finite automaton depends only on the number of types (which is constant) and is independent of $k+1$ and the size of the individual trees (\mathcal{T}_i) . Define $\ell_{\text{pump}} = O(1)$ as the number of states in this finite automaton. The following *pumping lemma* for bipolar trees is analogous to the pumping lemma for regular languages.

Lemma 7. *Let $\mathcal{H} = (\mathcal{T}_1, \dots, \mathcal{T}_k)$, with $k \geq \ell_{\text{pump}}$. We regard each \mathcal{T}_i in the string notation $\mathcal{H} = (\mathcal{T}_1, \dots, \mathcal{T}_k)$ as a character. Then \mathcal{H} can be decomposed into three substrings $\mathcal{H} = x \circ y \circ z$ such that (i) $|xy| \leq \ell_{\text{pump}}$, (ii) $|y| \geq 1$, and (iii) $\text{Type}(x \circ y^j \circ z) = \text{Type}(\mathcal{H})$ for each non-negative integer j .*

We will use Lemma 7 to expand the length of the core path of a bipolar tree to a desired *target length* w . The specification for the function `Pump` is as follows.

Pump Let $\mathcal{H} = (\mathcal{T}_i)_{i \in [k]}$ be a partially labeled bipolar tree with $k \geq \ell_{\text{pump}}$. `Pump` (\mathcal{H}, w) produces a partially labeled bipolar tree $\mathcal{H}' = (\mathcal{T}'_i)_{i \in [k']}$ such that (i) $\text{Type}(\mathcal{H}) = \text{Type}(\mathcal{H}')$, (ii) $k' \in [w, w + \ell_{\text{pump}}]$, and (iii) if we let $Z = \{\mathcal{T}_i\}_{i \in [k]}$ (resp., $Z' = \{\mathcal{T}'_i\}_{i \in [k']}$) be the set of rooted trees appearing in the tree list of \mathcal{H} (resp., \mathcal{H}'), then $Z' = Z$.

By Lemma 7, such a function `Pump` exists.

3.9 Rake & Compress Graph Decomposition

In this section we describe an $O(\log n)$ -round `DetLOCAL` algorithm to decompose the vertex set $V(G)$ of a tree into the disjoint union $V_1 \cup \dots \cup V_L$, $L = O(\log n)$. Our algorithm is inspired by Miller and Reif's *parallel tree contraction* [30]. We first describe the decomposition algorithm then analyze its properties.

Fix the constant $\ell = 2(r + \ell_{\text{pump}})$, where r, ℓ_{pump} depend on the LCL problem \mathcal{P} . In the *postprocessing* step of the decomposition algorithm we compute an $(\ell, 2\ell)$ -independent set, in $O(\log^* n)$ time [29], defined as follows.

Definition 1. *Let P be a path. A set $I \subset V(P)$ is called an (α, β) -independent set if the following conditions are met: (i) I is an independent set, and I does not contain either endpoint of P , and (ii) each connected component induced by $V(P) - I$ has at least α vertices and at most β vertices, unless $|V(P)| < \alpha$, in which case $I = \emptyset$.*

The Decomposition Algorithm. The algorithm begins with $U = V(G)$ and $i = 1$, repeats Steps 1–3 until $U = \emptyset$, then executes the *Postprocessing* step.

1. For each $v \in U$:
 - (a) **Compress.** If v belongs to a path P such that $|V(P)| \geq \ell$ and $\deg_U(u) = 2$ for each $u \in V(P)$, then tag v with i_C .

- (b) **Rake.** If $\deg_U(v) = 0$, then tag v with i_R . If $\deg_U(v) = 1$ and the unique neighbor u of v in U satisfies either (i) $\deg_U(u) > 1$ or (ii) $\deg_U(u) = 1$ and $\text{ID}(v) > \text{ID}(u)$, then tag v with i_R .

2. Remove from U all vertices tagged i_C or i_R .
3. $i \leftarrow i + 1$.

Postprocessing Step. Initialize V_i as the set of all vertices tagged i_C or i_R . At this point the graph induced by V_i consists of *unbounded length* paths, but we prefer constant length paths. For each edge $\{u, v\}$ such that v is tagged i_R and u is tagged i_C , promote v from V_i to V_{i+1} . For each path P that is a connected component induced by vertices tagged i_C , compute an $(\ell, 2\ell)$ -independent set I_P of P , and then promote every vertex in I_P from V_i to V_{i+1} .¹⁴

Properties of the Decomposition. As we show below, $L = O(\log n)$ iterations suffice, i.e., $V(G) = V_1 \cup \dots \cup V_L$. The following properties are easily verified.

- Define G_i as the graph induced by vertices at level i or above: $\bigcup_{j=i}^L V_j$. For each $v \in V_i$, $\deg_{G_i}(v) \leq 2$.
- Define \mathcal{P}_i as the set of connected components (paths) induced by vertices in V_i that contain more than one vertex. For each $P \in \mathcal{P}_i$, $\ell \leq |V(P)| \leq 2\ell$ and $\deg_{G_i}(v) = 2$ for each vertex $v \in V(P)$.
- The graph G_L contains only isolated vertices, i.e., $\mathcal{P}_L = \emptyset$.

As a consequence, each vertex $v \in V_i$ falls into exactly one of two cases: (i) v has $\deg_{G_i}(v) \leq 1$ and has no neighbor in V_i , or (ii) v has $\deg_{G_i}(v) = 2$ and is in some path $P \in \mathcal{P}_i$.

Analysis. We prove that for $L = O(\log_{1+1/\ell} n) = O(\log n)$, L iterations of the graph decomposition routine suffices to decompose any n -vertex tree. Each iteration of the routine takes $O(1)$ time, and the $(\ell, 2\ell)$ -independent set computation at the end takes $O(\log^* n)$ time, so $O(\log n)$ time suffices in **DetLOCAL**.

Let W be the vertices of a connected component induced by U at the beginning of the i th iteration.¹⁵ We claim that at least a constant $\Omega(1/\ell)$ fraction of vertices in W are eliminated (i.e., tagged i_C or i_R) in the i th iteration. The proof of the claim is easy for the special case of $\ell = 1$, as follows. If W is not a single edge, then all $v \in W$ with $\deg_U(v) \leq 2$ are eliminated. Since the degree of at least half of the vertices in a tree is at most 2, the claim follows. In general, degree-2 paths of length less than ℓ are not eliminated quickly. If one endpoint of such a path is a leaf, vertices in the path are peeled off by successive *Rake* steps.

Assume w.l.o.g. that $|W| > 2(\ell + 1)$. Define $W_1 = \{v \in W \mid \deg_U(v) = 1\}$, $W_2 = \{v \in W \mid \deg_U(v) = 2\}$, and $W_3 = \{v \in W \mid \deg_U(v) \geq 3\}$.

Case 1: $|W_2| \geq \frac{\ell|W|}{\ell+1}$. The number of connected components induced by vertices in W_2 is at most $|W_1| + |W_3| - 1 < \frac{|W|}{\ell+1}$. The number of vertices in W_2 that are not tagged i_C during **Compress** is less than $\frac{(\ell-1)|W|}{\ell+1}$. Therefore, at least $\frac{\ell|W|}{\ell+1} - \frac{(\ell-1)|W|}{\ell+1} = \frac{|W|}{\ell+1}$ vertices are tagged i_C by **Compress**.

Case 2: $|W_2| < \frac{\ell|W|}{\ell+1}$. In any tree $|W_1| > |W_3|$, so $|W_1| > \frac{|W_1| + |W_3|}{2} = \frac{|W| - |W_2|}{2} \geq \frac{|W|}{2(\ell+1)}$. Therefore, at least $\frac{|W|}{2(\ell+1)}$ vertices are tagged i_R by **Rake**.

Hence the claim follows.

3.10 Extend and Label Operations

In this section we define two operations **Extend** and **Label** which are used extensively in Sections 3.11–3.12. The operation **Extend** is parameterized by a target length $w \geq \ell = 2(r + \ell_{\text{pump}})$. The operation **Label** is parameterized by a function f which takes a partially labeled bipolar tree \mathcal{H} as input, and assigns output labels to the vertices in $v \in N^{r-1}(e)$, where e is the middle edge in the core path of \mathcal{H} .¹⁶

¹⁴The set V_i in the graph decomposition is analogous to (but clearly different from) the set V_i defined in the Hierarchical $2\frac{1}{2}$ -coloring problem from Section 2.

¹⁵In general, the graph induced by U is a forest. It is simpler to analyze a single connected component W .

¹⁶By definition, if $e = \{x, y\}$ then $N^{r-1}(e) = N^{r-1}(x) \cup N^{r-1}(y)$.




Figure 4: **Left:** A bipolar subtree \mathcal{H} is attached to the rest of the graph \mathcal{G} via edges $\{u, s\}, \{v, t\}$. The pink nodes have been pre-committed to output labels by **Label** ($r = 1$). **Right:** The **Duplicate-Cut** operation duplicates \mathcal{H} and attaches one copy to u and the other to v .

Label. Let $\mathcal{H} = (\mathcal{T}_1, \dots, \mathcal{T}_x)$ be a partially labeled bipolar tree with $x \geq \ell$. Let (v_1, \dots, v_x) be the core path of \mathcal{H} and $e = \{v_{\lfloor x/2 \rfloor}, v_{\lfloor x/2 \rfloor + 1}\}$ be the middle edge of the core path. It is guaranteed that all vertices in $N^{r-1}(e)$ in \mathcal{H} are not already assigned output labels. The partially labeled bipolar tree $\mathcal{H}' = \text{Label}(\mathcal{H})$ is defined as the result of assigning output labels to vertices in $N^{r-1}(e)$ by the function f .¹⁷

Extend. Let $\mathcal{H} = (\mathcal{T}_1, \dots, \mathcal{T}_x)$ be a partially labeled bipolar tree with $x \in [\ell, 2w]$. The partially labeled bipolar tree $\mathcal{H}' = \text{Extend}(\mathcal{H})$ is defined as follows. Consider the decomposition $\mathcal{H} = \mathcal{X} \circ \mathcal{Y} \circ \mathcal{Z}$, where $\mathcal{Y} = (\mathcal{T}_{\lfloor x/2 \rfloor - r + 1}, \dots, \mathcal{T}_{\lfloor x/2 \rfloor + r})$. Then $\mathcal{H}' = \text{Pump}(\mathcal{X}, w) \circ \mathcal{Y} \circ \text{Pump}(\mathcal{Z}, w)$.

Intuitively, the goal of the operation **Extend** is to extend the length of the core path of \mathcal{H} while preserving the type of \mathcal{H} , due to Lemma 5. Suppose that the number of vertices in the core path of \mathcal{H} is in the range $[\ell, 2\ell]$. The prefix \mathcal{X} and suffix \mathcal{Z} are stretched to lengths in the range $[w, w + \ell_{\text{pump}}]$, and the middle part \mathcal{Y} has length $2r$, so the core path of \mathcal{H}' has length in the range $[2(w + r), 2(w + r + \ell_{\text{pump}})]$.

The reason that the **Extend** operation does not modify the middle part \mathcal{Y} is to ensure that (given any labeling function f) the type of $\mathcal{H}' = \text{Extend}(\text{Label}(\mathcal{H}))$ is invariant over all choices of the parameter w .¹⁸ We have the following lemma.

Lemma 8. Let $\mathcal{G} = (G, \mathcal{L})$ be a partially labeled graph and $\mathcal{H} = (H, \mathcal{L})$ be a bipolar subtree of \mathcal{G} with poles (s, t) . Let $\tilde{\mathcal{H}}$ be another partially labeled bipolar tree with $\text{Type}(\tilde{\mathcal{H}}) = \text{Type}(\mathcal{H})$ and $\mathcal{H}' = \text{Extend}(\text{Label}(\tilde{\mathcal{H}}))$. If $\mathcal{G}' = \text{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$ admits a legal labeling \mathcal{L}'_{\diamond} , then \mathcal{G} admits a legal labeling \mathcal{L}_{\diamond} such that $\mathcal{L}_{\diamond}(v) = \mathcal{L}'_{\diamond}(v')$ for each vertex $v \in V(G) - V(H)$ and its corresponding vertex $v' \in V(G') - V(H')$.

Proof. Recall that the operation **Extend** guarantees that $\text{Type}(\text{Extend}(\tilde{\mathcal{H}})) = \text{Type}(\tilde{\mathcal{H}}) = \text{Type}(\mathcal{H})$. Define $\mathcal{H}'' = \text{Extend}(\tilde{\mathcal{H}})$ and $\mathcal{G}'' = \text{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}'')$. Observe that $\mathcal{H}' = \text{Extend}(\text{Label}(\tilde{\mathcal{H}}))$ can be seen as the result of fixing the output labels of some unlabeled vertices in $\mathcal{H}'' = \text{Extend}(\tilde{\mathcal{H}})$. Therefore, \mathcal{L}'_{\diamond} is also a legal labeling of \mathcal{G}'' . By Lemma 4, the desired legal labeling \mathcal{L}_{\diamond} of $\mathcal{G} = \text{Replace}(\mathcal{G}'', \mathcal{H}'', \mathcal{H})$ can be obtained from the legal labeling \mathcal{L}'_{\diamond} of \mathcal{G}'' . \square

In addition to **Extend** and **Label**, we also modify trees using the **Duplicate-Cut** operation, defined below.

Duplicate-Cut. Let $\mathcal{G} = (G, \mathcal{L})$ be a partially labeled graph and $\mathcal{H} = (H, \mathcal{L})$ be a bipolar subtree with poles (s, t) . Suppose that \mathcal{H} is connected to the rest of \mathcal{G} via two edges $\{u, s\}$ and $\{v, t\}$. The partially labeled graph $\mathcal{G}' = \text{Duplicate-Cut}(\mathcal{G}, \mathcal{H})$ is formed by (i) duplicating \mathcal{H} and the edges $\{u, s\}, \{v, t\}$ so that u and v are attached to both copies of \mathcal{H} , (ii) removing the edge that connects u to one \mathcal{H} , and removing the edge from v to the other copy of \mathcal{H} .

¹⁷Note that the neighborhood function is evaluated w.r.t. H . In particular, the set $N^{r-1}(e)$ contains the vertices $v_{\lfloor x/2 \rfloor - r + 1}, \dots, v_{\lfloor x/2 \rfloor + r}$ of the core path, and also contains parts of the trees $\mathcal{T}_{\lfloor x/2 \rfloor - r + 1}, \dots, \mathcal{T}_{\lfloor x/2 \rfloor + r}$.

¹⁸Notice that **Extend** is applied *after* **Label**. Thus, the vertices that are assigned output labels during **Label** must be within the middle part \mathcal{Y} , no part of which is modified during **Extend**.

Later on we will see that both poles of a bipolar tree are responsible for computing the labeling of the tree. On the other hand, we do not want the poles to have to communicate too much. As Lemma 9 shows, the Duplicate-Cut operation (in conjunction with Extend and Label) allows both poles to work independently and cleanly integrate their labelings afterward.

Lemma 9. *Suppose that $\mathcal{H} = \text{Extend}(\text{Label}(\tilde{\mathcal{H}}))$ for some partially labeled bipolar tree $\tilde{\mathcal{H}}$. If $\mathcal{G}' = \text{Duplicate-Cut}(\mathcal{G}, \mathcal{H})$ admits a legal labeling \mathcal{L}'_\diamond , then \mathcal{G} admits a legal labeling \mathcal{L}_\diamond such that $\mathcal{L}_\diamond(v) = \mathcal{L}'_\diamond(v')$ for each vertex $v \in V(G) - V(H)$ and a particular corresponding vertex v' in \mathcal{G}' .*

Proof. Let $\mathcal{G}' = (G', \mathcal{L}')$. We write $\mathcal{H} = (\mathcal{T}_1, \dots, \mathcal{T}_x)$. Let (v_1, \dots, v_x) be the core path of \mathcal{H} , where $s = v_1$ and $t = v_x$ are the two poles of \mathcal{H} . Let $\{u, s\}$ and $\{v, t\}$ be the two edges that connect H to the rest of G . Let $e = \{v_j, v_{j+1}\}$ be the edge in the core path of \mathcal{H} such that the output labels of vertices in $N^{r-1}(e)$ in \mathcal{H} were fixed by Label.¹⁹ We write \mathcal{H}_u (resp., \mathcal{H}_v) to denote the copy of \mathcal{H} in \mathcal{G}' that attaches to u (resp., v). Define a mapping ϕ from $V(G)$ to $V(G')$ as follows.

- For $z \in V(G) \setminus V(H)$, $\phi(z)$ is the corresponding vertex in G' .
- For $z \in \bigcup_{i=1}^j \mathcal{T}_i$, $\phi(z)$ is the corresponding vertex in \mathcal{H}_u .
- For $z \in \bigcup_{i=j+1}^x \mathcal{T}_i$, $\phi(z)$ is the corresponding vertex in \mathcal{H}_v .

We set $\mathcal{L}_\diamond(z) = \mathcal{L}'_\diamond(\phi(z))$ for each $z \in V(G)$. It is straightforward to verify that the distance- r neighborhood view (with output labeling \mathcal{L}_\diamond) of each vertex $z \in V(G)$ is the same as the distance- r neighborhood view (with output labeling \mathcal{L}'_\diamond) of its corresponding vertex $\phi(z)$ in G' . Thus, \mathcal{L}_\diamond is a legal labeling. \square

Notice that in the proof of Lemma 9, the only property of \mathcal{H} that we use is that $N^{r-1}(e)$ was assigned output labels in the application of Label($\tilde{\mathcal{H}}$).

3.11 A Hierarchy of Partially Labeled Trees

In this section we construct several sets of partially labeled unipolar and bipolar trees— $\{\mathcal{T}_i\}$, $\{\mathcal{H}_i\}$, and $\{\mathcal{H}_i^+\}$, $i \in \mathbb{Z}^+$ —using the operations Extend and Label. If each member of $\mathcal{T}^* = \bigcup_i \mathcal{T}_i$ admits a legal labeling, then we can use these trees to design an $O(\log n)$ -time DetLOCAL algorithm for \mathcal{P} . Each $\mathcal{T} \in \mathcal{T}^*$ is partially labeled in the following restricted manner. The tree $\mathcal{T} = (T, \mathcal{L})$ has a set of *designated edges* such that $\mathcal{L}(v) \neq \perp$ is defined if and only if $v \in N^{r-1}(e)$ for some designated edge e ; these vertices were issued labels by some invocation of Label.

The sets of bipolar trees $\{\mathcal{H}_i\}_{i \in \mathbb{Z}^+}$ and $\{\mathcal{H}_i^+\}_{i \in \mathbb{Z}^+}$ and unipolar trees $\{\mathcal{T}_i\}_{i \in \mathbb{Z}^+}$ are defined inductively. In the base case we have $\mathcal{T}_1 = \{\mathcal{T}\}$, where \mathcal{T} is the unique unlabeled, single-vertex, unipolar tree.

\mathcal{T} Sets: For each $i > 1$, \mathcal{T}_i consists of all partially labeled rooted trees \mathcal{T} formed in the following manner.

The root z of \mathcal{T} has degree $0 \leq \deg(z) \leq \Delta$. Each child of z is either (i) the root of a partially labeled rooted tree \mathcal{T}' from \mathcal{T}_{i-1} (having degree at most $\Delta - 1$), or (ii) one of the two poles of a bipolar tree \mathcal{H} from \mathcal{H}_{i-1}^+ .

\mathcal{H} Sets: For each $i \geq 1$, \mathcal{H}_i contains all partially labeled bipolar trees $\mathcal{H} = (\mathcal{T}_j)_{j \in [x]}$ such that $x \in [\ell, 2\ell]$, and for each $j \in [x]$, $\mathcal{T}_j \in \mathcal{T}_i$, where the root of \mathcal{T}_j has degree at most $\Delta - 2$. For example, since \mathcal{T}_1 contains only the single-vertex unlabeled tree, \mathcal{H}_1 is the set of all bipolar, unlabeled paths with between ℓ and 2ℓ vertices.

\mathcal{H}^+ Sets: For each $i \geq 1$, \mathcal{H}_i^+ is constructed by the following procedure. If $i = 1$, initialize $\mathcal{H}_1^+ \leftarrow \emptyset$, otherwise initialize $\mathcal{H}_i^+ \leftarrow \mathcal{H}_{i-1}^+$. Consider each $\mathcal{H} \in \mathcal{H}_i$ in some canonical order. If there does *not* already exist a partially labeled bipolar tree $\tilde{\mathcal{H}}$ such that $\text{Type}(\tilde{\mathcal{H}}) = \text{Type}(\mathcal{H})$ and $\text{Extend}(\text{Label}(\tilde{\mathcal{H}})) \in \mathcal{H}_i^+$, then update $\mathcal{H}_i^+ \leftarrow \mathcal{H}_i^+ \cup \{\text{Extend}(\text{Label}(\mathcal{H}))\}$.

¹⁹Because Pump usually does not extend \mathcal{X} and \mathcal{Z} by precisely the same amount, the edge e is generally not *exactly* in the middle.

Observe that whereas $\{\mathcal{T}_i\}$ and $\{\mathcal{H}_i\}$ grow without end, and contain arbitrarily large trees, the cardinality of \mathcal{H}_i^+ is at most the total number of types, which is constant.²⁰ This is due to the observation that whenever we add a new partially labeled bipolar tree $\text{Extend}(\text{Label}(\mathcal{H}))$ to \mathcal{H}_i^+ , it is guaranteed that there is no other partially labeled bipolar tree $\mathcal{H}' = \text{Extend}(\text{Label}(\mathcal{H})) \in \mathcal{H}_i^+$ such that $\text{Type}(\mathcal{H}') = \text{Type}(\mathcal{H})$. The property that $|\mathcal{H}_i^+|$ is constant is crucial in the proof of Lemma 16. Lemmas 10–12 reveal some useful properties of these sets.

Lemma 10. *We have (i) $\mathcal{T}_1 \subseteq \mathcal{T}_2 \subseteq \dots$, (ii) $\mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \dots$, and (iii) $\mathcal{H}_1^+ \subseteq \mathcal{H}_2^+ \subseteq \dots$.*

Proof. By construction, we already have $\mathcal{H}_1^+ \subseteq \mathcal{H}_2^+ \subseteq \dots$. Due to the construction of \mathcal{H}_i from the set \mathcal{T}_i , it is guaranteed that if $\mathcal{T}_j \subseteq \mathcal{T}_{j+1}$ holds then $\mathcal{H}_j \subseteq \mathcal{H}_{j+1}$ holds as well. Thus, it suffices to show that $\mathcal{T}_1 \subseteq \mathcal{T}_2 \subseteq \dots$. This is proved by induction.

For the base case, we have $\mathcal{T}_1 \subseteq \mathcal{T}_2$ because \mathcal{T}_2 also contains $\mathcal{T} \in \mathcal{T}_1$, the unlabeled, single-vertex, unipolar tree.

For the inductive step, suppose that we already have $\mathcal{T}_1 \subseteq \mathcal{T}_2 \subseteq \dots \subseteq \mathcal{T}_i$, $i \geq 2$. Then we show that $\mathcal{T}_i \subseteq \mathcal{T}_{i+1}$. Observe that the set \mathcal{T}_{i+1} contains all partially labeled rooted trees constructed by attaching partially labeled trees from the sets \mathcal{H}_i^+ and \mathcal{T}_i to the root vertex. We already know that $\mathcal{H}_{i-2}^+ \subseteq \mathcal{H}_{i-1}^+$, and by the inductive hypothesis we have $\mathcal{T}_{i-2} \subseteq \mathcal{T}_{i-1}$. Thus, each $\mathcal{T} \in \mathcal{T}_i$ must also appear in the set \mathcal{T}_{i+1} . \square

If \mathcal{T} and \mathcal{H} are arbitrary sets of unipolar and bipolar trees, we define $\text{Class}(\mathcal{T}) = \{\text{Class}(\mathcal{T}) \mid \mathcal{T} \in \mathcal{T}\}$ and $\text{Type}(\mathcal{H}) = \{\text{Type}(\mathcal{H}) \mid \mathcal{H} \in \mathcal{H}\}$ to be the set of classes and types appearing among them.

Lemma 11. *Define $k^* = |\mathcal{C}|$, where \mathcal{C} is the set of all classes. Then $\text{Class}(\mathcal{T}^*) = \text{Class}(\mathcal{T}_{k^*})$.*

Proof. For each $i > 1$, $\text{Class}(\mathcal{T}_i)$ depends only on $\text{Type}(\mathcal{H}_{i-1}^+)$ and $\text{Class}(\mathcal{T}_{i-1})$, due to Lemmas 2 and 3. Let i^* be the smallest index such that $\text{Class}(\mathcal{T}_{i^*}) = \text{Class}(\mathcal{T}_{i^*+1})$. Then we have $\text{Type}(\mathcal{H}_{i^*}) = \text{Type}(\mathcal{H}_{i^*+1})$ and as a consequence, $\mathcal{H}_{i^*}^+ = \mathcal{H}_{i^*+1}^+$. This implies that $\text{Class}(\mathcal{T}_{i^*+1}) = \text{Class}(\mathcal{T}_{i^*+2})$. By repeating the same argument, we conclude that for each $j \geq i^*$, we have $\text{Class}(\mathcal{T}_j) = \text{Class}(\mathcal{T}_{i^*}) = \text{Class}(\mathcal{T}^*)$. Since $\mathcal{T}_1 \subseteq \mathcal{T}_2 \subseteq \dots$ (Lemma 10), we have $i^* \leq |\mathcal{C}|$. \square

Lemma 12. *For each i , $\text{Class}(\mathcal{T}_i)$ does not depend on the parameter w used in the operation *Extend*.*

Proof. Let $\mathcal{H} = (\mathcal{T}_1, \dots, \mathcal{T}_x)$ be any partially labeled bipolar tree with $x \geq 2r + 2\ell_{\text{pump}}$. The type of $\mathcal{H}' = \text{Extend}(\mathcal{H})$ is invariant over all choices of the parameter w . Thus, by induction, the sets $\text{Class}(\mathcal{T}_i)$, $\text{Type}(\mathcal{H}_i)$, and $\text{Type}(\mathcal{H}_i^+)$ are also invariant over the choice of w . \square

Feasible Labeling Function. In view of Lemma 12, $\text{Class}(\mathcal{T}^*)$ depends only on the choice of the labeling function f used by *Label*. We call a function f *feasible* if implementing *Label* with f makes each tree in $\text{Class}(\mathcal{T}^*)$ good, i.e., its partial labeling can be extended to a complete and legal labeling. In Section 3.12 we show that given a feasible function, we can generate a *DetLOCAL* algorithm to solve \mathcal{P} in $O(\log n)$ -time. In Section 3.13, we show that (i) a feasible function can be derived from any given an $n^{o(1)}$ -time *RandLOCAL* algorithm for \mathcal{P} , and (ii) the existence of a feasible function is decidable. These results together imply the $\omega(\log n) - n^{o(1)}$ gap. Moreover, given an LCL problem \mathcal{P} on bounded degree trees, it is decidable whether the *RandLOCAL* complexity of \mathcal{P} is $n^{\Omega(1)}$ or the *DetLOCAL* complexity of \mathcal{P} is $O(\log n)$.

3.12 A $O(\log n)$ -time *DetLOCAL* Algorithm from a Feasible Labeling Function

In this section, we show that given a feasible function f for the LCL problem \mathcal{P} , it is possible to design a $O(\log n)$ -time *DetLOCAL* algorithm for \mathcal{P} on bounded degree trees.

Regardless of f , the algorithm begins by computing the graph decomposition $V(G) = V_1 \cup \dots \cup V_L$, with $L = O(\log n)$; see Section 3.9. We let the three infinite sequences $\{\mathcal{H}_i\}_{i \in \mathbb{Z}^+}$, $\{\mathcal{H}_i^+\}_{i \in \mathbb{Z}^+}$, and $\{\mathcal{T}_i\}_{i \in \mathbb{Z}^+}$ be constructed with respect to the feasible f and any parameter w .

²⁰However, it is not necessarily true that \mathcal{H}_i^+ contains at most one bipolar tree of each type. The *Extend* operation is type-preserving, but this is not true of *Label*: $\text{Type}(\text{Label}(H))$ may not equal $\text{Type}(H)$, so it is possible that \mathcal{H}_i^+ contains two members of the same type.

A Sequence of Partially Labeled Graphs. We define below a sequence of partially labeled graphs $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_L$, where \mathcal{R}_1 is the unlabeled tree G (the underlying communications network), and \mathcal{R}_{i+1} is constructed from \mathcal{R}_i using the graph operations **Extend**, **Label**, and **Duplicate-Cut**. An alternative, and helpful way to visualize \mathcal{R}_i is that it is obtained by stripping away some vertices of G , and then grafting on some *imaginary* subtrees to a subset of its remaining vertices. Formally, let \tilde{V}_i be the union of V_i and the subset of lower level vertices in $\bigcup_{j=1}^{i-1} V_j$ that are adjacent to a higher level vertex in $\bigcup_{j=i+1}^L V_j$. Let \tilde{G}_i be the subgraph of G induced by $\tilde{V}_i \cup \bigcup_{j=i+1}^L V_j$. The graph \mathcal{R}_i is formed by taking \tilde{G}_i , and identifying each vertex $u \in \tilde{V}_i$ with the root of a partially labeled imaginary tree \mathcal{T}_u (defined within the proof of Lemma 13). Moreover, if $u \in V_j$ then $\mathcal{T}_u \in \mathcal{T}_j$. Since G_L consists solely of isolated vertices, \mathcal{R}_L is the disjoint union of trees drawn from \mathcal{T}_L .

The official, inductive construction of \mathcal{R}_i is described in the proof of Lemma 13.

Lemma 13. *Suppose that a feasible function f is given. The partially labeled graphs $\mathcal{R}_1, \dots, \mathcal{R}_L$ and partially labeled trees $\{\mathcal{T}_v \mid v \in V(G)\}$ are constructed in $O(\log n)$ time such that:*

1. For each $i \in [1, L]$, each vertex $v \in V_i$ knows $\mathcal{T}_v \in \mathcal{T}_i$. (The graph \mathcal{R}_i is formed by grafting $\{\mathcal{T}_v \mid v \in \tilde{V}_i\}$ onto the corresponding vertices in \tilde{G}_i .)
2. For each $i \in [2, L]$, given a legal labeling of \mathcal{R}_i (where each vertex in $\bigcup_{j=i+1}^L V_j$ knows its label and each $v \in \tilde{V}_i$ knows the entire labeling of \mathcal{T}_v) a legal labeling of \mathcal{R}_{i-1} can be computed in $O(1)$ time.

Proof. Part (1) of the lemma is proved by induction.

Base Case. Define $\mathcal{R}_1 = G$. This satisfies the lemma since, if $v \in V_1$, $\mathcal{T}_v \in \mathcal{T}_1$ must be the unlabeled single-vertex tree.

Inductive Step. We can assume inductively that \mathcal{R}_{i-1} and $\{\mathcal{T}_v \mid v \in \bigcup_{j=1}^{i-1} V_j\}$ have been defined and satisfy the lemma. The set \mathcal{P}_{i-1} was defined in Section 3.9. Each $P \in \mathcal{P}_{i-1}$ is a path such that $\deg_{G_{i-1}}(v) = 2$ for each vertex $v \in V(P)$ and $|V(P)| \in [\ell, 2\ell]$. Fix a path $P = (v_1, \dots, v_x) \in \mathcal{P}_{i-1}$. The bipolar graphs \mathcal{H}_P and \mathcal{H}_P^+ are defined as follows.

- Define \mathcal{H}_P to be the partially labelled bipolar tree $(\mathcal{T}_{v_1}, \dots, \mathcal{T}_{v_x})$. Notice that \mathcal{H}_P is a subgraph of \mathcal{R}_{i-1} . Since $\mathcal{T}_{v_j} \in \mathcal{T}_{i-1}$, for each $j \in [x]$, it follows that $\mathcal{H}_P \in \mathcal{H}_{i-1}$.
- Construct \mathcal{H}_P^+ as follows. Select the unique member $\tilde{\mathcal{H}} \in \tilde{\mathcal{H}}_{i-1}$ such that $\text{Type}(\tilde{\mathcal{H}}) = \text{Type}(\mathcal{H}_P)$ and let $\mathcal{H}_P^+ = \text{Extend}(\text{Label}(\tilde{\mathcal{H}})) \in \mathcal{H}_{i-1}^+$. (Due to the definition of \mathcal{H}_{i-1}^+ , such an $\tilde{\mathcal{H}}$ must exist, since $\mathcal{H}_P \in \mathcal{H}_{i-1}$.)

The partially labeled graph \mathcal{R}_i is constructed from \mathcal{R}_{i-1} with the following three-step procedure. See Figure 5 for a schematic example of how these steps work.

Step 1. Define \mathcal{R}'_{i-1} as the result of applying the following operations on \mathcal{R}_{i-1} . For each $v \in V_{i-1}$ such that \mathcal{T}_v is a connected component of \mathcal{R}_{i-1} , remove \mathcal{T}_v .²¹

Step 2. Define \mathcal{R}^+_{i-1} by the following procedure. (i) Initialize $\tilde{\mathcal{G}} \leftarrow \mathcal{R}'_{i-1}$. (ii) For each $P \in \mathcal{P}_{i-1}$, do $\tilde{\mathcal{G}} \leftarrow \text{Replace}(\tilde{\mathcal{G}}, \mathcal{H}_P, \mathcal{H}_P^+)$. (iii) Set $\mathcal{R}^+_{i-1} \leftarrow \tilde{\mathcal{G}}$.

Step 3. Define \mathcal{R}_i by the following procedure. (i) Initialize $\tilde{\mathcal{G}} \leftarrow \mathcal{R}^+_{i-1}$. (ii) For each $P \in \mathcal{P}_{i-1}$, do $\tilde{\mathcal{G}} \leftarrow \text{Duplicate-Cut}(\tilde{\mathcal{G}}, \mathcal{H}_P^+)$. (iii) Set $\mathcal{R}_i \leftarrow \tilde{\mathcal{G}}$. For $v \in V_i$, \mathcal{T}_v is now defined to be the tree in $\mathcal{R}_i - (V_i \setminus \{v\} \cup V_{i+1} \cup \dots \cup V_L)$ rooted at v .

We now need to verify that \mathcal{R}_i satisfies all the claims of the lemma. Given the partially labeled graph \mathcal{R}_i , the partially labeled trees \mathcal{T}_v for all $v \in V_i$ are uniquely determined. According to the construction of \mathcal{R}_i , each connected component of $\mathcal{R}_i - \bigcup_{j=i}^L V_j$ must be either (i) some \mathcal{T}_v , where $v \in \bigcup_{j=1}^{i-1} V_j$ or (ii) a copy

²¹A tree \mathcal{T}_v is a connected component of \mathcal{R}_{i-1} if and only if v 's neighborhood in G contains only vertices at lower levels: V_1, \dots, V_{i-2} .

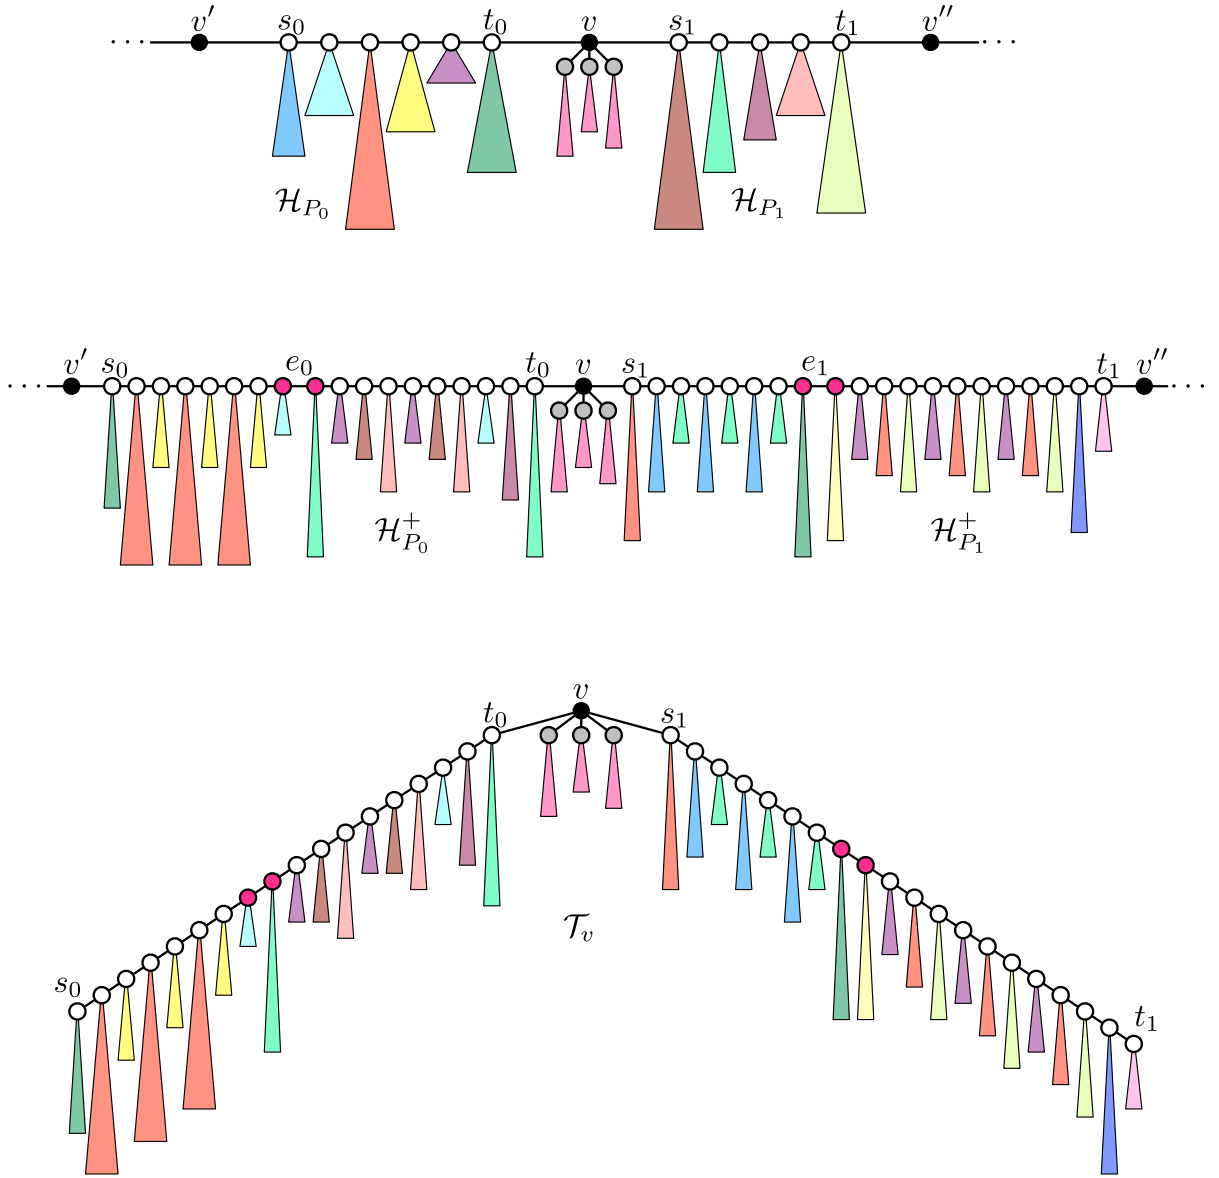


Figure 5: **Top:** In this example v was a vertex in a long degree-2 path tagged $(i-1)_C$ by the decomposition procedure, and subsequently promoted to V_i . Black vertices are in V_i (or above); white vertices are in V_{i-1} ; gray vertices are in V_{i-2} or below. The paths $P_0 = (s_0, \dots, t_0)$ and $P_1 = (s_1, \dots, t_1)$ adjacent to v have constant length, between ℓ and 2ℓ . The colored subtrees grafted onto white and gray vertices are imaginary subtrees formed in the construction of \mathcal{R}_{i-1} . **Middle:** The graph is transformed by finding the graph $\tilde{\mathcal{H}}_b \in \mathcal{H}_{i-1}^+$, $b \in \{0, 1\}$ that has the same type as \mathcal{H}_{P_b} , and replacing \mathcal{H}_{P_b} with $\mathcal{H}_{P_b}^+ = \text{Extend}(\text{Label}(\tilde{\mathcal{H}}_b))$. The vertices receiving pre-committed labels are indicated in pink ($r = 1$). **Bottom:** We duplicate $\mathcal{H}_{P_b}^+$, $b \in \{0, 1\}$, and attach one of the copies of each duplicate to v . (The copies of $\mathcal{H}_{P_b}^+$ attached to v', v'' are not shown.) The tree \mathcal{T}_v is the resulting tree rooted at v . Since each subtree of v is in \mathcal{T}_{i-1} or \mathcal{H}_{i-1}^+ , it follows that $\mathcal{T}_v \in \mathcal{T}_i$. In this case v had no neighbors at higher levels ($i+1$ and above), so \mathcal{T}_v is a connected component of \mathcal{R}_i . Thus, v can locally compute a legal labelling of \mathcal{T}_v .

of $\mathcal{H}_P^+ \in \mathcal{H}_{i-1}^+$, where $P \in \mathcal{P}_{i-1}$. According to the inductive definition of \mathcal{T}_i , for each $v \in V_i$ we have $\mathcal{T}_v \in \mathcal{T}_i$. This concludes the induction.

We now turn to the proof of Part (2) of the lemma. Suppose that we have a legal labeling of \mathcal{R}_i , where the labeling of \mathcal{T}_v is stored in $v \in V_i$. We show how to compute a legal labeling of \mathcal{R}_{i-1} in $O(1)$ time as follows.²² Starting with any legal labeling \mathcal{L}_1 of \mathcal{R}_i , we compute a legal labeling \mathcal{L}_2 of \mathcal{R}_{i-1}^+ , a legal labeling \mathcal{L}_3 of \mathcal{R}'_{i-1} , and finally a legal labeling \mathcal{L}_4 of \mathcal{R}_{i-1} . Throughout the process, the labels of all vertices in $\bigcup_{j=i}^L V_j$ are stable under $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$, and \mathcal{L}_4 .

From \mathcal{L}_1 to \mathcal{L}_2 . Let s, t be the poles of \mathcal{H}_P^+ and u, v be the vertices outside of \mathcal{H}_P^+ in \mathcal{R}_{i-1}^+ connected to s, t , respectively. At this point u and v have legal labelings of \mathcal{T}_u and \mathcal{T}_v , both trees of which contain a copy of \mathcal{H}_P^+ . Using Lemma 9 we integrate the labelings of \mathcal{T}_u and \mathcal{T}_v to fix a single legal labeling \mathcal{L}_2 of \mathcal{H}_P^+ in \mathcal{R}_{i-1}^+ .²³

From \mathcal{L}_2 to \mathcal{L}_3 . A legal labeling \mathcal{L}_3 of \mathcal{R}'_{i-1} is obtained by applying Lemma 8. For each $P \in \mathcal{P}_{i-1}$, the labeling \mathcal{L}_3 on \mathcal{H}_P in \mathcal{R}'_{i-1} can be determined from the labeling \mathcal{L}_2 of \mathcal{H}_P^+ in \mathcal{R}_{i-1}^+ . In greater detail, suppose s, t are the poles of $\mathcal{H}_P/\mathcal{H}_P^+$, which know \mathcal{L}_2 on the $(2r-1)$ -neighborhood of $\{s, t\}$ in \mathcal{H}_P^+ . By Lemma 8, there exists a legal labeling \mathcal{L}_3 on \mathcal{H}_P , which can be succinctly encoded by fixing \mathcal{L}_3 on the $(2r-1)$ -neighborhoods of the *roots* of each unipolar tree on the core path ($s = v_1, \dots, v_x = t$) of \mathcal{H}_P . Thus, once s, t calculate \mathcal{L}_3 , they can transmit the relevant information with constant-length messages to the roots v_1, \dots, v_x . At this point each $v_j \in V_{i-1}$ can locally compute an extension of its labeling to all of \mathcal{T}_{v_j} .

From \mathcal{L}_3 to \mathcal{L}_4 . Notice that \mathcal{R}_{i-1} is simply the disjoint union of \mathcal{R}'_{i-1} —for which we already have a legal labeling \mathcal{L}_3 —and each \mathcal{T}_v that is a connected component of \mathcal{R}_{i-1} . A legal labeling \mathcal{L}_4 of \mathcal{T}_v is computed locally at v , which is guaranteed to exist since $\mathcal{T}_v \in \mathcal{T}_{i-1}$.

This concludes the proof of the lemma. □

Lemma 14. *Let \mathcal{P} be any LCL problem on trees with $\Delta = O(1)$. Given a feasible function f , the LCL problem \mathcal{P} can be solved in $O(\log n)$ time in DetLOCAL.*

Proof. First compute a graph decomposition in $O(\log n)$ time. Given the graph decomposition, each vertex $v \in V(G)$ can compute the partially labeled rooted tree \mathcal{T}_v in $O(\log n)$ rounds. Since f is feasible, each partially labeled tree in \mathcal{T}^* admits a legal labeling. Therefore, \mathcal{R}_L admits a legal labeling, and such a legal labeling can be computed without communication by the vertices in V_L . Starting with any legal labeling of \mathcal{R}_L , legal labelings of $\mathcal{R}_{L-1}, \dots, \mathcal{R}_1 = G$ can be computed in $O(\log n)$ additional time, using Lemma 13(2). □

3.13 Existence of Feasible Labeling Function

In Lemmas 15 and 16 we show *two* distinct ways to arrive at a feasible labeling function. In Lemma 15 we assume that we are given the code of a RandLOCAL algorithm \mathcal{A} that solves \mathcal{P} in $n^{o(1)}$ time with at most $1/n$ probability of failure. Using \mathcal{A} we can extract a feasible labeling function f .²⁴ Lemma 15 suffices to prove our $n^{o(1)} \rightarrow O(\log n)$ speedup theorem but, because it needs the code of \mathcal{A} , it is insufficient to answer a more basic question. Given the description of an LCL \mathcal{P} , is \mathcal{P} solvable in $O(\log n)$ time on trees or not? Lemma 16 proves that this question is, in fact, decidable, which serves to highlight the delicate boundary between decidable and undecidable problems in LCL complexity [5, 33].

Lemma 15. *Suppose that there exists a RandLOCAL algorithm \mathcal{A} that solves \mathcal{P} in $n^{o(1)}$ time on n -vertex bounded degree trees, with local probability of failure at most $1/n$. Then there exists a feasible function f .*

²²Remember that \mathcal{R}_i and \mathcal{R}_{i-1} are both *imaginary* trees, only parts of which are embedded in the actual communications network G . “Time” refers to communications rounds in the *actual* network, not any imaginary graph.

²³It is not necessary to physically store the entire \mathcal{L}_2 on \mathcal{H}_P^+ . To implement the following steps, it suffices that s, t both know what \mathcal{L}_2 is on the subgraph induced by the $(2r-1)$ -neighborhood of $\{s, t\}$ in \mathcal{H}_P^+ .

²⁴The precise running time of \mathcal{A} influences the w parameter used by Extend. For example, if \mathcal{A} runs in $O(\log^2 n)$ time then w will be smaller than if \mathcal{A} runs in $n^{1/\log \log \log n}$ time.

Proof. Define $\beta = |\Sigma_{\text{out}}|^{\Delta^r}$ to be an upper bound on the number of distinct output labelings of $N^{r-1}(e)$, where e is any edge in any graph of maximum degree Δ . Define N as the maximum number of vertices of a tree in \mathcal{T}_{k^*} over *all* choices of labeling function f . As Δ, r, k^* are all constants, we have $N = w^{O(1)}$. Define t to be the running time of \mathcal{A} on a $(\beta N + 1)$ -vertex tree. Notice that t depends on N , which depends on w .

Choices of w and f . We select w to be sufficiently large such that $w \geq 4(r + t)$. Such a w exists since \mathcal{A} runs in $n^{o(1)}$ time on an n -vertex graph, and in our case n is polynomial in w . By our choice of w , the labeled parts of $\mathcal{T} = (T, \mathcal{L}) \in \mathcal{T}_{k^*}$ are spread far apart. In particular, (i) the sets $N^{(r-1)+t}(e)$ for all designated edges e in \mathcal{T} are disjoint, (ii) for each vertex $v \in V(T)$, there is at most one designated edge e such that the set $N^{r+t}(v)$ intersects $N^{r-1+t}(e)$.

Let the function f be defined as follows. Take any bipolar tree $\mathcal{H} = (H, \mathcal{L}')$ with middle edge e on its core path. The output labels of $N^{r-1}(e)$ are assigned by selecting the *most probable labeling* that occurs when running \mathcal{A} on the tree $\mathcal{H}' = \text{Extend}(\mathcal{H})$, while pretending that the underlying graph has $\beta N + 1$ vertices. Notice that the most probable labeling occurs with probability at least $1/\beta$.

Proof Idea. In what follows, consider *any* partially labeled rooted tree $\mathcal{T} = (T, \mathcal{L}) \in \mathcal{T}_{k^*}$, where the set \mathcal{T}_{k^*} is constructed with the parameter w and function f . All we need to prove is that \mathcal{T} admits a legal labeling \mathcal{L}_\circ . Suppose that we execute \mathcal{A} on T while pretending that the total number of vertices is $\beta N + 1$. Let v be any vertex in T . According to \mathcal{A} 's specs, the probability that the output labeling of $N^r(v)$ is inconsistent with \mathcal{P} is at most $1/(\beta N + 1)$. However, it is not guaranteed that the output labeling resulting from \mathcal{A} is also consistent with \mathcal{T} , since \mathcal{T} is partially labeled. To handle the partial labeling of \mathcal{T} , our strategy is to consider a modified distribution of random bits generated by vertices in T that forces any execution of \mathcal{A} to agree with \mathcal{L} , wherever it is defined. We will later see that with an appropriately chosen distribution of random bits, the outcome of \mathcal{A} is a legal labeling of \mathcal{T} with positive probability.

Modified Distribution of Random Bits. Suppose that an execution of \mathcal{A} on a $(\beta N + 1)$ -vertex graph needs a b -bit random string for each vertex. For each designated edge e , let U_e be the set of all assignments of b -bit strings to vertices in $N^{(r-1)+t}(e)$. Define S_e as the subset of U_e such that $\rho \in S_e$ if and only if the following is true. Suppose that the b -bit string of each $u \in N^{(r-1)+t}(e)$ is $\rho(u)$. Using the b -bit string $\rho(u)$ for each $u \in N^{(r-1)+t}(e)$, the output labeling of the vertices in $N^{r-1}(e)$ resulting from executing \mathcal{A} is the same as the output labeling specified by \mathcal{L} . According to our choice of f , we must have $|S_e|/|U_e| \geq 1/\beta$.

Define the modified distribution \mathcal{D} of b -bit random strings to the vertices in T as follows. For each designated edge e , the b -bit strings of the vertices in $N^{(r-1)+t}(e)$ are chosen uniformly at random from the set S_e . For the remaining vertices, their b -bit strings are chosen uniformly at random.

Legal Labeling \mathcal{L}_\circ Exists. Suppose that \mathcal{A} is executed on T with the modified distribution of random bits \mathcal{D} . Then it is guaranteed that \mathcal{A} outputs a complete labeling that is consistent with \mathcal{T} . Of course, the probability that \mathcal{A} outputs an *illegal* labeling under \mathcal{D} may be different. We need to show that \mathcal{A} nonetheless succeeds with non-zero probability.

Consider any vertex $v \in V(T)$. The probability that $N^r(v)$ is inconsistent with \mathcal{P} is at most $\beta/(\beta N + 1)$ under distribution \mathcal{D} , as explained below. Due to our choice of w , the set $N^{r+t}(v)$ intersects at most one set $N^{r-1+t}(e)$ where e is a designated edge. Let U_v be the set of all assignments of b -bit strings to vertices in $N^{r+t}(v)$. For each $\rho \in U_v$, the probability that ρ occurs in an execution of \mathcal{A} is $1/|U_v|$ if all random bits are chosen uniformly at random, and is at most $\beta/|U_v|$ under \mathcal{D} . Thus, the probability that \mathcal{A} (using distribution \mathcal{D}) labels $N^r(v)$ incorrectly is at most $\beta/(\beta N + 1)$. The total number of vertices in T is at most N . Thus, by the union bound, the probability that the output labeling of \mathcal{A} (using \mathcal{D}) is not a legal labeling is $\beta N/(\beta N + 1) < 1$. \square

Lemma 16. *Given an LCL problem \mathcal{P} on bounded degree graphs, it is decidable whether there exists a feasible function f .*

Proof. Throughout the construction of the three infinite sequences $\{\mathcal{H}_i\}_{i \in \mathbb{Z}^+}$, $\{\mathcal{H}_i^+\}_{i \in \mathbb{Z}^+}$, and $\{\mathcal{T}_i\}_{i \in \mathbb{Z}^+}$, the number of distinct applications of the operation **Label** is constant, as $|\mathcal{H}_i^+|$ is at most the total number of types.

Therefore, the number of distinct candidate functions f that need to be examined is finite. For each candidate labeling function f (with any parameter $w \geq \ell$), in bounded amount of time we can construct the set \mathcal{T}_{k^*} , as $k^* = |\mathcal{C}|$ is a constant. By examining the classes of the partially labeled rooted trees in \mathcal{T}_{k^*} we can infer whether the function f is feasible (Lemma 11). Thus, deciding whether there exists a feasible function f can be done in bounded amount of time. \square

Combining Lemmas 14, 15, and 16, we obtain the following theorem.

Theorem 5. *Let \mathcal{P} be any LCL problem on trees with $\Delta = O(1)$. If there exists a RandLOCAL algorithm \mathcal{A} that solves \mathcal{P} in $n^{o(1)}$ rounds, then there exists a DetLOCAL algorithm \mathcal{A}' that solves \mathcal{P} in $O(\log n)$ rounds. Moreover, given a description of \mathcal{P} , it is decidable whether the RandLOCAL complexity of \mathcal{P} is $n^{\Omega(1)}$ or the DetLOCAL complexity of \mathcal{P} is $O(\log n)$.*

Remark 2. *Observe that our graph decomposition algorithm also works on graphs of girth at least $c \log n$, where c is a sufficiently large constant depending on \mathcal{P} . This implies that Theorem 5 also applies to the class of n -vertex graphs with girth $\omega(\log n)$.*

4 A Gap in the RandLOCAL Complexity Hierarchy

Consider a set \mathcal{V} of independent random variables, and a set \mathcal{X} of *bad events*, where $A \in \mathcal{X}$ depends only on some subset $\text{vbl}(A) \subset \mathcal{V}$ of variables.²⁵ The *dependency graph* $G_{\mathcal{X}} = (\mathcal{X}, \{(A, B) \mid \text{vbl}(A) \cap \text{vbl}(B) \neq \emptyset\})$ joins events by an edge if they depend on at least one common variable. The *Lovász local lemma* (LLL) and its variants give criteria under which $\Pr(\bigcap_{A \in \mathcal{X}} \bar{A}) > 0$, i.e., it is possible that all bad events do not occur. We will narrow our discussion to *symmetric* criteria, expressed in terms of p and d , where $p = \max_A \Pr(A)$ and d is the maximum degree in $G_{\mathcal{X}}$. A standard version of the LLL states that if $ep(d+1) < 1$, then $\Pr(\bigcap \bar{A}) > 0$. Given that all bad events *can* be avoided, it is often desirable to constructively find a point in the probability space (i.e., an assignment to variables in \mathcal{V}) that avoids them. This problem has been thoroughly investigated in the sequential context [31, 20, 19, 24, 25, 22, 1], but somewhat less so from the point of view of parallel and distributed computation [7, 12, 4, 6, 18].

The *distributed* constructive LLL problem is the following. The communications network is precisely $G_{\mathcal{X}}$. Each vertex (event) A knows the number of bad events in $G_{\mathcal{X}}$ and the distribution of those variables appearing in $\text{vbl}(A) \subset \mathcal{V}$. Vertices communicate for some number of rounds, and collectively reach a consensus on an assignment to \mathcal{V} in which no bad event occurs. Moser and Tardos's [31] parallel resampling algorithm implies an $O(\log^2 n)$ time RandLOCAL algorithm under the LLL criterion $ep(d+1) < 1$. Chung, Pettie, and Su [7] gave an $O(\log_{1/epd^2} n)$ time algorithm under the LLL criterion $epd^2 < 1$ and an $O(\log n / \log \log n)$ time algorithm under criterion $p \cdot \text{poly}(d)2^d = O(1)$. They observed that under *any* criterion of the form $p \cdot f(d) < 1$, $\Omega(\log^* n)$ time is necessary. Ghaffari's [12] weak MIS algorithm, together with [7], implies an $O(\log d \cdot \log_{1/ep(d+1)} n)$ algorithm under LLL criterion $ep(d+1) < 1$. Brandt et al. [4] proved that $\Omega(\log_d \log n)$ time in RandLOCAL is necessary, even under the permissive LLL criterion $p2^d \leq 1$. Chang et al. [6]'s results imply that $\Omega(\log_d n)$ time is necessary in DetLOCAL; however, there are no known deterministic distributed LLL algorithms. It is conceivable that the distributed complexity of the LLL is very sensitive to the criterion used. We define $T_{LLL}(n, d)$ to be the RandLOCAL time to compute a point in the probability space avoiding all bad events (w.h.p.), under any "polynomial" LLL criterion of the form

$$pd^c = O(1), \tag{1}$$

where c can be an arbitrarily large constant. Prior results [7, 4] imply that $T_{LLL}(n, d)$ is $\Omega(\log_d \log n)$, $\Omega(\log^* n)$, and $O(\log_{1/epd^2} n)$. In this section we prove an automatic speedup theorem for sublogarithmic RandLOCAL algorithms. We do *not* assume that $\Delta = O(1)$ in this section.

Theorem 6. *Let \mathcal{A} be a RandLOCAL algorithm that solves some LCL problem \mathcal{P} w.h.p. Suppose that for any fixed $\epsilon > 0$, the running time $T_{\Delta}(n)$ of \mathcal{A} can be upper bounded by $C(\Delta) + \epsilon \log_{\Delta} n$, for some function C . It is possible to transform \mathcal{A} into a new RandLOCAL algorithm \mathcal{A}' that solves \mathcal{P} (w.h.p.) in $O(C(\Delta) \cdot T_{LLL}(n, \Delta^{O(C(\Delta))}))$ time.*

²⁵Each variable $V \in \mathcal{V}$ may have a different distribution and range, so long as the range is some finite set.

Proof. Suppose that \mathcal{A} has a local probability of failure $1/n$, that is, for any $v \in V(G)$, the probability that $N^r(v)$ is inconsistent with \mathcal{P} is $1/n$, where r is the radius of \mathcal{P} . Expressing the running time in terms of $\epsilon = O((2c)^{-1})$, where c is the constant in the LLL criterion, define n^* as the minimum value for which

$$t^* = T_\Delta(n^*) < (1/2c) \cdot \log_\Delta n^* - r.$$

It follows that $t^* = O(C(\Delta))$ and $n^* = \Delta^{O(C(\Delta))}$.

The algorithm \mathcal{A}' applied to an n -vertex graph G works as follows. Imagine an experiment where we run \mathcal{A} , but lie to the vertices, telling them that “ n ” = n^* . Any $v \in V(G)$ will see a t^* -neighborhood $N^{r+t^*}(v)$ that is consistent with some n^* -vertex graph. However, the *bad event* that $N^r(v)$ is incorrectly labeled is $1/n^*$, not $1/\text{poly}(n)$, as desired. We now show that this system of bad events satisfies the LLL criterion (1). Define the following events, graph, and quantities:

$$\begin{aligned} \mathcal{E}_v &: \text{the event that } N^r(v) \text{ is incorrectly labeled according to } \mathcal{P} \\ \mathcal{X} &= \{\mathcal{E}_v \mid v \in V(G)\} && \text{the set of bad events} \\ G_{\mathcal{X}} &= (\mathcal{X}, \{(\mathcal{E}_u, \mathcal{E}_v) \mid N^{r+t^*}(u) \cap N^{r+t^*}(v) \neq \emptyset\}) && \text{the dependency graph} \\ d &\leq \Delta^{2(r+t^*)} \\ p &= 1/n^* \end{aligned}$$

The event \mathcal{E}_v is determined by the labeling of $N^r(v)$ and the label of each $v' \in N^r(v)$ is determined by $N^{t^*}(v')$, hence \mathcal{E}_v is determined by (the data stored in, and random bits generated by) vertices in $N^{r+t^*}(v)$. Clearly \mathcal{E}_v is independent of any \mathcal{E}_u for which $N^{r+t^*}(u) \cap N^{r+t^*}(v) = \emptyset$, which justifies the definition of the edge set of $G_{\mathcal{X}}$. Since the maximum degree in G is Δ , the maximum degree d in $G_{\mathcal{X}}$ is less than $\Delta^{2(r+t^*)}$. By definition of \mathcal{A} , $\Pr(\mathcal{E}_v) \leq 1/n^* = p$. This system satisfies LLL criterion (1) since, by definition of t^* ,

$$pd^c = p\Delta^{2c(r+t^*)} < (1/n^*) \cdot n^* = 1.$$

The algorithm \mathcal{A}' now simulates a constructive LLL algorithm on $G_{\mathcal{X}}$ in order to find a labeling in $\bigcap_v \overline{\mathcal{E}_v}$. Since a virtual edge $(\mathcal{E}_u, \mathcal{E}_v)$ exists iff u and v are at distance at most $2(r+t^*) = O(C(\Delta))$, any RandLOCAL algorithm in $G_{\mathcal{X}}$ can be simulated in G with $O(C(\Delta))$ slowdown. Thus, \mathcal{A}' runs in $O(C(\Delta) \cdot T_{LLL}(n, \Delta^{O(C(\Delta))}))$ time. \square

Theorem 6 shows that when $\Delta = O(1)$, $o(\log n)$ -time RandLOCAL algorithms can be *sped up* to run in $O(T_{LLL}(n, O(1)))$ time. Another consequence of this same technique is that sublogarithmic RandLOCAL algorithms with *large messages* can be converted to (possibly slightly slower) algorithms with small messages. The statement of Theorem 7 reflects the use of a particular distributed LLL algorithm, namely [7]; it may be improvable using future distributed LLL technology.

Theorem 7. *Let \mathcal{A} be a $(C(\Delta) + \epsilon \log_\Delta n)$ -time RandLOCAL algorithm that solves some LCL problem \mathcal{P} with high probability. Each vertex locally generates $r_\Delta(n)$ random bits and sends $m_\Delta(n)$ -bit messages. It is possible to transform \mathcal{A} into a new RandLOCAL algorithm \mathcal{A}' that solves \mathcal{P} (w.h.p.) in $O(\log_\Delta n)$ time, where each vertex generates $O(\log n + r_\Delta(\zeta) \cdot \log_\zeta n)$ random bits, and sends $\min\{\log(|\Sigma_{\text{out}}|) \cdot \Delta^{O(1)} + m_\Delta(\zeta), r_\Delta(\zeta) \cdot \zeta\}$ -bit messages, where $\zeta = \Delta^{O(C(\Delta))}$ depends only on Δ .*

Proof. We continue to use the notation and definitions from Theorem 6, and fix $c = 3$ in the LLL criterion (1). Since $d = \Omega(\Delta^{O(C(\Delta))}) = \Omega(\zeta)$ and we selected t^* w.r.t. $c = 3$ (i.e., LLL criterion $pd^3 < 1$), we have $1/epd^2 = \Omega(\zeta)$. If \mathcal{A}' uses the LLL algorithm of [7], it will first generate $O(\log n)$ -bit unique IDs and generate $r_\Delta(n^*) \cdot O(\log_{1/epd^2} n) = O(r_\Delta(\zeta) \cdot \log_\zeta n)$ random bits throughout the computation. We give two methods to execute \mathcal{A}' that have different message complexities.

In the first method, before the LLL algorithm proper begins, each vertex v gathers up all IDs and random bits in its $(t^* + r)$ -neighborhood. This takes $O((\log n + r_\Delta(\zeta) \cdot \log_\zeta n) \cdot \zeta/b)$ time with b -bit messages. Each vertex u owns one random variable: an $r_\Delta(n^*)$ -bit string V_u . In order for v to tell whether \mathcal{E}_v occurs under the current variable assignment it only needs to know how many times each V_u , $u \in N^{r+t^*}(v)$, has been resampled. The actual execution of \mathcal{A} within $N^{r+t^*}(v)$ can then be simulated locally at v .²⁶ Thus, between

²⁶Whether the output labeling of v is locally consistent depends on the output labeling of vertices in $N^r(v)$. Given the graph topology and the random bits within $N^{r+t^*}(v)$, the vertex v can locally simulate \mathcal{A} and decides whether \mathcal{E}_v occurs.

resampling steps of the LLL algorithm, each u simply needs to alert its $(r + t^*)$ -neighborhood whether V_u is resampled or not. This can be accomplished in $O(C(\Delta))$ time with ζ -bit messages.

In the second method, vertices keep their random bits private. In order to tell which subset of bad events $\{\mathcal{E}_v\}_{v \in V(G)}$ occur under the current variable assignment, all vertices simulate \mathcal{A} for t^* rounds, sending $m_\Delta(n^*)$ -bit messages. After the simulation, for a vertex v to tell whether \mathcal{E}_v occurs, it needs to gather the output labeling of the vertices in $N^r(v)$. This can be done $O(1)$ rounds, sending $\log(|\Sigma_{\text{out}}|) \cdot \Delta^{O(1)}$ -bit messages.²⁷ The resampling steps themselves can be implemented with constant-size messages and $O(C(\Delta))$ time.

There are $O(\log_{1/\epsilon p d^2} n) = O(\log_\zeta n) = O(\frac{\log_\Delta n}{C(\Delta)})$ resampling steps, each of which is simulated in $O(C(\Delta))$ time, hence the total time is $O(\log_\Delta n)$, independent of the function C . \square

An interesting corollary of Theorem 7 is that when $\Delta = O(1)$, randomized algorithms with unbounded length messages can be simulated with 1-bit messages.

Corollary 1. *Let \mathcal{P} be any LCL problem. When $\Delta = O(1)$, any $o(\log n)$ algorithm solving \mathcal{P} w.h.p. using unbounded length messages can be made to run in $O(\log n)$ time with 1-bit messages.*

5 Conclusion

We now have a very good understanding of the LOCAL complexity landscape for cycles, tori, bounded degree trees, and to a lesser extent, general bounded degree graphs. See Figure 1. However, there are some very critical gaps in our understanding.

Our randomized speedup theorem of Section 4 depends on the complexity of a relatively weak version of the Lovász local lemma. Since the LLL is essentially a “complete” problem for sublogarithmic RandLOCAL algorithms, understanding the distributed complexity of the LLL is a significant open problem.

Conjecture 1. *There exists a sufficiently large constant c such that the distributed LLL problem can be solved in $O(\log \log n)$ time on bounded degree graphs, under the symmetric LLL criterion $pd^c < 1$.*

The new polynomial complexities introduced in Section 2 are of the form $\Theta(n^{1/k})$, $k \in \mathbb{Z}^+$. Is this set of polynomial complexities exhaustive? Is it possible to engineer problems with complexity $\Theta(n^q)$ for any given rational q ? We think the answer is no, and resolving Conjecture 2 would be the first step.

Conjecture 2. *Any $o(n)$ -time DetLOCAL algorithm can be automatically sped up to run in $O(\sqrt{n})$ time. In general, there is an $\omega(n^{1/(k+1)}) - o(n^{1/k})$ gap in the DetLOCAL complexity hierarchy.*

One advantage of working with graph classes with bounded degree Δ is that arbitrary LCLs can be encoded in $O(1)$ space by enumerating all acceptable configurations. To study LCLs on unbounded-degree graphs it is probably necessary to work with logical representations of LCLs, and here the expressive power of the logic may introduce a new measure of problem complexity. For example, the sentence “ I is an MIS” can be expressed using the predicate $[v \in I]$ and quantification over all vertices $v \in V(G)$, and all vertices $u \in N(v)$.

References

- [1] D. Achlioptas and F. Iliopoulos. Random walks that find perfect objects and the Lovász local lemma. In *Proceedings 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 494–503, 2014.
- [2] R. Bar-Yehuda, K. Censor-Hillel, and G. Schwartzman. A distributed $(2 + \epsilon)$ -approximation for vertex cover in $O(\log \Delta / \epsilon \log \log \Delta)$ rounds. In *Proceedings 35th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 3–8, 2016.

²⁷An output label can be encoded as a $\log(|\Sigma_{\text{out}}|)$ -bit string. We do not assume that Δ is constant so $|\Sigma_{\text{out}}|$, which may depend on Δ but not directly on n , is also not constant. E.g., consider the $O(\Delta)$ vertex coloring problem.

- [3] L. Barenboim, M. Elkin, S. Pettie, and J. Schneider. The locality of distributed symmetry breaking. *J. ACM*, 63, 2016. Article 20.
- [4] 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 *Proceedings 48th ACM Symposium on the Theory of Computing*, pages 479–488, 2016.
- [5] S. Brandt, J. Hirvonen, J. H. Korhonen, T. Lempiäinen, P. R. J. Östergård, C. Purcell, J. Rybicki, J. Suomela, and P. Uznanski. LCL problems on grids. *CoRR*, abs/1702.05456, 2017.
- [6] Y.-J. Chang, T. Kopelowitz, and S. Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. In *Proceedings 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 615–624, 2016.
- [7] K.-M. Chung, S. Pettie, and H.-H. Su. Distributed algorithms for the Lovász local lemma and graph coloring. *Distributed Computing*, 2017. to appear.
- [8] R. Cole and U. Vishkin. Deterministic coin tossing with applications to optimal parallel list ranking. *Information and Control*, 70(1):32–53, 1986.
- [9] L. Feuilloley and P. Fraigniaud. Survey of distributed decision. *CoRR*, abs/1606.04434, 2016.
- [10] P. Fraigniaud, A. Korman, and D. Peleg. Towards a complexity theory for local distributed computing. *J. ACM*, 60(5):35:1–35:26, 2013.
- [11] M. Fürer. Data structures for distributed counting. *J. Comput. Syst. Sci.*, 28(2):231–243, 1984.
- [12] M. Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 270–277, 2016.
- [13] M. Ghaffari, F. Kuhn, and Y. Maus. On the complexity of local distributed graph problems. In *Proceedings 49th ACM Symposium on Theory of Computing (STOC)*, 2017.
- [14] M. Ghaffari and H.-H. Su. Distributed degree splitting, edge coloring, and orientations. In *Proceedings 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2505–2523, 2017.
- [15] M. Göös, J. Hirvonen, and J. Suomela. Linear-in- Δ lower bounds in the LOCAL model. *Distributed Computing*, pages 1–14, 2015.
- [16] M. Göös and J. Suomela. Locally checkable proofs in distributed computing. *Theory of Computing*, 12(1):1–33, 2016.
- [17] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey Theory*. John Wiley and Sons, New York, 2nd edition, 1990.
- [18] B. Haeupler and D. G. Harris. Parallel algorithms and concentration bounds for the Lovász local lemma via witness-DAGs. In *Proceedings 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1170–1187, 2017.
- [19] D. G. Harris. Lopsidedependency in the Moser-Tardos framework: Beyond the lopsided Lovász local lemma. *ACM Trans. Algorithms*, 13(1):17:1–17:26, 2016.
- [20] D. G. Harris and A. Srinivasan. A constructive algorithm for the Lovász local lemma on permutations. In *Proceedings 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 907–925, 2014.
- [21] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Trans. Amer. Math. Soc.*, 117:285–306, 1965.

- [22] N. J. A. Harvey and J. Vondrák. An algorithmic proof of the Lovász local lemma via resampling oracles. In *Proceedings 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1327–1346, 2015.
- [23] D. Hefetz, F. Kuhn, Y. Maus, and A. Steger. Polynomial lower bound for distributed graph coloring in a weak LOCAL model. In *Proceedings 30th International Symposium on Distributed Computing (DISC)*, pages 99–113, 2016.
- [24] K. B. R. Kolipaka and M. Szegedy. Moser and Tardos meet Lovász. In *Proceedings 43rd ACM Symposium on Theory of Computing (STOC)*, pages 235–244, 2011.
- [25] V. Kolmogorov. Commutativity in the algorithmic Lovász local lemma. In *Proceedings 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 780–787, 2016.
- [26] A. Korman, J.-S. Sereni, and L. Viennot. Toward more localized local algorithms: removing assumptions concerning global knowledge. *Distributed Computing*, 26(5–6):289–308, 2013.
- [27] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Local computation: Lower and upper bounds. *J. ACM*, 63(2):17:1–17:44, 2016.
- [28] F. Kuhn and R. Wattenhofer. On the complexity of distributed graph coloring. In *Proceedings 25th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 7–15, 2006.
- [29] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [30] G. L. Miller and J. H. Reif. Parallel tree contraction—Part I: Fundamentals. *Advances in Computing Research*, 5:47–72, 1989.
- [31] R. A. Moser and G. Tardos. A constructive proof of the general lovász local lemma. *J. ACM*, 57(2), 2010.
- [32] M. Naor. A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM J. Discrete Mathematics*, 4(3):409–412, 1991.
- [33] M. Naor and L. J. Stockmeyer. What can be computed locally? *SIAM J. Comput.*, 24(6):1259–1277, 1995.
- [34] A. Panconesi and A. Srinivasan. On the complexity of distributed network decomposition. *J. Algor.*, 20(2):356–374, 1996.
- [35] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [36] J. Suomela. Survey of local algorithms. *ACM Comput. Surv.*, 45(2):24:1–24:40, 2013.

A Speedup Implications of Naor & Stockmeyer

Let \mathcal{A} be any $T(n)$ -round DetLOCAL algorithm. Let η and η' be any two *order-indistinguishable* assignments of distinct IDs to $N^{T(n)}(v)$, i.e., for $u, w \in N^{T(n)}(v)$, $\eta(u) > \eta(w)$ if and only if $\eta'(u) > \eta'(w)$. If, for every possible input graph fragment induced by $N^{T(n)}(v)$, the output label of v is identical under every pair of order-indistinguishable η, η' , then \mathcal{A} is *order-invariant*.

Suppose that there exists a number $n' = O(1)$ such that $\Delta^{T(n)+r} \leq n'$. If \mathcal{A} is order-invariant then it can be turned into an $O(1)$ -round DetLOCAL algorithm \mathcal{A}' , since we can pretend that the total number of vertices is n' instead of n .

In [33], Naor and Stockmeyer proved that any DetLOCAL algorithm that takes $\tau = O(1)$ rounds on a bounded degree graph can be turned into an order-invariant τ -round DetLOCAL algorithm. A more careful analysis shows that the proof still works when τ is a slowly growing function of n .

Requirement for automatic speedup. The multicolor hypergraph Ramsey number $R(p, m, c)$ is the minimum number such that the following holds. Let H be a p -uniform complete hypergraph of at least $R(p, m, c)$ vertices. Then any c -edge-coloring of H contains a monochromatic clique of size m .

Given the number $\tau \geq 2$, the three parameters p, m , and c are selected as follows (see the proof of [33, Lemma 3.2] for more details).

- The number p is the maximum number of vertices in $N^\tau(v)$, over all vertices $v \in V(G)$ and all graphs G under consideration. For rings, $p = 2\tau + 1$. For tori, $p \leq 2(\tau + 1)^2$. For trees or general graphs, $p \leq \Delta^\tau$.
- The number m is the maximum number of vertices in $N^{\tau+r}(v)$, over all vertices $v \in V(G)$ and all graphs G under consideration. E.g., for rings, $p = 2\tau + 2r + 1$ and for general graphs, $p \leq \Delta^{\tau+r}$.
- The number z counts the distinguishable radius- τ centered subgraphs, disregarding IDs. For example, for LCLs on the ring without input labels or port numbering, $z = 1$, whereas with input labels and port numbering it is $(2|\Sigma_{\text{in}}|)^{2\tau+1}$ since each vertex has one of $|\Sigma_{\text{in}}|$ input labels and 2 port numberings. In general z is certainly at most $2^{\binom{2\tau}{2}} \cdot (\Delta!|\Sigma_{\text{in}}|)^p$.
- The number $c = |\Sigma_{\text{out}}|^{p!z}$ counts the number of output labelings of the z radius- τ centered subgraphs, that are equipped with unique vertex IDs in the range $\{1, \dots, p\}$.

Suppose that the ID space is $[n^k]$. Naor and Stockmeyer's proof implies that, as long as $n^k \geq R(p, m, c)$, any DetLOCAL τ -round algorithm on a bounded degree graph can be turned into an order-invariant τ -round DetLOCAL algorithm, which then implies an $O(1)$ -round DetLOCAL algorithm.

The Ramsey number $R(p, m, c)$. According to the proof of [17, §1, Theorem 2], we have:

$$\text{For } p = 1, \quad R(p, m, c) = c(m - 1) + 1$$

$$\text{For } p > 1, \quad R(p, m, c) \leq 2c^x \quad \text{where } x = \sum_{i=p-1}^{R(p-1, m, c)-1} \binom{i+1}{p-1} < R(p-1, m, c)^p$$

Therefore, $\log^*(R(p, m, c)) \leq p + \log^* m + \log^* c + O(1)$.

Automatic speedup theorems. Observe that in all scenarios, if the running time $\tau = \tau(n) = \omega(1)$, we have $\log^* m + \log^* c = o(p)$. Therefore, having $p \leq \epsilon \log^* n$ for some small enough constant ϵ suffices to meet the condition $n^k \geq R(p, m, c)$. We conclude that the complexity of any LCL problem (with or without input labels and port numbering) in the LOCAL model never falls in the following gaps:

$$\begin{array}{ll} \omega(1) - o(\log^* n) & \text{for } n\text{-rings.} \\ \omega(1) - o(\sqrt{\log^* n}) & \text{for } (\sqrt{n} \times \sqrt{n})\text{-tori.} \\ \omega(1) - o(\log(\log^* n)) & \text{for bounded degree trees or general graphs.} \end{array}$$

Due to the “Stepping-Up Lemma” (see [17, §4, Lemma 17]), we have $\log^*(R(p, m, 2)) = \Omega(p)$. Thus, Naor and Stockmeyer’s approach *alone* cannot give an $\omega(1) - o(\log^* n)$ gap for general graphs.

However, for the case of LCL problems on $(\sqrt{n} \times \sqrt{n})$ -tori that does not use input label and port-numbering, the gap can be widened to $\omega(1) - o(\log^* n)$ [5, p. 2]. A proof sketch is as follows (due to Jukka Suomela, personal communication).

Given a $(\sqrt{n} \times \sqrt{n})$ -torus G . We associate each vertex $v \in V(G)$ with a coordinate (α, β) , where $\alpha, \beta \in [\sqrt{n}]$. We consider the following special way to generate unique IDs. Let ϕ_x and ϕ_y be two functions mapping integers in $[\sqrt{n}]$ to integers in $\{0, \dots, n^{k/2} - 1\}$. The ID of the vertex v (with coordinate (α, β)) is defined as $\phi_x(\alpha) \cdot n^{k/2} + \phi_y(\beta)$. Notice that the IDs of all vertices in $N^\tau(v)$ can be deduced from these $4\tau + 2$ numbers $\phi_x(i)$, $i \in [\alpha - \tau, \alpha + \tau]$, and $\phi_y(j)$, $j \in [\beta - \tau, \beta + \tau]$. Therefore, the parameter p is “effectively” $4\tau + 2$, not $\Theta(\tau^2)$.

Suppose that there is a $\tau = o(\log^* n)$ -round algorithm \mathcal{A} for \mathcal{P} on G . Then \mathcal{A} must work correctly with the above special ID assignment. Naor and Stockmeyer’s proof implies that there is an order-invariant algorithm \mathcal{A}' which takes τ rounds (under the aforementioned special ID assignment). Consider the assignment $\phi_x(i) = i$ for all i , and $\phi_y(j) = j$ for all j . Since \mathcal{A}' is order-invariant, there must exist an element $\sigma \in \Sigma_{\text{out}}$ such that for all (α, β) with $\tau + r < \alpha, \beta < \sqrt{n} + 1 - \tau - r$, the algorithm \mathcal{A}' labels the vertex with coordinate (α, β) by σ .²⁸ Thus, labeling all vertices by σ yields a legal labeling of a $(\sqrt{n} \times \sqrt{n})$ -torus.

We conclude that any LCL problem \mathcal{P} (without input labeling) on $(\sqrt{n} \times \sqrt{n})$ -tori solvable in $o(\log^* n)$ rounds must be *trivial* in the sense that there exists an element $\sigma \in \Sigma_{\text{out}}$ such that labeling all vertices by σ gives a legal labeling (assume $\sqrt{n} > 2r + 1$). Thus, \mathcal{P} can be solved in $O(1)$ rounds.

It still remains an outstanding open problem whether the gap for other cases can also be widened to $\omega(1) - o(\log^* n)$.

²⁸This is the place where we use the requirement that the LCL problem does not use input label and port-numbering. Even if ϕ_x and ϕ_y of the coordinates are monotonically increasing, an order-invariant algorithm is still possible to “break symmetry” by utilizing port-numbering or input labeling. Notice that if the $(\sqrt{n} \times \sqrt{n})$ -torus is *oriented*, then the requirement that the LCL problem does not refer to port-numbering is unnecessary.