SIAM J. COMPUT.
Vol. 48, No. 1, pp. 33–69
© 2019 Society for Industrial and Applied Mathematics

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

A TIME HIERARCHY THEOREM FOR THE LOCAL MODEL*

YI-JUN CHANG† AND SETH PETTIE†

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: (a) 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. (b) 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 [Y.-J. Chang, T. Kopelowitz, and S. Pettie, *Proceedings of FOCS*, 2016, pp. 615–624], this establishes that on trees, there are no natural deterministic complexities in the ranges $\omega(\log^* n)$—$o(\log n)$ or $\omega(\log n)$—$n^{o(1)}$. (c) We expose a new 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: the complexity of the distributed Lovász local lemma (LLL) problem. In other words, the LLL is *complete* for sublogarithmic time. 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 an $\omega(1)$—$o(\log^* n)$ 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.

Key words. distributed local model, local checkable labeling, Lovász local lemma, time hierarchy theorem

AMS subject classifications. 05C85, 68W15

DOI. 10.1137/17M1157957

1. Introduction.

The goal of this paper is to understand the spectrum of natural problem complexities that can exist in the LOCAL model [38, 43] of distributed computation and to quantify the value of randomness in this model. Whereas the time hierarchy of Turing machines is known[1] to be very "dense," recent work [9, 7] 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*.

*Received by the editors November 21, 2017; accepted for publication (in revised form) October 16, 2018; published electronically January 3, 2019. A preliminary version of this paper appeared in *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, 2017.
http://www.siam.org/journals/sicomp/48-1/M115795.html
Funding: This work was supported by NSF grants CCF-1514383 and CCF-1637546.
†University of Michigan, Ann Arbor, MI 48109 (cyijun@umich.edu, pettie@umich.edu).
[1]For *any* time-constructible function $T(n)$, there is a problem solvable in $O(T(n))$ but not $o(T(n))$ time [30, 17].

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

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 $\{1, \ldots, \deg(v)\}$, certain global parameters such as $n \stackrel{\text{def}}{=} |V|$, $\Delta \stackrel{\text{def}}{=} \max_{v \in V} \deg(v)$, and possibly other information. The assumption that global parameters are common knowledge can sometimes be removed; see Korman, Sereni, and Viennot [35]. 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 $\deg(v), \mathrm{ID}(v)$, 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, but 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 fix the maximum tolerable global probability of failure to be $1/n$.

Clearly RandLOCAL algorithms can generate distinct IDs (with high probability (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 [42] 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 $r = O(1)$, input and output alphabets $\Sigma_{\text{in}}, \Sigma_{\text{out}}$ (which can depend on Δ but not n), and a set $\mathcal{C}$ of acceptable configurations. 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 $\phi_{\text{in}} : V(G) \to \Sigma_{\text{in}}$, an acceptable output is any function $\phi_{\text{out}} : V(G) \to \Sigma_{\text{out}}$ 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 $\mathcal{C}$.

For bounded degree graphs, an LCL can be described explicitly by enumerating a finite number of acceptable configurations. For graph classes with unbounded degrees, LCLs can be defined through logic expression. 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 an almost perfect understanding of the complexity landscape for two simple topologies: n-cycles/paths [12, 38, 41, 42, 9] and $(\sqrt{n} \times \sqrt{n})$-grids/tori [42, 9, 7]. See

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Figure 1, top and middle. On the n-cycle/path, 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 [42] characterization of $O(1)$-time LCL algorithms actually implies that any $o(\log^* n)$-time algorithm on the n-cycle/path can be transformed to run in $O(1)$ time; see Appendix A. Chang, Kopelowitz, and Pettie [9] showed that any $o(n)$-time RandLOCAL algorithm can be made to run in $O(\log^* n)$ time in DetLOCAL.

The situation with $(\sqrt{n} \times \sqrt{n})$-grids/tori is almost identical [7]: 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 [42] is $\omega(1)$—$o(\log^* n)$ on the n-cycle/path, it is $\omega(1)$—$o(\sqrt{\log^* n})$ on the $(\sqrt{n} \times \sqrt{n})$-torus; see Appendix A.[2] Whereas randomness is known not to help in n-cyles/paths [42, 9], it is an open question on grids/tori [7]. Whereas the classification question (whether an LCL is $O(\log^* n)$ or $\Omega(n)$) is decidable on n-cycles/paths, the same question is *undecidable* on $(\sqrt{n} \times \sqrt{n})$-grids/tori [42, 7].

The gap theorems of Chang, Kopelowitz, and Pettie [9] 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 ($O(\log_\Delta \log n)$ versus $\Omega(\log_\Delta n)$) between their RandLOCAL and DetLOCAL complexities, such as Δ-coloring degree-Δ trees [6, 9, 44], sinkless orientation [6, 21], and $(2\Delta - 2)$-edge coloring trees [8]. More generally, Chang, Kopelowitz, and Pettie [9] 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 nonconstant Δ, the RandLOCAL lower bounds of Kuhn, Moscibroda, and Wattenhofer [36] $\Omega(\min\{\frac{\log \Delta}{\log \log \Delta}, \sqrt{\frac{\log n}{\log \log n}}\})$ 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 [4]; the best maximal matching [5] and MIS [18] 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 [5], e.g., trees or planar graphs.

New results. In this paper we study the LOCAL complexity landscape on bounded degree trees and bounded degree general graphs; see Figure 1. We establish a new (deterministic and randomized) complexity gap for bounded degree trees, a new randomized complexity gap for general graphs, and a new infinite hierarchy of coloring problems with polynomial time complexities:

- We prove that on the class of bounded degree 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. That is, we cannot extend it lower than $\omega(\log n)$ [38, 9] 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

[2] Suomela [46] has a proof that there is an $\omega(1)$—$o(\log^* n)$ complexity gap for grids/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.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

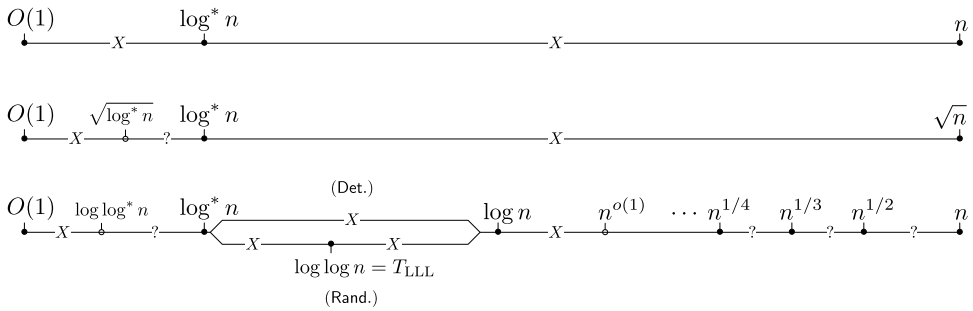

FIG. 1. *Top: The complexity landscape for LCL problems on the n-cycle/path. Middle: The complexity landscape for LCL problems on the* $(\sqrt{n} \times \sqrt{n})$*-grid/torus. Refer to* [42, 9, 7] *and Appendix* A *for proofs of the complexity gaps ("X") on paths/cycles and grids/tori. Bottom: The complexity 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* [9]*. The* $\omega(T_{LLL})$—$o(\log n)$ *and* $\omega(\log n)$—$n^{o(1)}$ *gaps are new. Recent results* [8, 15] *have put* $T_{LLL} = \Theta(\log \log n)$ *on trees. 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 on trees. With the exception of the* $\omega(\log n)$—$n^{o(1)}$ *gap, all known complexity gaps on bounded degree trees apply to bounded degree general graphs as well; however, the exact complexity of the Lovász local lemma (LLL) on general graphs has not been settled.*

upper bound holds in DetLOCAL on general graphs, and the lower bound holds even on degree-3 trees in RandLOCAL. Thus, in contrast to paths/cycles and grids/tori, trees and general graphs support an *infinite* number of natural problem complexities.

- Suppose we have a RandLOCAL algorithm for general graphs running in $C(\Delta) + \epsilon \log_\Delta n$ time for any desired $\epsilon > 0$ and some function C.[3] 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 LLL. At present, T_{LLL} is known to be $\Omega(\log \log n)$ [6] even on trees [8]. This establishes a new RandLOCAL complexity gap between $\omega(T_{LLL})$ and $o(\log n)$.

Finally, it seems to be folklore that Naor and Stockmeyer's work [42] implies *some kind* of complexity gap, which has been cited as $\omega(1)$—$o(\log^* n)$ [7, 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 cycles/paths,
- $\omega(1)$—$o(\sqrt{\log^* n})$ on grids/tori,
- $\omega(1)$—$o(\log(\log^* n))$ 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. Port 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. Suomela [46] gave a proof that the $\omega(1)$—$o(\log^* n)$ gap exists in grids/tori as well, for the class of LCL problems without input labels or port numbering. This proof is reproduced in Appendix A.

Commentary. All the existing automatic speedup theorems are quite different in terms of proof techniques. Naor and Stockmeyer's approach is based on Ramsey

[3] This is a convoluted way of saying "sublogarithmic time." Because of the nature of the proof, we care about what the time complexity is when n is small, not just when $n \to \infty$.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

theory. The speedup theorems of [9, 7] use the fact that $o(\log_\Delta n)$ algorithms on general graphs (and $o(n)$ algorithms on n-cycles/paths and $o(\sqrt{n})$ algorithms on $(\sqrt{n}\times\sqrt{n})$-grids/tori) cannot "see" the whole graph and can therefore be efficiently tricked into thinking the graph has constant size. Our $n^{o(1)} \to O(\log n)$ 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 [22] proved a sharp $\Omega(\Delta)$ lower bound for fractional maximal matching and Göös and Suomela [24] proved $\Omega(\log n)$ lower bounds on $(1+\delta)$-approximating the minimum vertex cover, $\delta > 0$, even on degree-3 graphs. See [37, 32] 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 in distributed computing. The gap theorems of [42, 9, 7] have already been discussed. Suomela surveys [45] the class of problems that can be computed with $O(1)$ time. Fraigniaud, Korman, and Peleg [16] defined a distributed model for locally deciding graph properties; see [13] for a survey of variants of the local distributed decision model. Göös and Suomela [23] considered the *proof* complexity (measured in terms of bits-per-vertex label) of locally verifying graph properties. Very recently, Ghaffari, Kuhn, and Maus [20] defined the SLOCAL model (sequential LOCAL) and exhibited several *complete* problems for this model, inasmuch as a polylog(n)-time DetLOCAL algorithm for any complete problem implies a polylog(n) DetLOCAL algorithm for every polylog(n)-time problem in SLOCAL.[4]

1.3. Recent developments. The preliminary version of this work [10] concluded with two conjectures, one on the complexity of the distributed LLL under a "polynomial" LLL criterion, and one on further gaps in the LOCAL complexity hierarchy. Subsequent work by Fischer and Ghaffari [15], Chang et al. [8], Ghaffari, Harris, and Kuhn [19], and Balliu et al. [3] has offered compelling evidence in favor of the first conjecture and disproved the second conjecture.

Fischer and Ghaffari [15] gave a deterministic LLL algorithm with complexity $O(n^{1/\lambda+O(1/\sqrt{\log n})})$ under criterion $p(ed)^\lambda < 1$ and a randomized algorithm with complexity $O(d^2 + (\log n)^{1/\lambda+O(1/\sqrt{\log\log n})})$ under criterion $p(ed)^{4\lambda} < 1$. (See section 4 for the definition of p, d and a discussion of LLL criteria.) When $d < (\log\log n)^{1/5}$, they improved their randomized algorithm to $2^{O(\sqrt{\log\log n})}$ time under criterion $p(ed)^{32} < 1$. Under criterion $p(ed)^{d^2} < 1$, the LLL can be solved in $O(d^2 + \log^* n)$ time [15], which matches the Chung–Pettie–Su [11] lower bound $\Omega(\log^* n)$ in terms of n and gives a new proof [15, Corollary 3] of the $\omega(\log^* n)$—

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

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

$o(\log\log n)$ RandLOCAL complexity gap [9] on bounded degree graphs. Chang et al. [8] designed special LLL algorithms for "tree structured" dependency graphs.[5] Under criterion $p(ed)^\lambda < 1$, they run in $O(\max\{\log_\lambda n, \log n/\log\log n\})$ time in DetLOCAL and $O(\max\{\log_\lambda \log n, \log\log n/\log\log\log n\})$ time in RandLOCAL, with no dependency on d. This work confirmed [10, Conjecture 1] for the special case of trees. In [19], the upper bound for T_{LLL} on bounded degree general graphs was further improved to $\exp^{(i)}(c_i\sqrt{\log^{(i+1)} n})$, for any i, where c_i depends only on i and $\exp^{(i)}, \log^{(i+1)}$ are iterated i-fold applications of exp and log, respectively.

Balliu et al. [3] disproved [10, Conjecture 2] and showed that on bounded degree general graphs, the complexity hierarchy is very dense in essentially every region left open by this work. In particular, there are an infinite number of LCL problem complexities between $\Omega(\log(\log^* n))$ and $O(\log^* n)$, an infinite number of complexities between $\Omega(\log n)$ and $n^{o(1)}$ (provably distinguishing the complexity hierarchies for trees and general graphs), and an infinite number of complexities of the form $\Theta(n^r)$ for rationals r not of the form $1/k$. Whether bounded degree trees can support the first and third categories is still open.

1.4. 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 $n^{o(1)} \to O(\log n)$ speedup theorem for bounded degree trees. In section 4 we discuss the constructive LLL and prove the $o(\log_\Delta n) \to T_{LLL}$ 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 $\Theta(n^{1/k})$.[6] 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 $\{\mathbf{a}, \mathbf{b}\}$) certain specific subgraphs of the input graph. Some vertices are *exempt* from being colored (in which case they are labeled $\mathbf{X}$), and in addition, it is possible to *decline* to 2-color certain subgraphs, by labeling them $\mathbf{D}$.

There are no input labels. The output label set is $\Sigma_{\text{out}} = \{\mathbf{a}, \mathbf{b}, \mathbf{D}, \mathbf{X}\}$. 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 level i, defined as follows:

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

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 $N^k(v)$.

[5] If T is a tree and $r = O(1)$, the graph $T^r = (V(T), \{\{u, v\} \mid \text{dist}_T(u, v) \le r\})$ is tree structured.

[6] Brandt et al. [7, 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 $\Theta(n^{1/k})$.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Exemption. A vertex labeled $\mathbf{X}$ is called *exempt*. No V_1 vertex is labeled $\mathbf{X}$; all V_{k+1} vertices are labeled $\mathbf{X}$. Any V_i vertex is labeled $\mathbf{X}$ if and only if it is adjacent to a lower level vertex labeled $\mathbf{a}, \mathbf{b}$, or $\mathbf{X}$. Define $X_i \subseteq V_i$ to be the set of level i exempt vertices.

Two-coloring. Vertices not covered by the exemption rule are labeled one of $\mathbf{a}, \mathbf{b}, \mathbf{D}$.

- Any vertex in V_i, $i \in [1, k]$, labeled $\mathbf{a}$ has no neighbor in V_i labeled $\mathbf{a}$ or $\mathbf{D}$.
- Any vertex in V_i, $i \in [1, k]$, labeled $\mathbf{b}$ has no neighbor in V_i labeled $\mathbf{b}$ or $\mathbf{D}$.
- Any vertex in $V_k - X_k$ with exactly 0 or 1 neighbors in $V_k - X_k$ must be labeled $\mathbf{a}$ or $\mathbf{b}$.

Commentary. The level rule implies that the graph induced by V_i consists of paths and cycles. The two-coloring rule implies that each component of nonexempt vertices in the graph induced by $V_i - X_i$ must either (a) be labeled uniformly by $\mathbf{D}$ or (b) be properly 2-colored by $\{\mathbf{a}, \mathbf{b}\}$. Every *path* in $V_k - X_k$ must be properly 2-colored, but *cycles* in $V_k - X_k$ are allowed to be labeled uniformly by $\mathbf{D}$. 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 $V_k - X_k$ are bipartite.

Remark 2.1. 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 $i \in [2, k]$ if and only if all but at most 2 neighbors have levels less than i.

THEOREM 2.2. *The* DetLOCAL *complexity of* $\mathcal{P}_k$ *on general graphs is* $O(n^{1/k})$.

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

- Compute X_i according to the exemption rule (e.g., $X_1 = \emptyset$, $X_{k+1} = V_{k+1}$).
- Each path in the subgraph induced by $V_i - X_i$ calculates its length. If it contains at most $\lceil 2n^{1/k} \rceil$ vertices, it properly 2-colors itself with $\{\mathbf{a}, \mathbf{b}\}$; longer paths and cycles in $V_i - X_i$ label themselves uniformly by $\mathbf{D}$.

This algorithm correctly solves $\mathcal{P}_k$ provided that it never labels a path in $V_k - X_k$ with $\mathbf{D}$. Let U_i be the subgraph induced by those vertices in $V_1 \cup \cdots \cup V_i$ labeled $\mathbf{D}$. 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 $\mathbf{D}$, 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 $\mathbf{D}$. ∎

THEOREM 2.3. *The* RandLOCAL *complexity of* $\mathcal{P}_k$ *on trees with maximum degree* $\Delta = 3$ *is* $\Omega(n^{1/k})$.

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*.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

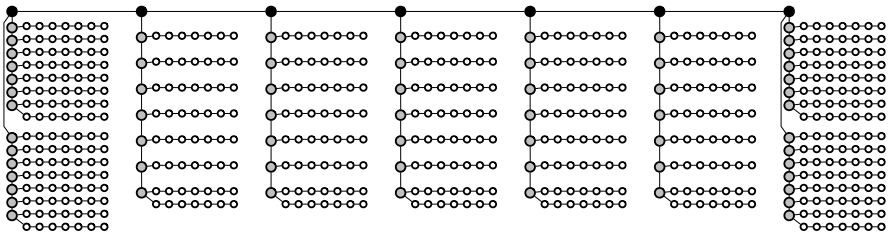

FIG. 2. *The graph H_k with parameters $k = 3, x = 7$. White vertices are in V_1, gray in V_2, and black in V_3. $V_4 = V_{k+1}$ is empty.*

- 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 $(v_1, v_2, \ldots, v_x)$, with head v_1 and tail v_x. Form $x+1$ copies $(H_{i-1}^{(j)})_{1 \le j \le x+1}$ 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 $k = 3$.

We 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 $(V_1, \ldots, V_k, V_{k+1})$, 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 $p_{\text{fail}} \le 1/|V(H_k)|$. However, we will now show that p_{fail} is constant, independent of $|V(H_k)|$.

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

Base case. We first prove that

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

Conditioning on $\mathcal{E}_2$ means that $X_2 \neq \emptyset$. Fix any $v \in X_2$ and let P be a copy of H_1 (a path) adjacent to v. In order for $v \in X_2$, it must be that P is properly 2-colored with $\{\mathbf{a}, \mathbf{b}\}$. 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 $p_{\mathbf{a}}$ and $p_{\mathbf{b}}$ be the probabilities that u/u' is labeled $\mathbf{a}$ and $\mathbf{b}$, respectively. A proper 2-coloring of P assigns u and u' different colors, and that occurs with probability $2p_{\mathbf{a}}p_{\mathbf{b}} \le 2p_{\mathbf{a}}(1 - p_{\mathbf{a}}) \le 1/2$. 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 $p_{\text{fail}} \ge p_2/2$, and hence $p_2 \le 2 \cdot p_{\text{fail}}$.

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

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

We are conditioning on $\mathcal{E}_i \backslash \mathcal{E}_{i-1}$. If this event is empty, then $p_i \le p_{i-1} \le 2(i-2) \cdot p_{\text{fail}}$ and the induction is complete. On the other hand, if $\mathcal{E}_i \backslash \mathcal{E}_{i-1}$ holds, then there is some

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

$v \in X_i$ adjacent to a copy of H_{i-1} with backbone path P, where $P \cap X_{i-1} = \emptyset$. In other words, if H_k is colored according to $\mathcal{P}_k$, then P must be properly 2-colored with $\{\mathbf{a}, \mathbf{b}\}$. The argument above shows this occurs with probability at least 1/2. Thus,

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

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 ($P \cap X_k \neq \emptyset$) is $p_k \leq 2(k-1) \cdot p_{\text{fail}}$. On the other hand, we have $\Pr(H_k \text{ not colored correctly} \mid \overline{\mathcal{E}_k}) \geq 1/2$ 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 w.h.p. ∎

3. A complexity gap on bounded degree trees. In this section we prove an $n^{o(1)} \to O(\log n)$ speedup theorem for LCL problems on bounded degree trees. The progression of definitions and lemmas in sections 3.2–3.13 is *logical* but obscures 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 [39]. 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 (see Figure 3): level-i vertices are arranged in an $O(1)$-length path, each the root of a subtree of level-($< i$) vertices (colored subtrees in the figure) that were removed in previous rake/compress steps, and bookended by level-($> i$) vertices (black in the figure). Call the subgraph between the bookends H.

In our approach it is the level-($> i$) 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 have only 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 the vertices of 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.

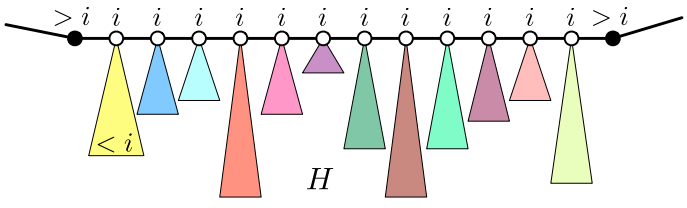

FIG. 3. *A constant length path resulting from spliting up long degree-2 paths of level-i vertices.*

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

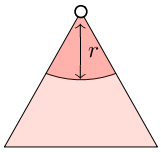

FIG. 4. *Class of a rooted tree.*

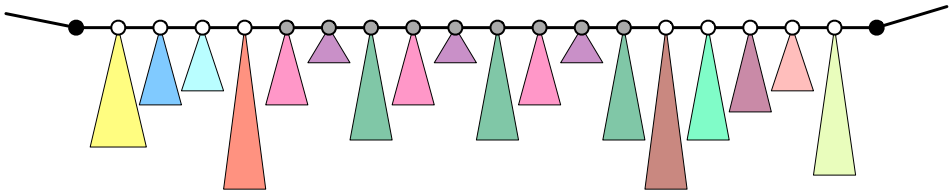

FIG. 5. *Pumping lemma for trees.*

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 $\mathcal{P}$ (see Figure 4).

In other words, the large and complex graph H can be succinctly encoded as a simple class vector $(c_1, c_2, \ldots, c_\ell)$, where c_j is the class of the jth 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 $(c_1, \ldots, c_\ell)$ one character at a time. By the pigeonhole principle, if ℓ is sufficiently large, then the automaton loops, meaning that $(c_1, \ldots, c_\ell)$ 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.

Figure 5 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 *precommit* to a labeling of a small (radius-r) subgraph of pumped-H. Given this precommitment, the left and right bookends no longer need to coordinate their activities: everything left (right) of the precommitted zone is now in the purview of the left (right) bookend. Interestingly, these manipulations (tree surgery and precommitments) 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

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

the notion of a partially labeled graph, i.e., one with certain vertices precommited 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 precommits a small partial labeling) 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 3.12 shows that $\mathcal{P}$ can be solved in $O(\log n)$ time, given a *feasible labeling function* f. Section 3.13 shows how a feasible f can be extracted from any $n^{o(1)}$-time algorithm $\mathcal{A}$.

3.2. Partially labeled graphs. A *partially labeled graph* $\mathcal{G} = (G, \mathcal{L})$ is a graph G together with a function $\mathcal{L} : V(G) \to \Sigma_{\text{out}} \cup \{\perp\}$. The vertices in $\mathcal{L}^{-1}(\perp)$ are *unlabeled*. A *complete labeling* $\mathcal{L}' : V(G) \to \Sigma_{\text{out}}$ 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 $S = (v_1, \ldots, v_p)$ be a list of the poles of $\mathcal{H}$ and let $S = (v'_1, \ldots, v'_p)$ be a designated set of poles in some partially labeled graph $\mathcal{H}'$. The partially labeled graph $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$ 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 $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$. 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 $V(G') - V(H')$. See Figure 6.

In the proof of our $n^{o(1)} \to O(\log n)$ speedup thereom 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 $\overset{\star}{\sim}$ on partially labeled graphs (with designated poles) so that the following is true: if $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$, then such a legal labeling $\mathcal{L}'_\diamond$ must exist, regardless of the

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

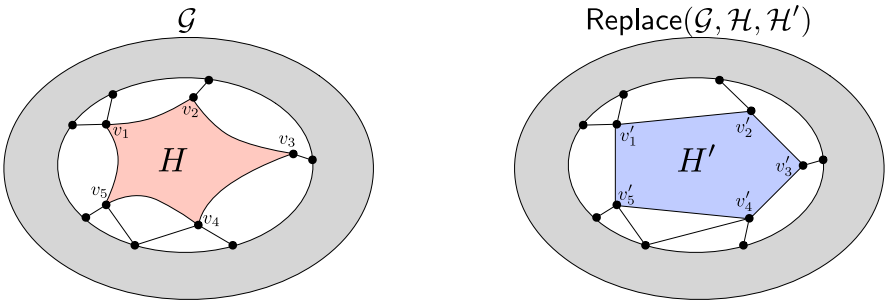

FIG. 6. *The operation* Replace.

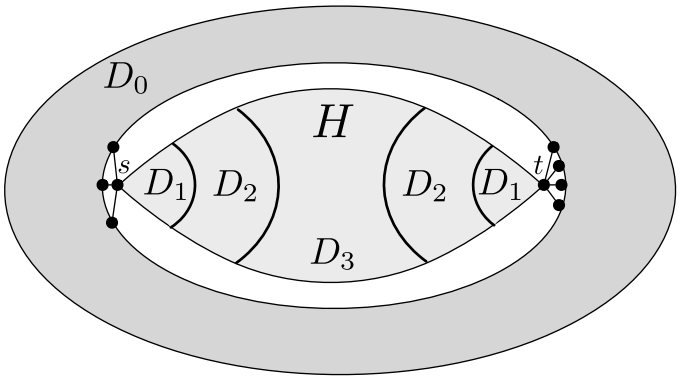

FIG. 7. *A partially labeled subgraph $\mathcal{H}$ with poles $S = (s, t)$, embedded in a larger graph $\mathcal{G}$. In the partition $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$, D_1 is the set of vertices in $V(H)$ within radius $r-1$ of S, D_2 are those within radius $2r-1$ of S, excluding D_1, and D_3 is the rest of $V(H)$. When $\mathcal{H}$ is embedded in some larger graph $\mathcal{G}$, D_0 denotes the remaining vertices in $V(G) - V(H)$.*

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 $S = (v_1, \ldots, v_p)$. Define $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ to be a tripartition of $V(H)$, where

$$D_1 = \bigcup_{v \in S} N^{r-1}(v),$$
$$D_2 = \bigcup_{v \in D_1} N^{r}(v) - D_1,$$
$$\text{and } D_3 = V(H) - (D_1 \cup D_2).$$

See Figure 7 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 $\mathcal{L}_* : D_1 \cup D_2 \to \Sigma_{\text{out}}$ assign output labels to $D_1 \cup D_2$. We say that $\mathcal{L}_*$ is *extendible* (to all of $V(H)$) if there exists a complete labeling $\mathcal{L}_\diamond$ of H such that $\mathcal{L}_\diamond$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

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 $D_2 \cup D_3$.[7]

3.5. An equivalence relation on graphs. Consider two partially labeled graphs $\mathcal{H}$ and $\mathcal{H}'$ with poles $S = (v_1, \ldots, v_p)$ and $S' = (v'_1, \ldots, v'_p)$, 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 $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$ holds if and only if there is a 1-1 correspondence $\phi : (D_1 \cup D_2) \to (D'_1 \cup D'_2)$ 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 ith pole in S if and only if v' is the ith 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 $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$.

3.6. Properties of the equivalence relation. Let us consider the graph $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$ and the two partitions $\xi(\mathcal{H}, S) = (D_1, D_2, D_3)$ and $\xi(\mathcal{H}', S') = (D'_1, D'_2, D'_3)$. Let $D_0 = V(G) - V(H)$ and $D'_0 = V(G') - V(H')$ be the remaining vertices in G and G', respectively.

If $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$, then there exists a 1-1 correspondence $\phi : (D_0 \cup D_1 \cup D_2) \to (D'_0 \cup D'_1 \cup D'_2)$ 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 $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$. Such a 1-1 correspondence ϕ is called *good*. We have the following lemma.

LEMMA 3.1. *Let $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$. Consider the two 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 $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$, so there is a* good 1-1 *correspondence $\phi : (D_0 \cup D_1 \cup D_2) \to (D'_0 \cup D'_1 \cup D'_2)$. 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.

[7] We are not concerned with 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 $V(G) - V(H)$, the labeling of H does not provide enough information to tell if these vertices' r-neighborhoods will be consistent with $\mathcal{P}$. See Figure 7.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

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 to all of $\mathcal{H}$. Since $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$, the labeling $\mathcal{L}'_*$ must also be extendible to all of $\mathcal{H}'$. 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 $v' \in D'_2$.[8] ∎

Theorem 3.2 provides a user-friendly corollary of Lemma 3.1, which does not mention the tripartition ξ.

THEOREM 3.2. *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 $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$ and let $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$. 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 3.2 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 3.2 is also legal. Observe that setting $\mathcal{G} = \mathcal{H}$ in Theorem 3.2 implies $\mathcal{G}' = \mathcal{H}'$. Suppose that $\mathcal{H}$ admits a legal labeling. For *any* $(\mathcal{H}', S')$ such that $(\mathcal{H}', S') \overset{\star}{\sim} (\mathcal{H}, S)$, 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 3.3 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 $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$.

THEOREM 3.3. *Let $\mathcal{G} = (G, \mathcal{L})$, and let $\mathcal{H} = (H, \mathcal{L})$ be a subgraph of $\mathcal{G}$. Suppose $\mathcal{H}'$ is a graph that satisfies $(\mathcal{H}, S) \overset{\star}{\sim} (\mathcal{H}', S')$ for some pole lists S, S'. Let $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, (\mathcal{H}, S), (\mathcal{H}', S'))$ 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 $(\mathcal{G}, X) \overset{\star}{\sim} (\mathcal{G}', X')$.*

Proof. Consider the 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 $(\mathcal{G}, X) \overset{\star}{\sim} (\mathcal{G}', X')$, it suffices to prove that ϕ (restricted to the domain $D_1 \cup D_2$) is a witness to the relation $(\mathcal{G}, X) \overset{\star}{\sim} (\mathcal{G}', X')$.

Let $\mathcal{L}_* : (D_1 \cup D_2) \to \Sigma_{\text{out}}$ 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 all of $V(G)$ if and only if $\mathcal{L}'_*$ is extendible to all of $V(G')$. Since we can also write $\mathcal{G} = \mathsf{Replace}(\mathcal{G}', (\mathcal{H}', S'), (\mathcal{H}, S))$, it suffices to show just one direction, i.e., if $\mathcal{L}_*$ is extendible, then $\mathcal{L}'_*$ is extendible.

[8] It is this lemma that motivates our definition of the *tri*partition $\xi(\mathcal{H}, S)$. It is not clear how an analogue of Lemma 3.1 could be proved using the seemingly more natural *bi*partition, i.e., by collapsing D_1, D_2 into one set.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

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 3.1, there exists a complete labeling $\mathcal{L}'_\diamond$ of $\mathcal{G}'$ such that the two conditions in Lemma 3.1 are met. We show that this implies that $\mathcal{L}'_*$ is extendible.

Lemma 3.1 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 3.1 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 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 $D'_2 \cup D'_3$. ∎

3.7. The number of equivalence classes. An important feature of $\overset{\star}{\sim}$ 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.[9]

Consider a partially labeled graph $\mathcal{H}$ with poles $S = (v_1, \ldots, v_p)$. 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 list $S = (z)$, $z \in V(T)$ being the root. Define $\mathsf{Class}(\mathcal{T})$ to be the equivalence class of $(\mathcal{T}, S)$ w.r.t. $\overset{\star}{\sim}$. Notice that whether a partially labeled rooted tree $\mathcal{T}$ admits a legal labeling is determined by $\mathsf{Class}(\mathcal{T})$ (Theorem 3.2). We say that a class is *good* if each partially labeled rooted tree in the class admits a legal labeling; otherwise the class is *bad*. We write $\mathscr{C}$ to denote the set of all classes. Notice that $|\mathscr{C}|$ is constant. The following lemma is a specialization of Theorem 3.3.

LEMMA 3.4. *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 $\mathsf{Class}(\mathcal{T}') = \mathsf{Class}(\mathcal{T}'')$. Then replacing $\mathcal{T}'$ with $\mathcal{T}''$ does not alter the class of $\mathcal{T}$.*

[9] For 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 p) is still a constant.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

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 $P = (v_1, \ldots, v_k)$, 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 $\mathsf{Type}(\mathcal{H})$ as the equivalence class of $(\mathcal{H}, S = (s, t))$ w.r.t. $\overset{\star}{\sim}$. The following lemma follows from Theorem 3.3.

LEMMA 3.5. *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 $\mathsf{Class}(\mathcal{T})$ is determined by $\mathsf{Type}(\mathcal{H})$. If we write $\mathcal{H} = (\mathcal{T}_i)_{i\in[k]}$, then $\mathsf{Type}(\mathcal{H})$ is determined by $\mathsf{Class}(\mathcal{T}_1), \ldots, \mathsf{Class}(\mathcal{T}_k)$.*

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 $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$ 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 3.2 and 3.3.

LEMMA 3.6. *Consider $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$. If $\mathsf{Type}(\mathcal{H}') = \mathsf{Type}(\mathcal{H})$ 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 3.7. *Suppose that $\mathcal{G} = (\mathcal{T}_i)_{i\in[k]}$ is a partially labeled bipolar tree, $\mathcal{H} = (\mathcal{T}_i, \ldots, \mathcal{T}_j)$ is a bipolar subtree of $\mathcal{G}$, and $\mathcal{H}'$ is some other partially labeled bipolar tree with $\mathsf{Type}(\mathcal{H}') = \mathsf{Type}(\mathcal{H})$. Then $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$ is a partially labeled bipolar tree and $\mathsf{Type}(\mathcal{G}') = \mathsf{Type}(\mathcal{G})$.*

LEMMA 3.8. *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 $\mathsf{Type}(\mathcal{H}')$ is a function of $\mathsf{Type}(\mathcal{H})$ and $\mathsf{Class}(\mathcal{T}_{k+1})$.*

Lemma 3.8 is what allows us to bring classical automata theory into play. Suppose that we somehow computed and stored $c_i = \mathsf{Class}(\mathcal{T}_i)$ at the root of $\mathcal{T}_i$. Lemma 3.8 implies that a finite automaton walking along the core path of $\mathcal{H}' = (\mathcal{T}_i)_{i\in[k+1]}$ can compute $\mathsf{Type}(\mathcal{H}')$, by reading the vector $(c_1, \ldots, c_{k+1})$ 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 3.9. *Let $\mathcal{H} = (\mathcal{T}_1, \ldots, \mathcal{T}_k)$ with $k \geq \ell_{\text{pump}}$. We regard each $\mathcal{T}_i$ in the string notation $\mathcal{H} = (\mathcal{T}_1, \ldots, \mathcal{T}_k)$ 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) $\mathsf{Type}(x \circ y^j \circ z) = \mathsf{Type}(\mathcal{H})$ *for each nonnegative integer j.*

We will use Lemma 3.9 to expand the length of the core path of a bipolar tree to be close 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 $k \geq \ell_{\text{pump}}$. The function $\mathsf{Pump}(\mathcal{H}, w)$ produces a partially labeled bipolar tree $\mathcal{H}' = (\mathcal{T}'_i)_{i\in[k']}$ such that (i) $\mathsf{Type}(\mathcal{H}) = \mathsf{Type}(\mathcal{H}')$, (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 3.9, such a function Pump exists.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

3.9. Rake and 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 $V_1 \cup \cdots \cup V_L$, $L = O(\log n)$. Our algorithm is inspired by Miller and Reif's *parallel tree contraction* [39]. We first describe the decomposition algorithm then analyze its properties.

Fix the constant $\ell = 2(r + \ell_{\text{pump}})$, where r and 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 [38], defined as follows.

DEFINITION 3.10. *Let P be a path. A subset $I \subset V(P)$ 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 $\mathrm{ID}(v) > \mathrm{ID}(u)$, 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}. Notice that the 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.

Properties of the decomposition. As we show below, $L = O(\log n)$ iterations suffice, i.e., $V(G) = V_1 \cup \cdots \cup V_L$. 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 $\mathscr{P}_i$ as the set of connected components (paths) induced by vertices in V_i that contain more than one vertex. For each $P \in \mathscr{P}_i$, $\ell \leq |V(P)| \leq 2\ell$ and $\deg_{G_i}(v) = 2$ for each vertex $v \in V(P)$.
- The graph G_L contains only isolated vertices, i.e., $\mathscr{P}_L = \emptyset$.

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 $P \in \mathscr{P}_i$.

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 ith iteration. In general, the graph induced by U is a forest, but it is simpler to analyze a single connected component W. We claim that at least a constant $\Omega(1/\ell)$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

fraction of vertices in W are eliminated (i.e., tagged i_C or i_R) in the ith 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 $\deg_U(v) \le 2$ 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 without loss of generality (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 $W_3 = \{v \in W \mid \deg_U(v) \ge 3\}$.

Case 1: $|W_2| \ge \frac{\ell|W|}{\ell+1}$. 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 $|W_1| > \frac{|W_1|+|W_3|}{2} = \frac{|W|-|W_2|}{2} \ge \frac{|W|}{2(\ell+1)}$. 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 three operations Extend, Label, and Duplicate-Cut which are used extensively in sections 3.11 and 3.12. All these operations are graph-theoretic operations, and they are not implemented in a distributed manner.

The operation Extend is parameterized by a target length $w \ge \ell = 2(r + \ell_{\text{pump}})$. 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 $\mathcal{H}$.[10] The function f will be constructed in section 3.13.

Label. Let $\mathcal{H} = (\mathcal{T}_1, \ldots, \mathcal{T}_x)$ be a partially labeled bipolar tree with $x \ge \ell$. Let $(v_1, \ldots, v_x)$ 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 $\mathcal{H}' = \mathsf{Label}(\mathcal{H})$ is defined as the result of assigning output labels to vertices in $N^{r-1}(e)$ by the function f.[11]

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

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 3.7. 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 $\mathcal{H}' = \mathsf{Extend}(\mathsf{Label}(\mathcal{H}))$ is invariant over all choices of the parameter w.[12] We have the following lemma.

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

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

[12] Notice 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.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

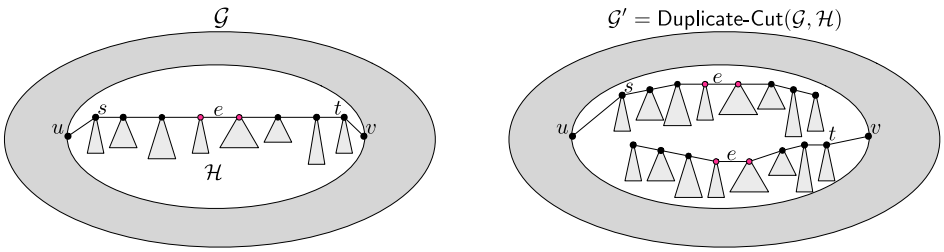

FIG. 8. *Left: A bipolar subtree $\mathcal{H}$ is attached to the rest of the graph $\mathcal{G}$ via edges $\{u, s\}, \{v, t\}$. The pink nodes have been precommitted to output labels by* Label *($r = 1$). Right: The* Duplicate-Cut *operation duplicates $\mathcal{H}$ and attaches one copy to u and the other to v.*

LEMMA 3.11. *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 $\mathsf{Type}(\tilde{\mathcal{H}}) = \mathsf{Type}(\mathcal{H})$ and $\mathcal{H}' = \mathsf{Extend}(\mathsf{Label}(\tilde{\mathcal{H}}))$. If $\mathcal{G}' = \mathsf{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}')$ 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

$$\mathsf{Type}(\mathsf{Extend}(\tilde{\mathcal{H}})) = \mathsf{Type}(\tilde{\mathcal{H}}) = \mathsf{Type}(\mathcal{H}).$$

Define $\mathcal{H}'' = \mathsf{Extend}(\tilde{\mathcal{H}})$ and $\mathcal{G}'' = \mathsf{Replace}(\mathcal{G}, \mathcal{H}, \mathcal{H}'')$. Observe that the graph $\mathcal{H}' = \mathsf{Extend}(\mathsf{Label}(\tilde{\mathcal{H}}))$ can be seen as the result of fixing the output labels of some unlabeled vertices in $\mathcal{H}'' = \mathsf{Extend}(\tilde{\mathcal{H}})$. Therefore, $\mathcal{L}'_\diamond$ is also a legal labeling of $\mathcal{G}''$. By Lemma 3.6, the desired legal labeling $\mathcal{L}_\diamond$ of $\mathcal{G} = \mathsf{Replace}(\mathcal{G}'', \mathcal{H}'', \mathcal{H})$ can be obtained from the legal labeling $\mathcal{L}'_\diamond$ of $\mathcal{G}''$. ∎

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 $\mathcal{G}' = \mathsf{Duplicate\text{-}Cut}(\mathcal{G}, \mathcal{H})$ 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 copy of $\mathcal{H}$, and removing the edge from v to the other copy of $\mathcal{H}$.

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 3.12 shows, the Duplicate-Cut operation (in conjunction with Extend and Label) allows both poles to work independently and cleanly integrate their labelings afterward.

LEMMA 3.12. *Let $\mathcal{H} = \mathsf{Extend}(\mathsf{Label}(\tilde{\mathcal{H}}))$ for some partially labeled bipolar tree $\tilde{\mathcal{H}}$. If $\mathcal{G}' = \mathsf{Duplicate\text{-}Cut}(\mathcal{G}, \mathcal{H})$ 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 $\mathcal{H} = (\mathcal{T}_1, \ldots, \mathcal{T}_x)$. Let $(v_1, \ldots, v_x)$ 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

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

the two edges that connect H two 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.[13] 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 $z \in V(G) - V(H)$, $\phi(z)$ is the corresponding vertex in G'.
- For $z \in \bigcup_{i=1}^{j} \mathcal{T}_i$, $\phi(z)$ is the corresponding vertex in H_u.
- For $z \in \bigcup_{i=j+1}^{x} \mathcal{T}_i$, $\phi(z)$ is the corresponding vertex in 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. ∎

Notice that in the proof of Lemma 3.12, the only property of $\mathcal{H}$ that we use is that $N^{r-1}(e)$ was assigned output labels in the application of $\mathsf{Label}(\tilde{\mathcal{H}})$.

3.11. A hierarchy of partially labeled trees. In this section we construct several sets of partially labeled unipolar and bipolar trees—$\{\mathscr{T}_i\}$, $\{\mathscr{H}_i\}$, and $\{\mathscr{H}_i^+\}$, $i \in \mathbb{Z}^+$—using the operations Extend and Label. If each member of $\mathscr{T}^\star = \bigcup_i \mathscr{T}_i$ admits a legal labeling, then we can use these trees to design an $O(\log n)$-time $\mathsf{DetLOCAL}$ algorithm for $\mathcal{P}$. Each $\mathcal{T} \in \mathscr{T}^\star$ 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 $\{\mathscr{H}_i\}_{i\in\mathbb{Z}^+}$ and $\{\mathscr{H}_i^+\}_{i\in\mathbb{Z}^+}$ and unipolar trees $\{\mathscr{T}_i\}_{i\in\mathbb{Z}^+}$ are defined inductively. In the base case we have $\mathscr{T}_1 = \{\mathcal{T}\}$, where $\mathcal{T}$ is the unique unlabeled, single-vertex, unipolar tree.

$\mathscr{T}$ Sets: For each $i > 1$, $\mathscr{T}_i$ consists of all partially labeled rooted trees $\mathcal{T}$ formed in the following manner. The root z of $\mathcal{T}$ has degree $0 \le \deg(z) \le \Delta$. Each child of z is either (i) the root of a partially labeled rooted tree $\mathcal{T}'$ from $\mathscr{T}_{i-1}$ (having degree at most $\Delta - 1$ in $\mathcal{T}'$) or (ii) one of the two poles of a bipolar tree $\mathcal{H}$ from $\mathscr{H}_{i-1}^+$.

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

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

Observe that whereas $\{\mathscr{T}_i\}$ and $\{\mathscr{H}_i\}$ grow without end, and contain arbitrarily large trees, the cardinality of $\mathscr{H}_i^+$ is at most the total number of types, which is constant.[14] This is due to the observation that whenever we add a new partially labeled bipolar tree $\mathsf{Extend}(\mathsf{Label}(\mathcal{H}))$ to $\mathscr{H}_i^+$, it is guaranteed that there is no other partially labeled bipolar tree $\mathsf{Extend}(\mathsf{Label}(\tilde{\mathcal{H}})) \in \mathscr{H}_i^+$ such that $\mathsf{Type}(\tilde{\mathcal{H}}) = \mathsf{Type}(\mathcal{H})$. The property

[13] Since 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.

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

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

that $|\mathscr{H}_i^+|$ is constant is crucial in the proof of Lemma 3.20. Lemmas 3.13–3.16 reveal some useful properties of these sets.

LEMMA 3.13. *We have* (i) $\mathscr{T}_1 \subseteq \mathscr{T}_2 \subseteq \cdots$, (ii) $\mathscr{H}_1 \subseteq \mathscr{H}_2 \subseteq \cdots$, *and* (iii) $\mathscr{H}_1^+ \subseteq \mathscr{H}_2^+ \subseteq \cdots$.

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

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

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

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

LEMMA 3.14. *Define* $k^\star = |\mathscr{C}|$, *where* $\mathscr{C}$ *is the set of all classes. Then we have* $\mathsf{Class}(\mathscr{T}^\star) = \mathsf{Class}(\mathscr{T}_{k^\star})$.

Proof. For each $i > 1$, $\mathsf{Class}(\mathscr{T}_i)$ depends only on $\mathsf{Type}(\mathscr{H}_{i-1}^+)$ and $\mathsf{Class}(\mathscr{T}_{i-1})$, due to Lemmas 3.4 and 3.5. Let i^* be the smallest index such that $\mathsf{Class}(\mathscr{T}_{i^*}) = \mathsf{Class}(\mathscr{T}_{i^*+1})$. Then we have $\mathsf{Type}(\mathscr{H}_{i^*}) = \mathsf{Type}(\mathscr{H}_{i^*+1})$ and as a consequence, $\mathscr{H}_{i^*}^+ = \mathscr{H}_{i^*+1}^+$. This implies that $\mathsf{Class}(\mathscr{T}_{i^*+1}) = \mathsf{Class}(\mathscr{T}_{i^*+2})$. By repeating the same argument, we conclude that for each $j \geq i^*$, we have $\mathsf{Class}(\mathscr{T}_j) = \mathsf{Class}(\mathscr{T}_{i^*}) = \mathsf{Class}(\mathscr{T}^\star)$. Since $\mathscr{T}_1 \subseteq \mathscr{T}_2 \subseteq \cdots$ (Lemma 3.13), we have $i^* \leq |\mathscr{C}|$. ∎

LEMMA 3.15. *For each* i, $\mathsf{Class}(\mathscr{T}_i)$ *does not depend on the parameter* w *used in the operation* Extend.

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

LEMMA 3.16. *The maximum number of vertices of a tree in* $\mathscr{T}_i$, *over all choices of labeling function* f, *is at most* λ^{i-1}, *where* $\lambda = 2\Delta(r + w + \ell_{\mathrm{pump}})$.

Proof. For any i, we write t_i (resp., h_i) to denote the maximum number of vertices of a tree in $\mathscr{T}_i$ (resp., $\mathscr{H}_i^+$). By the definition of these sets, we have the following formulas, which together imply that $t_i \leq \lambda^{i-1}$, where $\lambda = 2\Delta(r + w + \ell_{\mathrm{pump}})$:

$$t_1 = 1, \tag{3.1}$$

$$t_i \leq \Delta \max\{t_{i-1}, h_{i-1}\} \qquad \text{for } i > 1, \tag{3.2}$$

$$h_i \leq (2(w + \ell_{\mathrm{pump}}) + 2r)t_i \qquad \text{for } i \geq 1. \tag{3.3}$$

We explain the numbers in the upper bound on h_i. The operation Extend takes $\mathcal{H} = \mathcal{X} \circ \mathcal{Y} \circ \mathcal{Z}$ as an input and returns $\mathcal{H}' = \mathsf{Pump}(\mathcal{X}, w) \circ \mathcal{Y} \circ \mathsf{Pump}(\mathcal{Z}, w)$; the length of the core path of $\mathcal{Y}$ is $2r$; the length of the core path of both $\mathsf{Pump}(\mathcal{X}, w)$ and $\mathsf{Pump}(\mathcal{Z}, w)$ is at most $w + \ell_{\mathrm{pump}}$.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Notice that Formula 3.3 is not tight in the sense that we actually have $\mathscr{H}_{i^*}^+ = \mathscr{H}_{i^*+1}^+ = \cdots$, i.e., the sequence (h_i) stops growing as $i \geq i^*$. However, even for $i \geq i^*$, the sequence (t_i) still grows exponentially in view of Formula 3.2. ∎

Feasible labeling function. In view of Lemma 3.15, $\mathsf{Class}(\mathscr{T}^\star)$ 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 $\mathsf{Class}(\mathscr{T}^\star)$ 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 $n^{o(1)}$-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. An $O(\log n)$-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 an $O(\log n)$-time DetLOCAL algorithm for $\mathcal{P}$ on bounded degree trees.

Regardless of f, the algorithm begins by computing the graph decomposition $V(G) = V_1 \cup \cdots \cup V_L$ with $L = O(\log n)$; see section 3.9. We let the three infinite sequences $\{\mathscr{H}_i\}_{i\in\mathbb{Z}^+}$, $\{\mathscr{H}_i^+\}_{i\in\mathbb{Z}^+}$, and $\{\mathscr{T}_i\}_{i\in\mathbb{Z}^+}$ be constructed with respect to a feasible function f and a sufficiently large parameter w. We will choose w to be large enough so that a feasible function exists. Notice that the operation Extend already requires $w \geq \ell = 2(r + \ell_{\text{pump}})$.

A sequence of partially labeled graphs. We define below a sequence of partially labeled graphs $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_L$, 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 $\mathsf{Duplicate\text{-}Cut}$. 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 its remaining vertices. Formally, the graph $\mathcal{R}_i$ is formed by taking G_i (the subforest induced by $\bigcup_{j=i}^{L} V_j$, defined in section 3.9), and identifying each vertex $u \in V(G_i)$ with the root of a partially labeled imaginary tree $\mathcal{T}_{u,i} \in \mathscr{T}_i$ (defined within the proof of Lemma 3.17). Since G_L consists solely of isolated vertices, $\mathcal{R}_L$ is the disjoint union of trees drawn from $\mathscr{T}_L$.

Once each vertex $v \in V(G_i) = \bigcup_{j=i}^{L} V_j$ in the communication network G knows $\mathcal{T}_{v,i}$, we are able to simulate the imaginary graph $\mathcal{R}_i$ in the communication network G. In particular, a *legal labeling of* $\mathcal{R}_i$ is represented by storing the entire output labeling of the (imaginary) tree $\mathcal{T}_{v,i}$ at the (real) vertex $v \in V(G_i)$.

The official, inductive construction of $\mathcal{R}_i$ is described in the proof of Lemma 3.17. We remark that the "precommitment" of output labeling specified by the function f during the operation Label (in the construction of $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_L$) is used only in the *imaginary* trees. This does not directly lead to any *real* vertices committing to specific output labels.

LEMMA 3.17. *Suppose that a feasible function f is given. The partially labeled graphs $\mathcal{R}_1, \ldots, \mathcal{R}_L$ and partially labeled trees $\{\mathcal{T}_{v,i} \mid v \in V(G_i), i \in [L]\}$ can be constructed in $O(\log n)$ time meeting the following conditions:*

1. *For each $i \in [1, L]$, each vertex $v \in V(G_i) = \bigcup_{j=i}^{L} V_j$ knows $\mathcal{T}_{v,i} \in \mathscr{T}_i$.*
2. *For each $i \in [2, L]$, given a legal labeling of $\mathcal{R}_i$, a legal labeling of $\mathcal{R}_{i-1}$ can be computed in $O(1)$ time.*

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

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

Base case. Define $\mathcal{R}_1 = G$. This satisfies the lemma since $\mathcal{T}_{v,1} \in \mathscr{T}_1$ must be the unlabeled single-vertex tree, for each $v \in V(G)$.

Inductive step. We can assume inductively that $\mathcal{R}_{i-1}$ and $\{\mathcal{T}_{v,i-1} \mid v \in V(G_{i-1})\}$ have been defined and satisfy the lemma. The set $\mathscr{P}_{i-1}$ was defined in section 3.9. Each $P \in \mathscr{P}_{i-1}$ 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 $P = (v_1, \ldots, v_x) \in \mathscr{P}_{i-1}$. The bipolar graphs $\mathcal{H}_P$ and $\mathcal{H}_P^+$ are defined as follows:

- Define $\mathcal{H}_P$ to be the partially labeled bipolar tree $(\mathcal{T}_{v_1,i-1}, \ldots, \mathcal{T}_{v_x,i-1})$. Notice that $\mathcal{H}_P$ is a subgraph of $\mathcal{R}_{i-1}$. Since $\mathcal{T}_{v_j,i-1} \in \mathscr{T}_{i-1}$, for each $j \in [x]$, it follows that $\mathcal{H}_P \in \mathscr{H}_{i-1}$.
- Construct $\mathcal{H}_P^+$ as follows. Select the unique member $\tilde{\mathcal{H}} \in \mathscr{H}_{i-1}$ such that (i) $\mathsf{Type}(\tilde{\mathcal{H}}) = \mathsf{Type}(\mathcal{H}_P)$ and (ii) $\mathsf{Extend}(\mathsf{Label}(\tilde{\mathcal{H}})) \in \mathscr{H}_{i-1}^+$, and then set $\mathcal{H}_P^+ = \mathsf{Extend}(\mathsf{Label}(\tilde{\mathcal{H}})) \in \mathscr{H}_{i-1}^+$. Due to the way we define $\mathscr{H}_{i-1}^+$, such a graph $\tilde{\mathcal{H}} \in \mathscr{H}_{i-1}$ must exist, as $\mathcal{H}_P \in \mathscr{H}_{i-1}$.

The partially labeled graph $\mathcal{R}_i$ is constructed from $\mathcal{R}_{i-1}$ with the following three-step procedure. See Figure 9 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 $\mathcal{T}_{v,i-1}$ is a connected component of $\mathcal{R}_{i-1}$, remove $\mathcal{T}_{v,i-1}$. Notice that a tree $\mathcal{T}_{v,i-1}$ is a connected component of $\mathcal{R}_{i-1}$ if and only if v's neighborhood in G contains only vertices at lower levels: $V_1, \ldots, V_{i-2}$.

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

After Steps 1–3, for $v \in V(G_i)$, $\mathcal{T}_{v,i}$ is now defined to be the tree in $\mathcal{R}_i - (V(G_i) - \{v\})$ rooted at v. Notice that the two copies of $\mathcal{H}_P^+$ generated during Step 3(ii) become subtrees of $\mathcal{T}_{u,i}$ and $\mathcal{T}_{v,i}$, where u and v are the two vertices in $V(G_i)$ adjacent to the two endpoints of P in the graph G. See Figure 9.

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 $\mathcal{T}_{v,i}$ for all $v \in V(G_i)$ are uniquely determined. According to the construction of $\mathcal{R}_i$, each connected component of $\mathcal{R}_i - V(G_i)$ must be an imaginary tree that is either (i) some $\mathcal{T}_{v,j}$, where $v \in V_j$ and $j \in \{1, \ldots, i-1\}$, or (ii) a copy of $\mathcal{H}_P^+$, where $P \in \mathscr{P}_j$ and $j \in \{1, \ldots, i-1\}$. By induction (and Lemma 3.13), for $v \in V_1 \cup \cdots \cup V_j$ and $j \in \{1, \ldots, i-1\}$, we have $\mathcal{T}_{v,j} \in \mathscr{T}_j \subseteq \mathscr{T}_{i-1}$; for each $P \in \mathscr{P}_j$ where $j \in \{1, \ldots, i-1\}$, we have $\mathcal{H}_P^+ \in \mathscr{H}_j^+ \subseteq \mathscr{H}_{i-1}^+$. According to the inductive definition of $\mathscr{T}_i$, for each $v \in V(G_i)$ we have $\mathcal{T}_{v,i} \in \mathscr{T}_i$. This concludes the induction of part 1.

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 $\mathcal{T}_{v,i}$ is stored in $v \in V(G_i)$. We show how to compute a legal labeling of $\mathcal{R}_{i-1}$ in $O(1)$ time as follows. 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 $\mathcal{L}_4$. Recall that $\mathcal{R}_i$, $\mathcal{R}_{i-1}^+$, $\mathcal{R}'_{i-1}$, and $\mathcal{R}_{i-1}$ are all *imaginary*. "Time" refers to communications rounds in the *actual* network G, not any imaginary graph.

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 $\mathcal{R}_{i-1}^+$ adjacent to s, t, respectively. At this point u and v have legal labelings

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

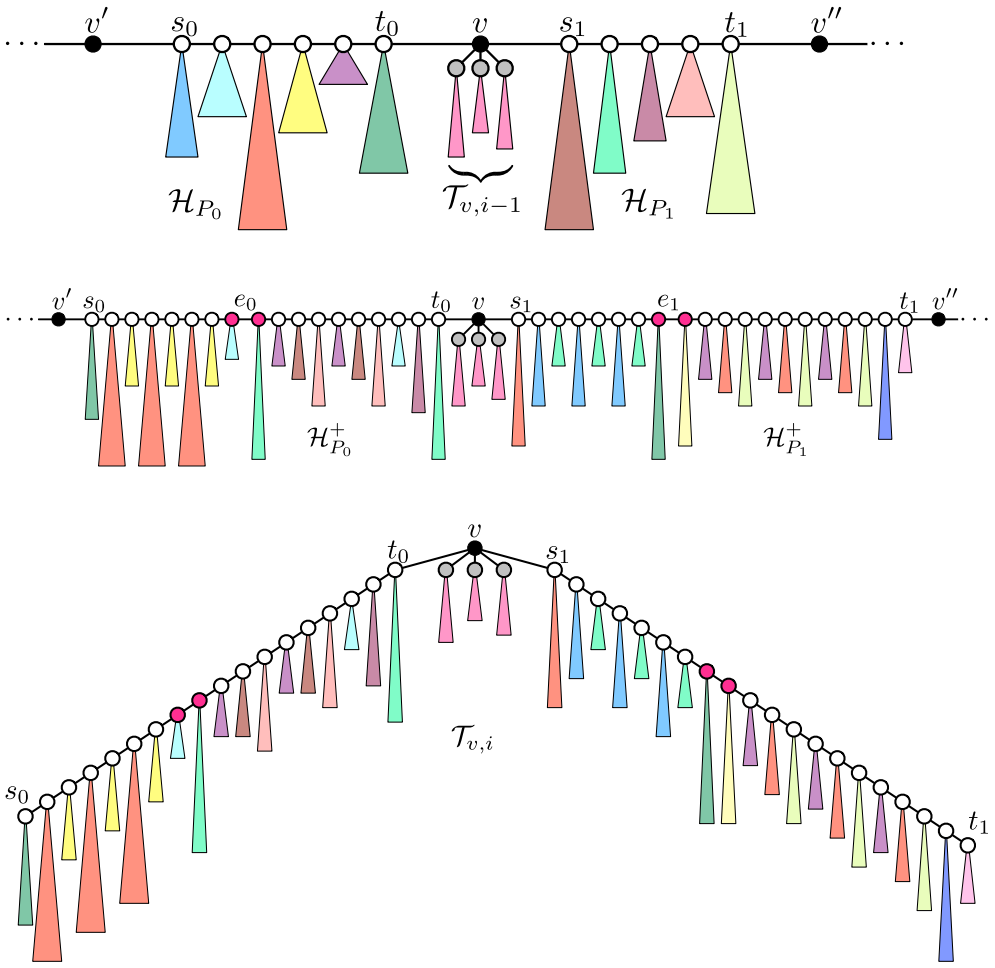

FIG. 9. *Top: In this example v was a vertex in a long degree-2 path tagged $(i-1)_C$ by the decomposition procedure and subsequently promoted to V_i. Black vertices are in V_i (or above); white vertices are in V_{i-1}; gray vertices are in V_{i-2} or below. The paths $P_0 = (s_0, \ldots, t_0)$ and $P_1 = (s_1, \ldots, t_1)$ 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 $\tilde{\mathcal{H}}_b \in \mathscr{H}^+_{i-1}$, $b \in \{0,1\}$ that has the same type as $\mathcal{H}_{P_b}$, and replacing $\mathcal{H}_{P_b}$ with $\mathcal{H}^+_{P_b} = \mathsf{Extend}(\mathsf{Label}(\tilde{\mathcal{H}}_b))$. The vertices receiving precommitted 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 $\mathcal{T}_{v,i}$ is the resulting tree rooted at v. Since each subtree of v is in $\mathscr{T}_{i-1}$ or $\mathscr{H}^+_{i-1}$, it follows that $\mathcal{T}_{v,i} \in \mathscr{T}_i$. In this case v had no neighbors at higher levels ($i+1$ and above), so $\mathcal{T}_{v,i}$ is a connected component of $\mathcal{R}_i$. Thus, v can locally compute a legal labeling of $\mathcal{T}_{v,i}$.*

of $\mathcal{T}_{u,i}$ and $\mathcal{T}_{v,i}$, both trees of which contain a copy of $\mathcal{H}^+_P$. Using Lemma 3.12 we integrate the labelings of $\mathcal{T}_{u,i}$ and $\mathcal{T}_{v,i}$ to fix a single legal labeling $\mathcal{L}_2$ of $\mathcal{H}^+_P$ in $\mathcal{R}^+_{i-1}$.[15]

[15] It 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 $\mathcal{H}^+_P$.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

From $\mathcal{L}_2$ to $\mathcal{L}_3$. A legal labeling $\mathcal{L}_3$ of $\mathcal{R}'_{i-1}$ is obtained by applying Lemma 3.11. For each $P \in \mathscr{P}_{i-1}$, 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 and t are the poles of $\mathcal{H}_P/\mathcal{H}_P^+$, and s and t know $\mathcal{L}_2$ on the $(2r-1)$-neighborhood of s and t in $\mathcal{H}_P^+$. By Lemma 3.11, 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 $\mathcal{T}_{v_j,i-1}$.

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 $\mathcal{T}_{v,i-1}$ that is a connected component of $\mathcal{R}_{i-1}$. A legal labeling $\mathcal{L}_4$ of $\mathcal{T}_{v,i-1}$ is computed locally at v, which is guaranteed to exist since $\mathcal{T}_{v,i-1} \in \mathscr{T}_{i-1}$.

This concludes the proof of the lemma. ∎

LEMMA 3.18. *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, for each $i \in [L]$, each vertex $v \in V_i$ computes the partially labeled rooted trees $\mathcal{T}_{v,j}$ for all $j \in [1, i]$; this can be done in $O(\log n)$ rounds. Since f is feasible, each partially labeled tree in $\mathscr{T}^\star$ 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 3.17(2). ∎

3.13. Existence of feasible labeling function. In Lemmas 3.19 and 3.20 we show *two* distinct ways to arrive at a feasible labeling function. In Lemma 3.19 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.[16] Lemma 3.19 suffices to prove our $n^{o(1)} \to O(\log n)$ 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 3.20 proves that this question is, in fact, decidable, which serves to highlight the delicate boundary between decidable and undecidable problems in LCL complexity [7, 42].

We briefly discuss some ideas behind the way we construct f. One natural attempt to assigning the output labels during Label is by simulating the given $n^{o(1)}$-time RandLOCAL algorithm $\mathcal{A}$. If we choose w to be sufficiently large (depending on n), then we can still force the runtime of the simulation to be less than w. This gives us a feasible function f that is *randomized*, which is enough for the purpose of establishing the $\omega(\log n)$—$n^{o(1)}$ gap in RandLOCAL.

In Lemma 3.19, we derandomize the above process with a choice of w *independent* of the size of the underlying graph n, thereby establishing the $\omega(\log n)$—$n^{o(1)}$ gap in DetLOCAL. In Lemma 3.20, we show that our construction of f leads to a decidability result.

[16] The 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.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

LEMMA 3.19. *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.*

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 $\mathscr{T}_{k^\star}$ over *all* choices of labeling function f. As $\Delta, r, \ell_{\text{pump}}$, and $k^\star$ are all constants, we have $N = w^{O(1)}$; see Lemma 3.16. 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 $\mathcal{T} = (T, \mathcal{L}) \in \mathscr{T}_{k^\star}$ 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 $\mathcal{H}' = \mathsf{Extend}(\mathcal{H})$, while pretending that the underlying graph has $\beta N + 1$ vertices. Notice that even though $\mathcal{A}$ is a randomized algorithm, there is no randomness involved in the definition of the labeling function f; that is, given the description of $\mathcal{A}$, the function f is defined *deterministically*. In the subsequent discussion, we will use the fact that the most probable labeling occurs with probability at least $|\Sigma_{\text{out}}|^{-\Delta^r} = 1/\beta$.

Proof idea. To show that f is good, all we need is to show that each member of $\mathscr{T}_{k^\star}$ admits a legal labeling. In what follows, consider *any* partially labeled rooted tree $\mathcal{T} = (T, \mathcal{L}) \in \mathscr{T}_{k^\star}$, where the set $\mathscr{T}_{k^\star}$ is constructed with the parameter w and function f. We prove that $\mathcal{T}$ admits a legal labeling $\mathcal{L}_\diamond$.

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.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Legal labeling $\mathcal{L}_\diamond$ 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 larger than under uniform randomness. 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 $\beta N/(\beta N+1) < 1$. Thus, $\mathcal{T} = (T, \mathcal{L})$ admits a legal labeling $\mathcal{L}_\diamond$. ∎

LEMMA 3.20. *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 $\{\mathscr{H}_i\}_{i\in\mathbb{Z}^+}$, $\{\mathscr{H}_i^+\}_{i\in\mathbb{Z}^+}$, and $\{\mathscr{T}_i\}_{i\in\mathbb{Z}^+}$, the number of distinct applications of the operation Label is constant, as $|\mathscr{H}_i^+|$ 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 $\mathscr{T}_{k^\star}$, as $k^\star = |\mathscr{C}|$ is a constant. By examining the classes of the partially labeled rooted trees in $\mathscr{T}_{k^\star}$ we can infer whether the function f is feasible (Lemma 3.14). Thus, deciding whether there exists a feasible function f can be done in bounded amount of time. ∎

Combining Lemmas 3.18, 3.19, and 3.20, we obtain the following theorem.

THEOREM 3.21. *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 $O(\log n)$.*

Discussion. To better understand Theorem 3.21, we consider some concrete examples. What would happen if we tried to apply the speedup theorem to the hierarchical $2\frac{1}{2}$-coloring $\mathcal{P}_2$ defined in section 2? Since the complexity of $\mathcal{P}_2$ is $\Theta(\sqrt{n})$, there does not exist a feasible function f for $\mathcal{P}_2$. In principle, one can write a program to test whether a feasible function f exists for a given LCL, but it is not hard to see that there is no feasible function for $\mathcal{P}_2$. Recall that $\mathscr{H}_1$ is the set of all bipolar, unlabeled paths with between ℓ and 2ℓ vertices. The partial labeling in $\mathscr{H}_1^+$ must not involve **a** and **b**, since the usage of these colors will make some members in $\mathscr{T}_2$ to have no legal labeling, due to the two-coloring rule. For example, consider a path $\mathcal{H} = \mathcal{H}_1 \circ \mathcal{H}_2 \circ \mathcal{H}_3$, where both $\mathcal{H}_1$ and $\mathcal{H}_3$ are colored by **a** and **b**, and $\mathcal{H}_2$ is unlabeled. Let $\mathcal{H}_2'$ be the path resulting from contracting one edge in $\mathcal{H}_2$, and let $\mathcal{H}' = \mathcal{H}_1 \circ \mathcal{H}_2' \circ \mathcal{H}_3$. If $\mathcal{H}$ admits a legal labeling, then $\mathcal{H}'$ must not have a legal labeling. Therefore, if there is a feasible function f for $\mathcal{P}_2$, then it must color all level 1 vertices **D**, since no V_1 vertex can be labeled **X** by the exemption rule. This coloring strategy clearly does not work (i.e., this does not give us an $O(\log n)$ time algorithm), since this requires each level 2 path (whose length can be $\Theta(n)$) to solve a 2-coloring problem.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Let us consider another problem. The problem of 3-coloring a 3-regular tree can be solved in $O(\log n)$ time, and so it admits a feasible function f. It is not hard to see that *any* function f that does a proper 3-coloring is feasible, i.e., the partial proper 3-coloring of any trees in $\mathscr{T}^\star$ can be completed to a full proper 3-coloring. For example, consider the above paths $\mathcal{H}$ and $\mathcal{H}'$, but here $\mathcal{H}_1$ and $\mathcal{H}_3$ are properly 3-colored. As long as $\mathcal{H}_2$ contain at least two vertices, both $\mathcal{H}$ and $\mathcal{H}'$ admits a proper 3-coloring.

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 $\mathrm{vbl}(A) \subset \mathcal{V}$ of variables. Each variable $V \in \mathcal{V}$ may have a different distribution and range, so long as the range is some finite set. The *dependency graph* $G_{\mathcal{X}} = (\mathcal{X}, \{(A,B) \mid \mathrm{vbl}(A) \cap \mathrm{vbl}(B) \neq \emptyset\})$ joins events by an edge if they depend on at least one common variable. The LLL and its variants give criteria under which $\Pr(\bigcap_{A\in\mathcal{X}} \overline{A}) > 0$, i.e., it is possible that none of the bad events occurs. We will narrow our discussion to *symmetric* criteria, expressed in terms of p and d, where $p = \max_{A\in\mathcal{X}} \Pr(A)$ and $d \geq 2$ is the maximum degree in $G_{\mathcal{X}}$. A standard version of the LLL states that if $ep(d+1) < 1$, then $\Pr(\bigcap \overline{A}) > 0$. 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 investigated in the sequential context [40, 29, 28, 33, 34, 31, 1] and from the point of the view of parallel and distributed computation [11, 18, 6, 9, 26, 15, 8, 19].

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 $\mathrm{vbl}(A) \subset \mathcal{V}$. 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 [40] parallel resampling algorithm implies an $O(\log^2 n)$ time RandLOCAL algorithm under the LLL criterion $ep(d+1) < 1$. Chung, Pettie, and Su [11] 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 $p \cdot \mathrm{poly}(d)2^d < 1$. They observed that under *any* criterion of the form $p \cdot f(d) < 1$, $\Omega(\log^* n)$ time is necessary. Ghaffari's [18] weak MIS algorithm, together with [11], implies an $O(\log d \cdot \log_{1/ep(d+1)} n)$ algorithm under LLL criterion $ep(d+1) < 1$. Brandt et al. [6] proved that $\Omega(\log_{\log(1/p)} \log n)$ time in RandLOCAL is necessary, even under the permissive LLL criterion $p2^d \leq 1$. Chang, Kopelowitz, and Pettie's [9] results imply that $\Omega(\log_d n)$ time is necessary in DetLOCAL, again, under the LLL criterion $p2^d \leq 1$.

We *define* $T_{LLL}(n,d,c)$ to be the RandLOCAL time to compute a point in the probability space avoiding all bad events (w.h.p.), under a "polynomial" LLL criterion of the form

$$pd^c < 1. \tag{4.1}$$

It is conceivable that the distributed complexity of the LLL is sensitive to the criterion used and depends on c. However, for our purpose (Theorem 4.1), *any* constant c is enough. In the subsequent discussion, we slightly abuse the notation to denote $T_{LLL}(n,d)$ as the distributed complexity of the LLL, where c is allowed to be an arbitrary constant. Earlier results [11, 6] imply that $T_{LLL}(n,d)$ is $\Omega(\log_{\log(1/p)} \log n)$, $\Omega(\log^* n)$, and $O(\log_{1/epd^2} n)$.

In this section we prove an automatic speedup theorem for RandLOCAL sublogarithmic algorithms. We do *not* assume that $\Delta = O(1)$ in this section. Theorem 4.1

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

considers algorithms that run in "sublogarithmic" time in RandLOCAL. The term *sublogarithmic* is insufficiently detailed, for two reasons. First, asymptotic notation is not always well defined when there are multiple free parameters (e.g., n and Δ). Second, and more importantly, the proof of Theorem 4.1 considers what happens when n gets very small, rather than $n \to \infty$. It is for these reasons that Theorem 4.1 assumes the running time can be written in a specific form.

THEOREM 4.1. *Suppose that $\mathcal{A}$ is a* RandLOCAL *algorithm that solves some LCL problem $\mathcal{P}$ (w.h.p.) in $T_\Delta(n)$ time. For any sufficiently small constant $\epsilon > 0$ and some function C, suppose $T_\Delta(n)$ is upper bounded by $C(\Delta) + \epsilon \log_\Delta n$. 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.*

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 $\mathcal{P}$. Once we settle on the LLL criterion exponent c in (4.1), we fix $\epsilon = O((2c)^{-1})$. Define $n^\star$ as the minimum value for which

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

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

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" $= n^\star$. Any $v \in V(G)$ will see a $t^\star$-neighborhood $N^{t^\star}(v)$ that is consistent with some $n^\star$-vertex graph. However, the probability of the *bad event* that $N^r(v)$ is incorrectly labeled is $1/n^\star$, not $1/\mathrm{poly}(n)$, as desired. We now show that this system of bad events satisfies the LLL criterion (4.1). Define the following events, graph, and quantities:

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

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 $N^{t^\star}(v')$, hence $\mathcal{E}_v$ is determined by (the data stored in, and random bits generated by) vertices in $N^{r+t^\star}(v)$. Clearly $\mathcal{E}_v$ is independent of any $\mathcal{E}_u$ for which $N^{r+t^\star}(u) \cap N^{r+t^\star}(v) = \emptyset$, 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 $\Delta^{2(r+t^\star)}$. By definition of $\mathcal{A}$, $\Pr(\mathcal{E}_v) \leq 1/n^\star = p$. This system satisfies LLL criterion (4.1) since, by definition of $t^\star$,

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

The algorithm $\mathcal{A}'$ now simulates a constructive LLL algorithm on $G_\mathcal{X}$ in order to find a labeling such that no bad event occurs. Since a virtual edge $(\mathcal{E}_u, \mathcal{E}_v)$ exists if

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

and only if u and v are at distance at most $2(r + t^\star) = O(C(\Delta))$, 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. ∎

Theorem 4.1 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 4.2 reflects the use of a particular distributed LLL algorithm, namely, [11, Corollary 1 and Algorithm 2]. It may be improvable using future distributed LLL technology.

The LLL algorithm of [11] works under the assumption that $epd^2 < 1$ and that each bad event $A \in \mathcal{X}$ is associated with a unique ID. The algorithm starts with a random assignment to the variables $\mathcal{V}$. In each iteration, let $\mathcal{F}$ be the set of bad events that occur under the current variable assignment; let $\mathcal{I}$ be the subset of $\mathcal{F}$ such that $A \in \mathcal{I}$ if and only if $\mathrm{ID}(A) < \mathrm{ID}(B)$ for each $B \in \mathcal{F}$ such that $\mathrm{vbl}(A) \cap \mathrm{vbl}(B) \neq \emptyset$. The next variable assignment is obtained by *resampling* all variables in $\bigcup_{A \in \mathcal{I}} \mathrm{vbl}(A)$. After $O(\log_{1/epd^2} n)$ iterations, no bad event occurs with probability $1 - 1/\mathrm{poly}(n)$.

THEOREM 4.2. *Let $\mathcal{A}$ be a $(C(\Delta) + \epsilon \log_\Delta n)$-time* RandLOCAL *algorithm that solves some LCL problem $\mathcal{P}$ w.h.p., where $\epsilon > 0$ is a sufficiently small constant. 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 $O(\min\{\log(|\Sigma_{\mathrm{out}}|) \cdot \Delta^{O(1)} + m_\Delta(\zeta) + \zeta,\ r_\Delta(\zeta) \cdot \zeta\})$-bit messages, where $\zeta = \Delta^{O(C(\Delta))}$ depends on Δ.*

Proof. We continue to use the notation and definitions from Theorem 4.1 and fix $c = 3$ in the LLL criterion (4.1). Since $d = \Omega(\Delta^{O(C(\Delta))}) = \Omega(\zeta)$ and we selected $t^\star$ 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 [11], each vertex $v \in V(G)$ will first generate an $O(\log n)$-bit unique identifier $\mathrm{ID}(\mathcal{E}_v)$ (which costs $O(\log n)$ random bits) and generate $r_\Delta(n^\star) \cdot O(\log_{1/epd^2} n) = O(r_\Delta(\zeta) \cdot \log_\zeta n)$ random bits throughout the computation. Thus, the total number of random bits per vertex is $O(\log n + r_\Delta(\zeta) \cdot \log_\zeta n)$.

In each resampling step of $\mathcal{A}'$, in order for v to tell whether $\mathcal{E}_v \in \mathcal{I}$, it needs the following information: (i) $\mathrm{ID}(\mathcal{E}_u)$ for all $u \in N^{2(r+t^\star)}(v)$, and (ii) whether $\mathcal{E}_u$ occurs under the current variable assignment, for all $u \in N^{2(r+t^\star)}(v)$. We now present two methods to execute one resampling step of $\mathcal{A}'$; they both take $O(C(\Delta))$ time using a message size that depends on Δ but is independent of n. There are $O(\log_{1/epd^2} n) = O(\log_\zeta n) = O(\frac{\log_\Delta n}{C(\Delta)})$ resampling steps, so the total time is $O(\log_\Delta n)$, independent of the function C.

Method 1. Before the LLL algorithm proper begins, we do the following preprocessing step. Each vertex v gathers up all IDs and random bits in its $3(t^\star + r)$-neighborhood. This takes $O((\log n + r_\Delta(\zeta) \cdot \log_\zeta n) \cdot \zeta/b)$ time with b-bit messages (recall that $\Delta^{O(t^\star + r)} = \Delta^{O(C(\Delta))} = \zeta$). In particular, the runtime can be made $O(\log_\Delta n)$ if we set $b = O(r_\Delta(\zeta) \cdot \zeta)$.

During the LLL algorithm, each vertex u owns one random variable: an $r_\Delta(n^\star)$-bit string V_u. In order for v to tell whether $\mathcal{E}_u$ occurs for each $u \in N^{2(r+t^\star)}(v)$ under the current variable assignment, it only needs to know how many times each V_u, $u \in$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

$N^{3(r+t^\star)}(v)$, has been resampled. Whether the output labeling of $u \in N^{2(r+t^\star)}(v)$ is locally consistent depends on the output labeling of vertices in $N^r(u)$, which depends on the random bits and the graph topology within $N^{r+t^\star}(u) \subseteq N^{3(r+t^\star)}(v)$. Given the graph topology, IDs, and the random bits within $N^{3(r+t^\star)}(v)$, the vertex v can locally simulate $\mathcal{A}$ and decides whether $\mathcal{E}_v \in \mathcal{I}$.

Thus, in each iteration of the LLL algorithm, each vertex v simply needs to alert its $3(r+t^\star)$-neighborhood whether V_v is resampled or not. This can be accomplished in $O(r+t^\star) = O(C(\Delta))$ time with ζ-bit messages.

Method 2. In the second method, vertices keep their random bits private. Similar to the first method, we do a preprocessing step to let each vertex gather up all IDs in its $2(t^\star+r)$-neighborhood. This can be done in $O(\log_\Delta n)$ time using ζ-bit messages.

During the LLL algorithm, 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 $t^\star$ rounds, sending $m_\Delta(n^\star)$-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 in $r = O(1)$ rounds, sending $\log(|\Sigma_{\text{out}}|)\cdot\Delta^{O(1)}$-bit messages.[17] Next, for a vertex v to tell whether $\mathcal{E}_v \in \mathcal{I}$, it needs to know whether $\mathcal{E}_u$ occurs for all $u \in N^{2(r+t^\star)}(v)$. This information can be gathered in $O(C(\Delta))$ time using messages of size $O(\zeta)$. To summarize, the required message size is $O(\log(|\Sigma_{\text{out}}|)\cdot\Delta^{O(1)} + m_\Delta(\zeta) + \zeta)$. ∎

An interesting corollary of Theorem 4.2 is that when $\Delta = O(1)$, randomized algorithms with unbounded length messages can be simulated with 1-bit messages.

COROLLARY 4.3. *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 paths/cycles, grids/tori, and, to a lesser extent, bounded degree trees and bounded degree general graphs. After the preliminary publication of this paper [10], an impressive body of work [15, 19, 3, 8, 2] has improved our understanding of the complexity hierarchy on bounded degree graphs and the complexity of the distributed LLL. We restate a more detailed version of Conjecture 1 from [10].

CONJECTURE 5.1. *There exists a sufficiently large constant c such that the complexity of the distributed LLL problem under criterion $pd^c < 1$ is $O(\log\log n)$ in* RandLOCAL *and $O(\log n)$ in* DetLOCAL.

According to [9, Theorem 3], proving the DetLOCAL complexity of the LLL is $O(\log n)$ is a necessary (but not sufficient) precondition for showing its RandLOCAL complexity is $O(\log\log n)$. To prove Conjecture 5.1 we also need to show that LLL instances can be *shattered* in $O(\log\log n)$ time. Conjecture 5.1 has been confirmed for tree-structured dependency graphs (of any degree d); see [8, 15].

The results of Balliu et al. [3] imply that the complexity hierarchies for bounded degree trees and general graphs are definitely different. Whereas trees have no natural complexities between $\omega(\log n)$ and $n^{o(1)}$ (Theorem 3.21), there are an infinite number of such complexities on general graphs [3]. It is an open question whether the other parts of the complexity spectrum addressed in [3] are the same for trees and general graphs. In particular, are there any LCL problems whose complexity on bounded

[17] An 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. For example, consider the $O(\Delta)$ vertex coloring problem.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

degree trees is in the range $\Omega(\log(\log^* n))$—$o(\log^* n)$? Can complexities of the form $\Theta(n^r)$ be achieved for LCL problems on bounded degree trees, where r is *not* of the form $1/k$? Balliu et al. [2] demonstrated an $\omega(\sqrt{n})$—$o(n)$ gap for bounded degree trees.

Given a description of any LCL problem $\mathcal{P}$, Theorem 3.21 shows that it is decidable whether the RandLOCAL complexity of $\mathcal{P}$ is $n^{\Omega(1)}$ or the DetLOCAL complexity of $\mathcal{P}$ is $O(\log n)$. For other gaps on bounded degree trees (e.g., $\omega(\log^* n)$—$o(\log n)$), the decidability problem is still open.

Appendix A. Speedup implications of Naor and 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 $\Delta^{T(n')+r} < n'$. 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.

Naor and Stockmeyer [42] 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.

A.1. Requirements 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 complete p-uniform 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 [42, 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 paths/cycles, $p = 2\tau+1$. For grids/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. For example, for paths/cycles, $p = 2\tau + 2r + 1$, 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 n-cycle 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 less than $2^{\binom{\Delta^\tau}{2}} \cdot (\Delta!|\Sigma_{\text{in}}|)^p$.
- The number c is defined as $|\Sigma_{\text{out}}|^{p!z}$. Intuitively, we can use a number in $[c]$ to encode a function that maps a radius-τ centered subgraph, whose vertices are equipped with distinct vertex IDs drawn from some set S with cardinality p, to an output label in Σ_{out}.

Recall that vertices in DetLOCAL have $O(\log n)$-bit IDs, i.e., they can be viewed as elements of $[n^k]$ for some $k = O(1)$. 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.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

A.2. Automatic speedup theorems. According to the proof of [25, section 1, Theorem 2], we have

$$\begin{aligned} &\text{for } p = 1, \quad && R(p,m,c) = c(m-1)+1, \\ &\text{for } p > 1, \quad && R(p,m,c) \le 2c^x, \\ & && \text{where } x = \sum_{i=p-1}^{R(p-1,m,c)-1} \binom{i+1}{p-1} < R(p-1,m,c)^p. \end{aligned}$$

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

Observe that in all scenarios described in section A.1, if the running time τ satisfies $\tau = \tau(n) = \omega(1)$, we have $\log^* m + \log^* c = o(p)$. Therefore, having $p \le \epsilon \log^* n$ for some small enough constant ϵ suffices to meet the condition $n^k \ge R(p,m,c)$. 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:

$\omega(1)$—$o(\log^* n)$	for n-paths/cycles,
$\omega(1)$—$o\Big(\sqrt{\log^* n}\Big)$	for $(\sqrt{n} \times \sqrt{n})$-grids/tori,
$\omega(1)$—$o(\log(\log^* n))$	for bounded degree trees or bounded degree general graphs.

By [9, Corollary 1], the DetLOCAL and RandLOCAL complexities of any LCL problem are asymptotically the same if they are at most $2^{O(\log^* n)}$. Therefore, the above gaps apply not only to DetLOCAL but also to RandLOCAL.

Due to the "stepping-up lemma" (see [25, section 4, Lemma 17]), we have a lower bound $\log^*(R(p,m,2)) = \Omega(p)$ (for any p, m). Therefore, Naor and Stockmeyer's approach *alone* cannot give an $\omega(1)$—$o(\log^* n)$ gap for bounded degree trees. However, for a certain class of LCL problems on $(\sqrt{n} \times \sqrt{n})$-grids/tori, the gap can be widened to $\omega(1)$—$o(\log^* n)$ [7, p. 2]. The following proof is due to Suomela [46].

THEOREM A.1 (Suomela). *Let $\mathcal{P}$ be any LCL problem on $(\sqrt{n} \times \sqrt{n})$-grids/tori that does not refer to input labels or port numbering. The* DetLOCAL *and* RandLOCAL *complexity of $\mathcal{P}$ is either $O(1)$ or $\Omega(\log^* n)$.*

Proof. Given a $(\sqrt{n} \times \sqrt{n})$-torus G, we associate each vertex $v \in V(G)$ with a coordinate (α, β), where $\alpha, \beta \in \{0, \dots, \sqrt{n}-1\}$. We consider the following special way to generate unique $2k \log n$-bit IDs. Let ϕ_x and ϕ_y be two functions mapping integers in $\{0, \dots, \sqrt{n}-1\}$ to integers in $\{0, \dots, n^k - 1\}$. We additionally require that $\phi_x(0) < \cdots < \phi_x(\sqrt{n}-1) < \phi_y(0) < \cdots < \phi_y(\sqrt{n}-1)$. If v is at position (α, β), it has ID $\phi_x(\alpha) \cdot n^k + \phi_y(\beta)$. Notice that the IDs of all vertices in $N^\tau(v)$ can be deduced from just $4\tau + 2$ numbers: $\phi_x(i)$, $i \in [\alpha - \tau, \alpha + \tau]$, and $\phi_y(j)$, $j \in [\beta - \tau, \beta + \tau]$.

Suppose that the complexity of $\mathcal{P}$ is $o(\log^* n)$. Let $\mathcal{A}$ be any τ-round DetLOCAL algorithm for solving $\mathcal{P}$, where $\tau = o(\log^* n)$. Notice that the algorithm $\mathcal{A}$ works correctly even when we restrict ourselves to the above special ID assignment. Our goal is to show that $\mathcal{P}$ is actually *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, assuming w.l.o.g. that $\sqrt{n} > 2r + 1$. Thus, $\mathcal{P}$ can be solved in $O(1)$ rounds.

In subsequent discussion, we let v be any vertex whose position is (α, β), where $\tau + r \le \alpha \le (\sqrt{n}-1) - (\tau + r)$ and $\tau + r \le \beta \le (\sqrt{n}-1) - (\tau + r)$. That is, v is sufficiently far from the places where the coordinates wrap around.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Given $\mathcal{A}$, we construct a function f as follows. Let $S = (s_1, \ldots, s_{4\tau+2})$ be a vector of $4\tau + 2$ numbers in $\{0, \ldots, n^k - 1\}$ such that $s_l < s_{l+1}$ for each $l \in [4\tau + 2]$. Then $f(S) \in \Sigma_{\text{out}}$ is defined as the output labeling of v resulting from executing $\mathcal{A}$ with the following ID assignment of vertices in $N^\tau(v)$. We set $\phi_x(\alpha - \tau - 1 + i) = s_i$ for each $i \in [2\tau + 1]$ and set $\phi_y(\beta - \tau - 1 + j) = s_{j+2\tau+1}$ for each $j \in [2\tau + 1]$. Recall that $\mathcal{P}$ does not use port numbering and input labeling, so the output labeling of v depends only on IDs of vertices in $N^\tau(v)$.

We set $p = 4\tau + 2$, $m = 4\tau + 4r + 2$, and $c = |\Sigma_{\text{out}}|$. Notice that the calculation of the parameter c here is different from the original proof of Naor and Stockmeyer. Since we already force that $\phi_x(0) < \cdots < \phi_x(\sqrt{n} - 1) < \phi_y(0) < \cdots < \phi_y(\sqrt{n} - 1)$, we do not need to consider all $p!$ permutations of the set S.

We have $R(p, m, c) \ll n^k$ (since $p = o(\log^* n)$). Thus, there exists a set S' of m distinct numbers in $\{0, \ldots, n^k\}$ such that the following is true. We label these m numbers $\phi_x(i)$, $i \in [\alpha - \tau - r, \alpha + \tau + r]$, and $\phi_y(j)$, $j \in [\beta - \tau - r, \beta + \tau + r]$ by the set S' such that $\phi_x(\alpha - \tau - r) < \cdots < \phi_x(\alpha + \tau + r) < \phi_y(\beta - \tau - r) < \cdots < \phi_y(\beta + \tau + r)$. Then the output labels of all vertices in $N^r(v)$ assigned by $\mathcal{A}$ are identical.

Therefore, there exists an element $\sigma \in \Sigma_{\text{out}}$ such that labeling all vertices by σ yields a legal labeling of G. Thus, $\mathcal{P}$ can be solved in $O(1)$ rounds.

On grids, the proof above shows that the LCL $\mathcal{P}$ admits a labeling where all interior vertices (those at distance greater than r from the boundary) can be labeled uniformly by some $\sigma \in \Sigma_{\text{out}}$ and every other vertex can be labeled according to an $O(1)$-round order-invariant algorithm.

Similarly, by [9, Corollary 1], the $\omega(1)$—$o(\log^* n)$ gap given in this proof applies to both DetLOCAL and RandLOCAL. ∎

A.3. Discussion. It still remains an outstanding open problem whether the gap for other cases can also be widened to $\omega(1)$—$o(\log^* n)$.

The proof of Theorem A.1 extends easily to d-dimensional tori but does not extend to bounded degree trees, since there is a nontrivial problem that can be solved in $O(1)$ rounds on a subset of bounded degree trees (see the proof of Theorem A.1 for the definition of a trivial problem). A *weak coloring* is a coloring in which every vertex is colored differently than at least one neighbor. Naor and Stockmeyer [42] showed that on any graph class in which all vertex degrees are odd, *weak* $2^{O(\Delta \log \Delta)}$*-coloring* can be solved in two rounds and *weak* 2*-coloring* can be solved in $O(\log^* \Delta)$ rounds in DetLOCAL. This problem is *nontrivial* in the sense that coloring all vertices by the same color is not a legal solution. Since the d-dimensional torus is Δ-regular, $\Delta = 2d$, we infer that the complexity of weak $O(1)$-coloring on Δ-regular graphs is $\Theta(\log^* n)$ for every fixed even number $\Delta \geq 2$.

Theorem A.1 also does not extend to LCL problems that use input labels or port numbering. If either input labels or port numbering is allowed, then one can construct a nontrivial LCL problem that can be solved in $O(1)$ rounds even on cycle graphs. An *orientation* of a vertex $v \in V(G)$ is defined as a port number in $[\deg(v)]$, indicating a vertex in $N(v)$ that v is pointed toward. An ℓ*-orientation* of a cycle G is an orientation of all vertices in G meeting the following conditions. If $|V(G)| \leq \ell$, then all vertices in G are oriented to the same direction, i.e., no two vertices point toward each other. If $|V(G)| > \ell$, then each vertex $v \in V(G)$ belongs to a path P such that (i) all vertices in P are oriented to the same direction (no two point to each other) and (ii) the number of vertices in P is at least ℓ. Notice that ℓ-orientation, $\ell = O(1)$, is an LCL that refers to port numbering. We show that in $O(1)$ rounds we can compute an ℓ-orientation of G for any constant ℓ.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

THEOREM A.2. *Let G be a cycle graph and ℓ be a constant. There is a* DetLOCAL *algorithm that computes an ℓ-orientation of G in $O(1)$ rounds.*

Proof. This is a known result. See [27, Fact 5.2] or [14, Lemma 14 (Rounding Lemma), Case B] for a sketch of the proof. For the sake of completeness, we present a full proof. We first show how to compute a 2-orientation of a cycle G in $O(1)$ rounds, and then we extend it to any constant ℓ.

Computing a 2-orientation. We assume $|V(G)| \geq 3$. A DetLOCAL $O(1)$-round algorithm to compute a 2-orientation is described as follows: First, each vertex $v \in V(G)$ computes an arbitrary orientation. With respect to this orientation of G, define sets V_1, V_2, V_3 as follows.

- $v \in V_1$ if and only if there exists $u \in N(v)$ such that u and v are oriented to the same direction.
- $v \in V_2$ if and only if there exists $u \in N(v) \setminus V_1$ such that u and v are oriented toward each other.
- $V_3 = V(G) \setminus (V_1 \cup V_2)$. Observe that for each $v \in V_3$, there exists $u \in N(v) \cap V_1$.

A 2-orientation is obtained by reorienting the vertices in V_2 and V_3. The vertices in V_2 are partitioned into unordered pairs such that $u, v \in V_2$ are paired up if and only if (i) $\{u, v\} \in E(G)$ and (ii) u and v are oriented toward each other. For each pair $\{u, v\}$, reverse the orientation of any one of $\{u, v\}$. For each vertex $v \in V_3$, let u be any neighbor of v such that $u \in V_1$, and reorient v to the orientation of u.

Computing an ℓ-orientation. We define an $O(1)$-round DetLOCAL algorithm $\mathcal{A}_\ell$ that computes an ℓ-orientation. It makes recursive calls to $\mathcal{A}_{\lceil \ell/2 \rceil}$. In what follows, we assume $\ell \geq 3$ and $|V(G)| \geq 3$.

First, execute $\mathcal{A}_{\lceil \ell/2 \rceil}$ to obtain an $\lceil \ell/2 \rceil$-orientation of G. With respect to this orientation of G, define the following terminologies. Let $\mathscr{P}$ be the set of all maximal-size connected subgraphs in G such that all constituent vertices are oriented to the same direction. Notice that if $\mathscr{P}$ contains a cycle, then $\mathscr{P} = \{G\}$. Otherwise $\mathscr{P}$ contains only paths. Define $\mathscr{P}_1$ as the subset of $\mathscr{P}$ such that $P \in \mathscr{P}_1$ if and only if the number of vertices in P is at least ℓ. Define $\mathscr{P}_2$ as the subset of $\mathscr{P} \setminus \mathscr{P}_1$ such that $P \in \mathscr{P}_2$ if and only if there exists another path $P' \in \mathscr{P} \setminus \mathscr{P}_1$ meeting the following condition. There exist an endpoint u of P and an endpoint v of P' such that $\{u, v\} \in E(G)$, and u and v are oriented toward each other. Define $\mathscr{P}_3 = \mathscr{P} \setminus (\mathscr{P}_1 \cup \mathscr{P}_2)$. Observe that each $P \in \mathscr{P}_3$ is adjacent to a path in $\mathscr{P}_1$.

The paths in $\mathscr{P}_2$ are partitioned into unordered pairs such that $P, P' \in \mathscr{P}_2$ are paired up if and only if there exist an endpoint u of P and an endpoint v of P' such that $\{u, v\} \in E(G)$, and u and v are oriented toward each other. For each pair $\{P, P'\}$, reverse the orientation of all the vertices in any one of $\{P, P'\}$. For each path $P \in \mathscr{P}_3$, let $P' \in \mathscr{P}_1$ be any path adjacent to P, and re-orient P to the orientation of P'.

The round complexity of $\mathcal{A}_\ell$ satisfies the recurrence $T(\ell) = T(\lceil \ell/2 \rceil) + O(\ell)$, which is $O(\ell)$. ∎

Notice that even though orienting all vertices in the cycle to the same direction gives a legal labeling, ℓ-orientation is still a nontrivial LCL problem. Consider a subpath (v_1, v_2, v_3, v_4) in the cycle. Suppose that the port number of (v_2, v_3) stored at v_2 is 1, but the port number of (v_3, v_4) stored at v_3 is 2. Then we need to label v_2 and v_3 differently (1 and 2, respectively) in order to orient them in the same direction "$\rightarrow$".

Last, we remark that for the case the given $(\sqrt{n} \times \sqrt{n})$-torus is *oriented* in the sense that the input port numberings all agree with a fixed N/S/E/W orientation [7]; then there is no nontrivial LCL problem solvable in $O(1)$ time.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

REFERENCES

[1] D. Achlioptas and F. Iliopoulos, *Random walks that find perfect objects and the Lovász local lemma*, in Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2014, pp. 494–503, https://doi.org/10.1109/FOCS.2014.59.

[2] A. Balliu, S. Brandt, D. Olivetti, and J. Suomela, *Almost global problems in the LOCAL model*, in Proceedings of the 32nd International Symposium on Distributed Computing, 2018.

[3] A. Balliu, J. Hirvonen, J. H. Korhonen, T. Lempiäinen, D. Olivetti, and J. Suomela, *New classes of distributed time complexity*, in Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC), New York, 2018, pp. 1307–1318, https://doi.org/10.1145/3188745.3188860.

[4] 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*, J. ACM, 64 (2017), 23.

[5] L. Barenboim, M. Elkin, S. Pettie, and J. Schneider, *The locality of distributed symmetry breaking*, J. ACM, 63 (2016), 20.

[6] 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 of the 48th ACM Symposium on the Theory of Computing (STOC), 2016, pp. 479–488.

[7] 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*, in Proceedings of the 36th Annual ACM Symposium on Principles of Distributed Computing (PODC), 2017, pp. 101–110.

[8] Y.-J. Chang, Q. He, W. Li, S. Pettie, and J. Uitto, *The complexity of distributed edge coloring with small palettes*, in Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms (SODA), 2018, pp. 2633–2652.

[9] Y.-J. Chang, T. Kopelowitz, and S. Pettie, *An exponential separation between randomized and deterministic complexity in the LOCAL model*, in Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2016, pp. 615–624, https://doi.org/10.1109/FOCS.2016.72.

[10] Y.-J. Chang and S. Pettie, *A time hierarchy theorem for the LOCAL model*, in Proceedings of the 58th IEEE Symposium on Foundations of Computer Science (FOCS), 2017, pp. 156–167.

[11] K.-M. Chung, S. Pettie, and H.-H. Su, *Distributed algorithms for the Lovász local lemma and graph coloring*, Distrib. Comput., 30 (2017), pp. 261–280.

[12] R. Cole and U. Vishkin, *Deterministic coin tossing with applications to optimal parallel list ranking*, Inform. Control, 70 (1986), pp. 32–53.

[13] L. Feuilloley and P. Fraigniaud, *Survey of distributed decision*, Bull. Eur. Assoc. Theor. Comput. Sci. EATCS, 119 (2016), pp. 41–65.

[14] M. Fischer, *Improved deterministic distributed matching via rounding*, in Proceedings of the 31st International Symposium on Distributed Computing (DISC), 2017, pp. 17:1–17:15.

[15] M. Fischer and M. Ghaffari, *Sublogarithmic distributed algorithms for Lovász local lemma, and the complexity hierarchy*, in Proceedings of the 31st International Symposium on Distributed Computing (DISC), 2017, 18.

[16] P. Fraigniaud, A. Korman, and D. Peleg, *Towards a complexity theory for local distributed computing*, J. ACM, 60 (2013), 35, https://doi.org/10.1145/2499228.

[17] M. Fürer, *Data structures for distributed counting*, J. Comput. System Sci., 28 (1984), pp. 231–243, https://doi.org/10.1016/0022-0000(84)90067-9.

[18] M. Ghaffari, *An improved distributed algorithm for maximal independent set*, in Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2016, pp. 270–277, https://doi.org/10.1137/1.9781611974331.ch20.

[19] M. Ghaffari, D. G. Harris, and F. Kuhn, *On Derandomizing Local Distributed Algorithms*, in Proceedings of the 59th IEEE Symposium on Foundations of Computer Science (FOCS), 2018.

[20] M. Ghaffari, F. Kuhn, and Y. Maus, *On the complexity of local distributed graph problems*, in Proceedings of the 49th ACM Symposium on Theory of Computing (STOC), 2017, pp. 784–797.

[21] M. Ghaffari and H.-H. Su, *Distributed degree splitting, edge coloring, and orientations*, in Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2017, pp. 2505–2523, https://doi.org/10.1137/1.9781611974782.166.

[22] M. Göös, J. Hirvonen, and J. Suomela, *Linear-in-*Δ *lower bounds in the LOCAL model*, Distrib. Comput., 30 (2015), pp. 325–338, https://doi.org/10.1007/s00446-015-0245-8.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

[23] M. Göös and J. Suomela, *Locally checkable proofs in distributed computing*, Theory Comput., 12 (2016), pp. 1–33, https://doi.org/10.4086/toc.2016.v012a019.
[24] M. Göös and J. Suomela, *No sublogarithmic-time approximation scheme for bipartite vertex cover*, Distrib. Comput., 27 (2014), pp. 435–443, https://doi.org/10.1007/s00446-013-0194-z.
[25] R. L. Graham, B. L. Rothschild, and J. H. Spencer, *Ramsey Theory*, 2nd ed., John Wiley and Sons, New York, 1990.
[26] B. Haeupler and D. G. Harris, *Parallel algorithms and concentration bounds for the Lovász local lemma via witness-DAGs*, in Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2017, pp. 1170–1187, https://doi.org/10.1137/1.9781611974782.76.
[27] M. Hańćkowiak, M. Karoński, and A. Panconesi, *On the distributed complexity of computing maximal matchings*, SIAM J. Discrete Math., 15 (2001), pp. 41–57.
[28] D. G. Harris, *Lopsidependency in the Moser-Tardos framework: Beyond the lopsided Lovász local lemma*, ACM Trans. Algorithms, 13 (2016), 17, https://doi.org/10.1145/3015762.
[29] D. G. Harris and A. Srinivasan, *A constructive algorithm for the Lovász local lemma on permutations*, in Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2014, pp. 907–925, https://doi.org/10.1137/1.9781611973402.68.
[30] J. Hartmanis and R. E. Stearns, *On the computational complexity of algorithms*, Trans. Amer. Math. Soc., 117 (1965), pp. 285–306.
[31] N. J. A. Harvey and J. Vondrák, *An algorithmic proof of the Lovász local lemma via resampling oracles*, in Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2015, pp. 1327–1346, https://doi.org/10.1109/FOCS.2015.85.
[32] D. Hefetz, F. Kuhn, Y. Maus, and A. Steger, *Polynomial lower bound for distributed graph coloring in a weak LOCAL model*, in Proceedings of the 30th International Symposium on Distributed Computing (DISC), 2016, pp. 99–113, https://doi.org/10.1007/978-3-662-53426-7_8.
[33] K. B. R. Kolipaka and M. Szegedy, *Moser and Tardos meet Lovász*, in Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC), 2011, pp. 235–244, https://doi.org/10.1145/1993636.1993669.
[34] V. Kolmogorov, *Commutativity in the algorithmic Lovász local lemma*, in Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2016, pp. 780–787, https://doi.org/10.1109/FOCS.2016.88.
[35] A. Korman, J.-S. Sereni, and L. Viennot, *Toward more localized local algorithms: removing assumptions concerning global knowledge.*, Distrib. Comput., 26 (2013), pp. 289–308.
[36] F. Kuhn, T. Moscibroda, and R. Wattenhofer, *Local computation: Lower and upper bounds*, J. ACM, 63 (2016), 17, https://doi.org/10.1145/2742012.
[37] F. Kuhn and R. Wattenhofer, *On the complexity of distributed graph coloring*, in Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing (PODC), 2006, pp. 7–15.
[38] N. Linial, *Locality in distributed graph algorithms*, SIAM J. Comput., 21 (1992), pp. 193–201.
[39] G. L. Miller and J. H. Reif, *Parallel tree contraction—Part* I: *Fundamentals*, Adv. Comput. Res., 5 (1989), pp. 47–72.
[40] R. A. Moser and G. Tardos, *A constructive proof of the general Lovász local lemma*, J. ACM, 57 (2010), 11, https://doi.org/10.1145/1667053.1667060.
[41] M. Naor, *A lower bound on probabilistic algorithms for distributive ring coloring*, SIAM J. Discrete Math., 4 (1991), pp. 409–412, https://doi.org/10.1137/0404036.
[42] M. Naor and L. J. Stockmeyer, *What can be computed locally?*, SIAM J. Comput., 24 (1995), pp. 1259–1277, https://doi.org/10.1137/S0097539793254571.
[43] D. Peleg, *Distributed Computing: A Locality-Sensitive Approach*, Discrete Math. Appl. 5, SIAM, Philadelphia, 2000.
[44] S. Pettie and H.-H. Su, *Distributed algorithms for coloring triangle-free graphs*, Inform. and Comput., 243 (2015), pp. 263–280.
[45] J. Suomela, *Survey of local algorithms*, ACM Comput. Surv., 45 (2013), 24, https://doi.org/10.1145/2431211.2431223.
[46] J. Suomela, *private communication*, 2017.

Downloaded 08/08/19 to 141.211.4.224. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.