



Brief Announcement: Distributed Derandomization Revisited

Sameep Dahal  

Aalto University, Finland

Francesco d'Amore  

Aalto University, Finland

Henrik Lievonon  

Aalto University, Finland

Timothé Picavet  

Aalto University, Finland

ENS de Lyon, France

Jukka Suomela  

Aalto University, Finland

Abstract

One of the cornerstones of the distributed complexity theory is the derandomization result by Chang, Kopelowitz, and Pettie [FOCS 2016]: any randomized LOCAL algorithm that solves a locally checkable labeling problem (LCL) can be derandomized with at most exponential overhead. The original proof assumes that the number of random bits is bounded by some function of the input size. We give a new, simple proof that does not make any such assumptions – it holds even if the randomized algorithm uses infinitely many bits. While at it, we also broaden the scope of the result so that it is directly applicable far beyond LCL problems.

2012 ACM Subject Classification Theory of computation → Distributed computing models; Theory of computation → Pseudorandomness and derandomization

Keywords and phrases Distributed algorithm, Derandomization, LOCAL model

Digital Object Identifier 10.4230/LIPIcs.DISC.2023.40

Funding This work was supported in part by the Research Council of Finland, Grant 333837.

Acknowledgements We thank the participants in our reading group at Aalto University for helpful discussions.

1 Introduction

Distributed derandomization. A long line of recent work has led to a near-complete understanding of the distributed computational complexity of *locally checkable labeling problems* (LCLs) [5]. These are graph problems that can be defined by giving a finite list of feasible local neighborhoods [3]; for example, c -coloring in graphs of maximum degree Δ (for some fixed c and Δ) is an LCL problem.

We are in particular interested in the round complexity of LCLs in two standard models of distributed computing: deterministic and randomized versions of the LOCAL model [2, 4]. One of the cornerstones of the distributed complexity theory is the derandomization result by Chang, Kopelowitz, and Pettie [1, Theorem 3.1]:

► **Theorem 1** (Chang, Kopelowitz, and Pettie). *Let $\mathcal{A}_{\text{rand}}$ be a randomized LOCAL algorithm that solves an LCL problem \mathcal{P} in $T_{\text{rand}}(n)$ communication rounds in n -node graphs with probability at least $1 - 1/n$. Then there is a deterministic LOCAL algorithm \mathcal{A}_{det} that solves \mathcal{P} in $T_{\text{det}}(n)$ rounds, where $T_{\text{det}}(n) = T_{\text{rand}}(2^{n^2})$.*



© Sameep Dahal, Francesco d'Amore, Henrik Lievonon, Timothé Picavet, and Jukka Suomela; licensed under Creative Commons License CC-BY 4.0

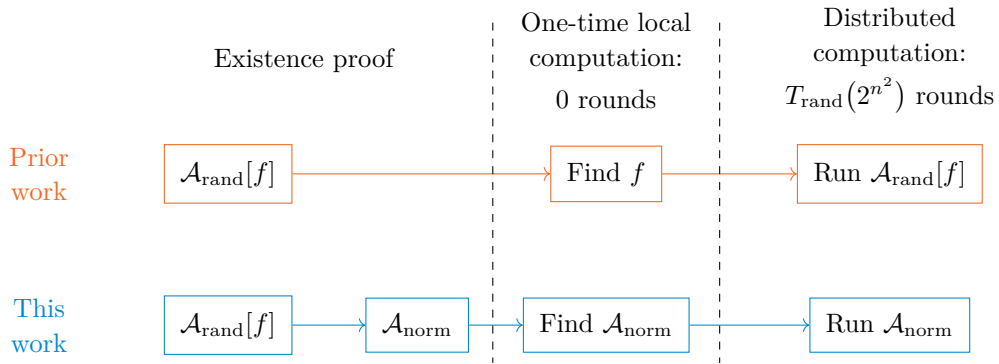
37th International Symposium on Distributed Computing (DISC 2023).

Editor: Rotem Oshman; Article No. 40; pp. 40:1–40:5



Leibniz International Proceedings in Informatics

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



■ **Figure 1** Proof strategy in this work and prior work [1].

But what do we mean, precisely, when we say that $\mathcal{A}_{\text{rand}}$ is a randomized algorithm in the LOCAL model? Chang, Kopelowitz, and Pettie [1] assume that there is some upper bound $r(n)$ on the number of random bits used by a node. This is a non-standard definition; while many reasonable algorithms naturally satisfy this, formally speaking it is not compatible with e.g. a randomized algorithm in which each node picks a number from a geometric distribution by repeated Bernoulli trials. All other results that build on Theorem 1 are also influenced by this assumption; the foundations of the field are on a bit shaky ground.

New result: unbounded randomness. In this short note we prove a stronger version of Theorem 1. Our proof does not need to assume anything about the number of random bits consumed by a node. Hence, we can now safely conclude that all corollaries of Theorem 1 also hold in the standard randomized LOCAL model, in which local computation – including the number of random bits generated – is unbounded.

Similar to [1], we assume that n and T_{rand} (or sufficiently tight bounds on them) are known. Similar to [1], the proof is constructive and \mathcal{A}_{det} is a uniform, computable, deterministic algorithm. The only difference is that we assume less about $\mathcal{A}_{\text{rand}}$.

Key new ideas. Exactly like Chang, Kopelowitz, and Pettie [1, Theorem 3.1], we start by defining $N = 2^{n^2}$. Then even though we are working in an n -node graph, we lie to $\mathcal{A}_{\text{rand}}$ that we have a graph with N nodes. The running time increases to $T_{\text{rand}}(N)$, but the success probability improves to $1 - 1/N$, which is large enough to show that there exists a mapping f from unique identifiers to random bits that works for every n -node graph.

At this point our paths deviate – see Figure 1 for an illustration. In [1, Theorem 3.1], \mathcal{A}_{det} is constructed as follows: Each node checks each possible mapping f , and picks the first one that works for every n -node graph; then \mathcal{A}_{det} simply simulates $\mathcal{A}_{\text{rand}}$ with random bits from f . This is where they make use of bounded randomness: for a fixed n there are only finitely many possible functions f to check.

We proceed as follows – instead of looking at the *internal behavior* of the algorithm we look at its *external behavior*:

1. Since a good mapping f exists, we could in principle hard-code this specific mapping to obtain a deterministic algorithm $\mathcal{A} = \mathcal{A}_{\text{rand}}[f]$. At this point we merely know that \mathcal{A} exists – this step is non-constructive, and \mathcal{A} might not even be computable.

2. However, any deterministic LOCAL algorithm can be represented in a *normal form* as a function $\mathcal{A}_{\text{norm}}$ that maps each possible $T_{\text{rand}}(N)$ -radius neighborhood to a local output. Since \mathcal{A} exists, we know that such a function $\mathcal{A}_{\text{norm}}$ also exists and solves \mathcal{P} correctly in all n -node graphs.
3. Now \mathcal{A}_{det} simply finds the first valid $\mathcal{A}_{\text{norm}}$, and then simulates $\mathcal{A}_{\text{norm}}$. This way we can construct a computable, uniform, deterministic algorithm \mathcal{A}_{det} even if we merely know that $\mathcal{A}_{\text{rand}}$ exists, and even if $\mathcal{A}_{\text{rand}}$ is non-computable or non-uniform.

Two extensions. While Theorem 1 was originally presented for LCL problems, our new proof works for a broader class of problems: we show how to handle labeling problems that are defined component-wise. The proof is given in Section 3; Theorem 1 then follows as a special case.

We also briefly discuss in Section 4 one extension: how to derandomize algorithms that are only guaranteed to work in connected graphs. A bit more care is needed when we lie about the number of nodes in that case.

2 Preliminaries

Let $G = (V, E)$ denote a simple undirected graph. For any two nodes $u, v \in V$, we denote their distance by $d(u, v)$, i.e., the number edges in a shortest path connecting u to v ; if such path does not exist, then $d(u, v) = +\infty$. Furthermore, by $\deg(v)$ we denote the degree of v , i.e., the number of incident edges.

LOCAL model. Let $G = (V, E)$ be any graph with n nodes. In the *deterministic LOCAL model*, each node $v \in V$ is given a unique identifier $\text{id}(v) \in \{1, 2, \dots, n^c\}$ for some constant $c \geq 1$. The initial knowledge of a node consists of its own identifier, its degree, the number of nodes n and (possibly) an input label. Each node runs the same algorithm and computation proceeds in synchronous rounds. In each round, nodes send messages of arbitrary size to their neighbors, then receive some messages, and then perform local computations of arbitrary complexity. After some number of rounds, a node must terminate its computation and decide on its local output. The running time (or complexity) of a distributed algorithm is defined as the number of rounds needed by all nodes to decide the local output.

In the *randomized LOCAL model*, each node is also given access to an infinite random bit stream, and the bit streams of the nodes are mutually independent. We say that an algorithm is *uniform* if the size of the description of the algorithm does not depend on n .

For any fixed locality T , the LOCAL model can also be viewed as a mapping from each radius- T neighborhood $N_T[v]$ of each node v to a local output. Here by $N_T[v]$ we mean the graph (V', E') , where $V' \subseteq V$ is the set of all nodes $u \in V(G)$ with $d(v, u) \leq T$ and E' is the set of edges $\{s, t\} \in E$ with $d(v, s) \leq T - 1$ and $d(v, t) \leq T$. Each node of $N_T[v]$ is also labeled with its original degree $\deg(u)$, unique identifier $\text{id}(u)$, local input, and – for randomized algorithms – its stream of random bits. This is exactly the information node v can gather in T rounds.

Labeling problems. Let Σ_{in} be a finite set of input labels and Σ_{out} be a finite set of admissible output labels. An *input labeling* of a graph $G = (V, E)$ is a function $\lambda_{\text{in}}: V \rightarrow \Sigma_{\text{in}}$, and an *output labeling* is a function $\lambda_{\text{out}}: V \rightarrow \Sigma_{\text{out}}$. A *labeling problem* \mathcal{P} specifies for each graph and each input labeling a set of feasible output labelings.

We say that \mathcal{P} is a *component-wise verifiable problem* if for each graph G and each connected component C of G , the set of valid output labelings restricted to C only depends on C .

Let $r \in \mathbb{N}$ be a constant. We say that \mathcal{P} is a *locally verifiable problem* with verification radius r if for each graph G and each node v of G , the set of valid output labelings restricted to $N_r[v]$ only depends on $N_r[v]$.

We note that LCL problems [3] are a special case of locally verifiable problems with a constant bound on the degree of the nodes. Locally verifiable problems are in turn a special case of component-wise verifiable problems.

3 Main result

We give the derandomization result directly for component-wise verifiable problems; Theorem 1 then follows as a corollary.

► **Theorem 2.** *Let $\mathcal{A}_{\text{rand}}$ be a randomized LOCAL algorithm that solves a component-wise verifiable problem \mathcal{P} in $T_{\text{rand}}(n)$ communication rounds in n -node graphs with probability at least $1 - 1/n$. Then there is a deterministic LOCAL algorithm \mathcal{A}_{det} that solves \mathcal{P} in $T_{\text{det}}(n)$ rounds, where $T_{\text{det}}(n) = T_{\text{rand}}(2^{n^2})$.*

Proof. Consider any sufficiently large n , and let $N = 2^{n^2}$. In what follows, we lie to algorithm $\mathcal{A}_{\text{rand}}$ that the input graph consists of N nodes. Hence, it runs in time $T := T_{\text{rand}}(N) = T_{\text{det}}(n)$ and succeeds with probability $1 - 1/N$.

Let $\mathcal{R}_n = \{f: \{0, 1\}^{c \log n} \rightarrow \{0, 1\}^{\mathbb{N}}\}$ be the family of all possible assignments of random bits streams to unique identifiers. For $f \in \mathcal{R}_n$, we write $\mathcal{A}_{\text{rand}}[f]$ to denote the *deterministic* LOCAL algorithm in which node v runs $\mathcal{A}_{\text{rand}}$ but uses $f(\text{id}(v))$ as its random bit stream. Note that $\mathcal{A}_{\text{rand}}$ is equivalent to the following process: choose $f \in \mathcal{R}_n$ uniformly at random and apply $\mathcal{A}_{\text{rand}}[f]$.

Let \mathcal{G}_n be the set of all possible inputs $(G, \text{id}, \lambda_{\text{in}})$, where G is an n -node graph, id is a unique identifier assignment, and λ_{in} is an input labeling. We know that

$$|\mathcal{G}_n| \leq 2^{\binom{n}{2}} \cdot 2^{cn \log n} \cdot |\Sigma_{\text{in}}|^n < N = 2^{n^2}$$

for a large enough n . We say that f is *good* if $\mathcal{A}_{\text{rand}}[f]$ outputs a valid solution for *every* input in family \mathcal{G}_n .

Now, we show there exists a good f . Let F be a uniform random variable over \mathcal{R}_n . Then

$$\Pr(F \text{ is bad}) \leq \sum_{G \in \mathcal{G}_n} \Pr(\mathcal{A}_{\text{rand}}[F] \text{ fails on } G) = \sum_{G \in \mathcal{G}_n} \Pr(\mathcal{A}_{\text{rand}} \text{ fails on } G) \leq \frac{|\mathcal{G}_n|}{N} < 1.$$

Therefore, $\Pr(F \text{ is good}) > 0$. Hence, there exists a good function; let f be any such function. Thus, there is a deterministic algorithm $\mathcal{A} = \mathcal{A}_{\text{rand}}[f]$ that solves \mathcal{P} on all inputs in \mathcal{G}_n in at most T rounds.

Any deterministic T -round algorithm in the LOCAL model defines a mapping $\mathcal{A}_{\text{norm}}$ from radius- T neighborhoods to local outputs. Conversely, such a mapping $\mathcal{A}_{\text{norm}}$ can be interpreted as a T -round algorithm. Furthermore, for a fixed n , there are only finitely many such mappings.

Now \mathcal{A}_{det} works as follows: Given n , each node first enumerates all candidate mappings $\mathcal{A}_{\text{norm}}$ in lexicographic order, checks if $\mathcal{A}_{\text{norm}}$ solves \mathcal{P} for every \mathcal{G}_n , and stops once the first such $\mathcal{A}_{\text{norm}}$ is found. Then \mathcal{A}_{det} uses T rounds so that each node v learns its radius- T neighborhood $N_T[v]$, and finally each node applies mapping $\mathcal{A}_{\text{norm}}$ to $N_T[v]$ to determine its local output. ◀

4 Technicality: connected graphs

In the proof of Theorem 2, a key step was that we lied about n . The algorithm cannot catch us lying, as the n -node input graph G is indistinguishable from some hypothetical N -node input graph G' in which one connected component is isomorphic to G . As \mathcal{P} was assumed to be component-wise verifiable, an algorithm that succeeds globally in G' also has to succeed locally when restricted to G .

The proof heavily exploited graphs that may consist of multiple connected components. In this section we briefly note that this is *not necessary*. We can prove the following version of Theorem 2 that holds even if $\mathcal{A}_{\text{rand}}$ only works correctly in connected graphs. However, component-wise problems are too broad class of problems in this case, and we consider locally verifiable problems instead:

► **Theorem 3.** *Let $\mathcal{A}_{\text{rand}}$ be a randomized LOCAL algorithm that solves a locally verifiable problem \mathcal{P} in $T_{\text{rand}}(n)$ communication rounds in n -node connected graphs with probability at least $1 - 1/n$. Then there is a deterministic LOCAL algorithm \mathcal{A}_{det} that solves \mathcal{P} in $O(T_{\text{det}}(n))$ rounds, where $T_{\text{det}}(n) = T_{\text{rand}}(2^{n^2})$.*

Proof. Let $t = T_{\text{det}}(n) + r$, where r is the verification radius of problem \mathcal{P} . In algorithm \mathcal{A}_{det} , each node v first explores its radius- t neighborhood to determine if the entire input graph G is contained in $N_t[v]$. If yes, we spend another t rounds to inform all nodes about G . In this case all nodes have learned G , and we can solve \mathcal{P} by brute force and stop.

Otherwise, we can proceed as we did in the proof of Theorem 2. We can now safely lie about N . To see this, assume that $\mathcal{A}_{\text{rand}}$ fails in some n -node graph G with probability more than n/N if we lie that G has N nodes. Then the algorithm also has to fail locally in the radius- r neighborhood of some node v with probability more than $1/N$. Now it is possible to construct an N -node graph G' with node v' such that radius- t neighborhood of v in G is isomorphic to the radius- t neighborhood of v' in G' (here we exploit the fact that radius- t neighborhood of v does not contain the entire graph G). As radius- t neighborhoods of v and v' agree, and the running time of $\mathcal{A}_{\text{rand}}$ is $t - r$ rounds, the output distributions of $N_r[v]$ and $N_r[v']$ also agree. Now it follows that $\mathcal{A}_{\text{rand}}$ fails locally in the radius- r neighborhood of v' in G' with probability more than $1/N$, and hence it also fails globally in G' with probability more than $1/N$, which is a contradiction with the assumption that $\mathcal{A}_{\text{rand}}$ solves \mathcal{P} in connected N -node graphs with probability at least $1 - 1/N$.

Now as long as we choose a large enough n such that $|\mathcal{G}_n| < N/n$, the rest of the proof of Theorem 2 goes through. ◀

References

- 1 Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the local model. *SIAM Journal on Computing*, 48(1):122–143, 2019. doi:10.1137/17M1117537.
- 2 Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992. doi:10.1137/0221015.
- 3 Moni Naor and Larry J. Stockmeyer. What can be computed locally? *SIAM Journal on Computing*, 24(6):1259–1277, 1995. doi:10.1137/S0097539793254571.
- 4 David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, 2000. doi:10.1137/1.9780898719772.
- 5 Jukka Suomela. Landscape of locality (invited talk). In *17th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2020)*, 2020. doi:10.4230/LIPIcs.SWAT.2020.2.