

# Distributed Testing of Distance- $k$ Colorings

Pierre Fraigniaud<sup>1\*</sup>, Magnús M. Halldórsson<sup>2</sup>, and Alexandre Nolin<sup>2\*\*</sup>

<sup>1</sup> IRIF, CNRS and Université de Paris, pierre.fraigniaud@irif.fr

<sup>2</sup> ICE-TCS, Department of Computer Science, Reykjavik University, Iceland,  
mmh@ru.is, alexandren@ru.is

**Abstract.** We study the distributed decision problem related to checking distance- $k$  coloring, defined as color assignments to the nodes such that every pair of vertices at distance at most  $k$  must receive distinct colors. While checking the validity of a distance- $k$  coloring only requires  $\lceil k/2 \rceil$  rounds in the LOCAL model, and a single round in the CONGEST model when  $k \leq 2$ , the task is extremely costly for higher  $k$ 's in CONGEST — there is a lower bound of  $\Omega(\Delta^{k/2})$  rounds in graphs with maximum degree  $\Delta$ . We therefore explore the ability of checking distance- $k$  coloring via distributed property testing. We consider several fairness criteria for measuring the distance to a valid coloring, and we derive upper and lower bounds for each of them. In particular, we show that for one natural fairness measure, significantly better algorithms are possible for testing distance-3 coloring than for testing distance- $k$  coloring for  $k \geq 4$ .

**Keywords:** distributed property testing, graph coloring, distributed decision.

## 1 Introduction

We study problems related to checking whether a given distance- $k$  coloring is proper, in the distributed CONGEST model. A valid (or proper) distance- $k$  coloring of a graph  $G = (V, E)$ , for  $k \geq 1$ , is a coloring of each node  $v$  with integer  $c_v$  so that any two nodes  $u, v$  of distance at most  $k$  are colored differently, i.e.,  $c_u \neq c_v$ . This is equivalent to the usual vertex coloring of the graph  $G^k = (V, E^k)$ , where two nodes are adjacent if they are within distance  $k$  in  $G$ .

Classical distance-1 colorings have been extensively studied in distributed computing as a tool of breaking symmetry. Let us denote by  $n$  the number of nodes, by  $m$  the number of edges, and by  $\Delta$  the maximum degree of  $G$ . For the core problem of finding a  $(\Delta+1)$ -coloring, there is a simple folklore  $O(\log n)$ -round randomized algorithm, and recent polylog( $n$ )-round deterministic algorithm by Bamberger, Kuhn and Maus [4] that

---

\* Additional support from ANR Projects DESCARTES, QuDATA, and FREDDA.

\*\* Partially supported by Icelandic Research Foundation grant 174484-051.

works in CONGEST (leveraging the recent breakthrough of Rozhoň and Ghaffari [31]). The corresponding distance-2 coloring questions has recently been addressed in [22], where an  $O(\log \Delta \log n)$ -round randomized algorithm is given that uses  $\Delta^2 + 1$  colors, as well as a  $\text{polylog}(n)$ -round deterministic algorithm that uses  $(1 + \epsilon)\Delta^2$  colors, for any  $\epsilon > 0$ . This opens the question about distance- $k$  coloring problems, for  $k \geq 3$ , which appear considerably harder.

*Why distance- $k$  coloring?* Distributed distance- $k$  colorings are interesting for various reasons. They appear naturally when constant-round randomized algorithms are derandomized using the method of conditional expectation [19]. They also appear in certain models of wireless models, where senders must be sufficiently separated, to limit interference. More abstractly, we can view distance- $k$  coloring problems as a way of studying *communication capacity constraints* on nodes, where communication must go through intermediate relays. Given that distance-2 colorings can be efficiently computed, distance-3 colorings appear to lie at the frontier of what can be solved efficiently by distributed algorithms.

*Deciding distance- $k$  coloring.* Given the apparently challenging task of *finding* an efficient distance-3 coloring, a natural question that arises is if we can at least *check* that a given coloring is valid. We can quickly dispose of that hope, as there is an easy reduction to Set Disjointness that shows that verifying a distance- $k$  coloring requires  $\Omega(\Delta^{\lfloor (k-1)/2 \rfloor})$  rounds in CONGEST. We provide a proof of this fact, for completeness, in App. A. Observe that the question is trivially answered in  $k$  rounds of the LOCAL model.

*Testing distance- $k$  coloring.* Distributed property testing is a relaxation of distributed decision, where we seek a CONGEST algorithm that can distinguish whether the given graph satisfies a given property (e.g., having a distance- $k$  coloring), or being *far* from having such a property. The most common notion for this is  $\epsilon$ -*farness* in the *sparse* model, when the addition or deletion of up to  $\epsilon \cdot m$  arbitrary edges to/from the graph  $G = (V, E)$  does not result in the property being satisfied. This notion is renamed  $\epsilon$ -*edge* in this paper, so as to avoid confusion as we use alternative notions of being far to a valid coloring. Distributed algorithms testing a property (here distance- $k$  coloring) are compared according to the *error rate*  $\epsilon(r)$  they can tolerate if restricted to  $r$  rounds, or equivalently, the *round complexity*  $r(\epsilon)$  to distinguish between legal instances and instances  $\epsilon$ -far from being legal.

## 1.1 Summary of main results

We consider several measures of distance from a valid coloring to define various notion of  $\epsilon$ -farness, deduce their relationship, and bound the efficiency of testing distance- $k$  colorings in CONGEST under these measures. As examples of such measures, we consider  $\epsilon$ -*edge*, where deleting up to  $\epsilon m$  edges cannot result in a valid distance- $k$  coloring, and  $\epsilon$ -*middle*, where there exist more than  $\epsilon m$  paths of length at most  $k-2$  between two nodes with distinct neighbors of the same color. We present the following results:

1. An algorithm for any constant  $k \geq 3$ , with round complexity  $O(1/\epsilon)$ , for all our measures but  $\epsilon$ -middle. We provide a matching lower bound for any algorithm under two of our considered measures.
2. An improved algorithm for distance-3 colorings under the  $\epsilon$ -middle measure. The round complexity is  $O(\epsilon^{-3/2}m^{-1/2})$ , for  $\epsilon \geq m^{-1/3}$ . We prove a matching lower bound, and as well as an  $\tilde{\Omega}(\epsilon^{-1})$  lower bound for distance-4. This shows that distance-4 is strictly harder than distance-3.
3. A communication complexity lower bound of  $\tilde{\Omega}(\epsilon^{-1}(\epsilon m)^{-1})$ , for any  $k \geq 3$ , under the  $\epsilon$ -edge measure.

The results suggest that distance-3 colorings are easier to test than for larger distances. This reinforces the role of distance-3 coloring on the frontier of what is computable efficiently in CONGEST.

## 1.2 Related work

Property testing has an extensive history in the sequential setting [20]. Distributed property testing was recently introduced by Brakerski and Patt-Shamir [6], and later revisited and formalized more broadly by Censor-Hillel et al. [8]. As in the centralized setting, different variants of farness can be considered, but most of the efforts on distributed property testing has been carried out in the *sparse* model, that is, the model of this paper, where farness is measured by the fraction of the *number of edges* that must be added or removed for satisfying the property under consideration. In this framework, most previous work has been dedicated to checking the absence of a specific graph pattern (e.g., a cycle  $C_k$ , or a clique  $K_k$ , for some  $k \geq 3$ ) as a subgraph of the actual network [11, 15, 17]. To our knowledge, this paper is the first to consider distributed testing proper distance- $k$  coloring.

More generally, distributed property testing falls into the wide class of *distributed decision* problems, initially motivated by *fault-tolerant* distributed computing [2, 3, 25]. Since these early work, there has been a

large body of work on distributed decision, with a range of models — see [12] for a survey. The closest to ours are *local decision* [13, 14], and *local verification* [14, 16, 21, 28]. In both cases, the nodes perform a constant number of rounds of communication before reaching a decision. Distributed property testing is a relaxed version of randomized distributed local decision, as nodes are not bounded to detect illegal instances that are “close to be legal”. In distributed verification, every node is also supplied with a *certificate* string, and the collection of certificates is supposed to form a distributed *proof* that the instance is legal. Distributed property testing performs in absence of such certificates. Recently, distributed verification has been extended to distributed *interactive proofs* [27, 29], involving interactions between the nodes and a powerful centralized oracle. Such mechanisms are obviously much more powerful than distributed property testing.

Overall, distributed property testing offers a tradeoff between simplicity (no need of certificates, nor of any interactions with an external entity), and efficiency (configurations that are “slightly” illegal may not be detected). It is thus an appealing lightweight alternative to complex mechanisms for distributed systems that can tolerate to be slightly faulty. This is typically the case of wireless systems, which are able to tolerate a certain level of interference, as long as these interferences do not exceed a certain threshold.

## 2 Model and Definitions

The input of our algorithms is a graph  $G$ , and a proposed coloring  $C = (c_v)_{v \in V}$ . Given an underlying distance metrics between solutions, we say that a solution is  $\epsilon$ -far from being correct (or valid, or legal) if it is of distance at least  $\epsilon$  from any valid solution for  $G$ . We seek a CONGEST protocol running on  $G$  to distinguish valid solutions from  $\epsilon$ -far solutions. The protocol should have 1-sided error:

If  $C$  is valid, then, with probability 1, all nodes output “yes”.

If  $C$  is  $\epsilon$ -far from being valid, then, with probability at least  $2/3$ , some node outputs “no”.

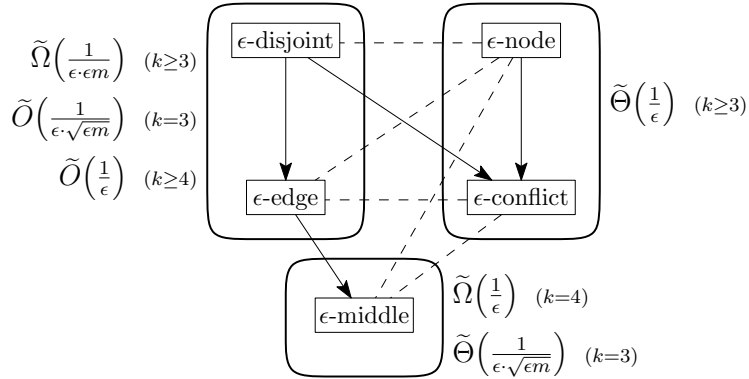
We explore different types of solution distances. In particular, we can divide them into two types: distance to a graph for which the given solution is valid, and distance to a valid solution for the given graph. We call two distinct nodes with the same color at distance at most  $k$  a *bad pair*, and call a path connecting a bad pair a *bad path*.

**Definition 1.** An  $n$ -node  $m$ -edge graph  $G = (V, E)$  and a coloring of its vertices  $(c_v)_{v \in V}$  are said to be

- $\epsilon$ -edge, when deleting up to  $\epsilon m$  arbitrary edges does not result in a valid distance- $k$  coloring.
- $\epsilon$ -disjoint, when there exist more than  $\epsilon m$  distinct pairs of similarly colored vertices linked by edge-disjoint paths of length at most  $k$ .
- $\epsilon$ -middle, when there exist more than  $\epsilon m$  paths of length at most  $k-2$  between two nodes with distinct neighbors of the same color.
- $\epsilon$ -node, when recoloring up to  $\epsilon n$  arbitrary vertices does not result in a valid distance- $k$  coloring.
- $\epsilon$ -conflict, when more than  $\epsilon n$  vertices have the same color as one of their distance- $k$  neighbors.

The  $\epsilon$ -edge measure is the classical one from property testing literature [6, 8, 15]. The  $\epsilon$ -disjoint measure is a variation that requires there to be many conflict pairs, not just one vertex that conflicts with many nodes (that might not conflict between themselves).

The  $\epsilon$ -middle measure has the appearance of being contrived, but actually captures the essence of the problem. In the first round, each node learns of the colors of *all* its neighbors. Thus, what we really need is to somehow connect the second node on a bad  $k$ -path to the second-to-last node, and see if the sets of colors in their neighborhoods intersect.




**Fig. 1.** The relationships between our notions of distance from a valid solution as well as our upper and lower bounds on the costs of testing for them. An arrow from  $\epsilon$ -notion<sub>1</sub> to  $\epsilon$ -notion<sub>2</sub> indicates that a solution that is  $\epsilon$ -notion<sub>1</sub> is also  $\Omega(\epsilon)$ -notion<sub>2</sub> away from a valid solution. Dashed lines indicate incomparability.

The last two definitions correspond to natural measures of invalidity of colorings. The  $\epsilon$ -conflict measure counts how many nodes are improperly colored (i.e., have a same-colored distance- $\leq k$  neighbor), while  $\epsilon$ -node is more conservative, bounding the number of recolorings needed to turn the coloring into a valid one.

We say that a measure  $\mu$  is more *strict* than measure  $\mu'$  if  $\mu(G, c) = O(\mu'(G, c))$ , for all graphs  $G$  and colorings  $c$ . Thus, if  $(G, c)$  is  $\epsilon$ -far in terms of measure  $\mu'$ , then it is  $O(\epsilon)$ -far in terms of measure  $\mu$  (but could be much less far).

It is easy to see that  $\epsilon$ -disjoint is more strict than  $\epsilon$ -edge, and  $\epsilon$ -node is more strict than  $\epsilon$ -conflict. It also holds that  $\epsilon$ -disjoint is more strict than  $\epsilon$ -conflict on sparse graphs, when  $|E(G)| = O(|V(G)|)$ . This is illustrated in Fig. 2, where solid arrows are drawn from a stricter measures to a less stricter one.

We can also verify that other pairs of measures can be arbitrarily divergent. The examples in Fig. 2 show that for any pair with a dotted line, there is a graph where one is constant and the other is  $O(1/n)$  (or  $O(1/m)$ ), and vice versa. The same holds for the inverse direction of the solid edges.



**Fig. 2.** Three colored graphs showing the incomparability of some of the measures of Definition 1. A number indicates a node’s color, unnumbered nodes each receive a color unique to them. For each graph, the second line indicates values of  $\epsilon$  for which the graph is (in order)  $\epsilon$ -disjoint,  $\epsilon$ -node,  $\epsilon$ -edge,  $\epsilon$ -conflict, and  $\epsilon$ -middle.

The property of  $\epsilon$ -edge and  $\epsilon$ -disjoint assignments that we shall use is that there is a set of at least  $\epsilon m$  edges, each of which is the first edge of a bad path. For  $\epsilon$ -node or  $\epsilon$ -conflict assignment, it follows from the definition that there is a set of  $\epsilon n$  nodes that have a same-colored node within distance at most  $k$ .

### 3 Preliminaries: Set Disjointness

The *set disjointness* problem is a two-party communication complexity decision problem where two players each receive a subset of an universe  $[N]$  and must decide whether their subsets are disjoint. This problem is known to require  $\Omega(N)$  communication – as large as the players’ inputs – to solve with bounded error by a randomized communication complexity protocol [26, 30, 5]. Doing a reduction from set disjointness to a task in the CONGEST model has been a fruitful source of lower bounds [32, 18, 24, 10, 1, 9].

In this paper, we will use slight variations of the original set disjointness problem. We consider a subset of the original problem, where the players have two additional promises: that their sets are of size at most  $s$ , and that their sets’ intersection is either empty or contains at least  $t$  elements, where  $s$  and  $t$  are two integer parameters.

**Definition 2 (Large intersection, bounded size set disjointness).**

Let  $N, s, t$  be three integers such that  $N \geq s \geq t > 0$ ,  $\mathcal{X} = \mathcal{Y} = [N]$ , and the players’ set of admissible inputs  $\mathcal{I}_{s,t} \subseteq \mathcal{X} \times \mathcal{Y}$  be:

$$\mathcal{I}_{s,t} = \{(X, Y) : |X| \leq s, |Y| \leq s, |X \cap Y| \in \{0\} \cup [t, +\infty)\}$$

The large intersection, bounded size set disjointness problem  $\mathbf{DISJ}_{s,t}^N : \mathcal{I}_{s,t} \rightarrow \{0, 1\}$  is defined as:

$$\mathbf{DISJ}_{s,t}^N(X, Y) = \begin{cases} 1 & \text{if } X \cap Y = \emptyset, \\ 0 & \text{otherwise.} \end{cases}$$

The standard set disjointness problem corresponds to the choice of parameters  $s = N, t = 1$ . A commonly studied variant bounds the size of the player’s sets but promises nothing about the intersection ( $t = 1$ ). This problem is known to have randomized communication complexity  $\Theta(s)$  [23]. Computing the intersection of the two sets also has randomized communication complexity only  $\Theta(s)$  [7]. In both cases, the lower bound is a simple consequence of the lower bound for the standard set disjointness problem.

**Lemma 1.**  $R_{1/3}(\mathbf{DISJ}_{s,t}^N) \in \Omega\left(\frac{s}{t}\right)$

*Proof.* Consider the  $\mathbf{DISJ}_{s/t,1}^{N/t}$  set disjointness problem. This is known to require  $\Omega(s/t)$  communication as the standard set disjointness problem  $\mathbf{DISJ}_{s/t,1}^{s/t}$  reduces to it (it is the same problem but on a subset of its input

space). Now remark that  $\mathbf{DISJ}_{s/t,1}^{N/t}$  reduces to  $\mathbf{DISJ}_{s,t}^N$ , as the players can construct a valid input to  $\mathbf{DISJ}_{s,t}^N$  from a  $\mathbf{DISJ}_{s/t,1}^{N/t}$  input by making  $t$  copies of each of their set elements, which concludes the proof.  $\square$

Our lower bounds on testing a distance-3 coloring use the  $\mathbf{DISJ}_{s,t}^N$  problem with parameters  $s \in \Theta(m)$  and  $t \in \Theta(\epsilon m)$  (Theorem 3),  $s \in \Theta(m)$  and  $t \in \Theta(\epsilon m)$  (Theorem 4), and  $s \in \Theta(\sqrt{m/\epsilon})$  and  $t = 1$  (Theorem 5), while our lower bound on testing a distance-4 coloring (Theorem 6) uses parameters  $s = m$  and  $t = 1$ . Our lower bound on verifying a distance- $k$  coloring for an arbitrary  $k$  (Theorem 7) uses parameters  $s \in \Theta(\Delta^{\lfloor (k-1)/2 \rfloor})$  and  $t = 1$ . Notice that the complexity of the set disjointness problem does not depend on  $N$  – the size of the universe – but only on  $s$  and  $t$ , the sizes of the input sets and their potential intersection.

## 4 Testing Distance- $k$ Colorings

For all the measures previously introduced (Definition 1), we give upper and lower bounds on detecting being  $\epsilon$ -far from a solution. Our first result is a protocol for all measures except  $\epsilon$ -middle, and for any constant  $k$  (Theorem 1). This protocol is later shown to be tight for the  $\epsilon$ -node and  $\epsilon$ -conflict models, even for  $k = 3$ . In the case  $k = 3$ , we give a more efficient algorithm in the  $\epsilon$ -middle,  $\epsilon$ -edge and  $\epsilon$ -disjoint models (Theorem 2). We prove that the algorithm is optimal for the  $\epsilon$ -middle measure (Theorem 5) and also prove a non-matching lower bound for the  $\epsilon$ -far and  $\epsilon$ -disjoint measures (Theorem 3). For  $k = 4$ , we prove an  $\tilde{\Omega}(\epsilon^{-1})$  lower bound in the  $\epsilon$ -middle model. This last lower bound is strictly higher than the complexity of the same problem when  $k = 3$ , demonstrating that the complexity of the problem can keep increasing as we increase  $k$  beyond 3 not just when doing verification, but also property testing.

All the lower bounds use the set disjointness problem (Definition 2). For a graphical summary of the results of this section, see Figure 2.

### 4.1 A general algorithm for any $k$

**Theorem 1.** *There exists a randomized CONGEST algorithm running in  $O(\frac{1}{\epsilon})$  rounds for testing an  $\epsilon$ -edge distance- $k$  coloring. By extension this also applies to  $\epsilon$ -disjoint, and a slight modification yields the same result for  $\epsilon$ -node and  $\epsilon$ -conflict.*

*Proof.* Consider the following basic algorithm BFS that runs for  $k$  rounds, which we then repeat to obtain success probability  $2/3$ . The edges are



independently assigned a random priority (such as a random value from  $[|E|^3]$ , with higher values receiving precedence). Nodes use the max of the priorities of their incident edges as their own priority. In the initial round, the each node transmits its color (along with its ID) to all its neighbors, along with its priority. In each subsequent round, the algorithm transmits to each neighbor the color and priority of the two highest priority color it received in the previous round. Effectively, the color from a highest priority node gets forwarded along a breadth-first-search tree. If at the end of round  $k$ , the algorithm has received a color (from another node) that matches its own, it outputs 'invalid'; otherwise, it outputs 'valid'.

In an  $\epsilon$ -edge graph, there are at least  $\epsilon m$  edges that are the first edge of a bad path. If any of those edges receives the highest priority in a round of the basic algorithm, a color conflict gets detected. So with probability at least  $\epsilon$ , the basic algorithm detects an  $\epsilon$ -edge graph

This basic algorithm is then repeated to increase the success probability to at least  $2/3$ . It suffices to repeat it  $t$  times, where  $t$  satisfies  $(1 - \epsilon)^t \leq 1/3$ . The time complexity is then  $t \cdot k$ . Since  $(1 - \epsilon)^t \leq e^{-\epsilon t}$ , setting  $t = \ln(3)/\epsilon$  achieves the desired result, yielding an  $O(k/\epsilon)$ -round algorithm.

We can simplify and adapt this algorithm for the  $\epsilon$ -conflict model: Each *node* picks a random priority. There are now  $\epsilon n$  improperly colored nodes, and if any of them gets selected, the coloring will be found to be invalid. The rest of the argument is the same.  $\square$

## 4.2 A better running time for $k = 3$

**Theorem 2.** *There exists a randomized CONGEST algorithm running in  $O\left(\frac{1}{\epsilon \cdot \sqrt{\epsilon m}}\right)$  rounds for testing an  $\epsilon$ -middle distance-3 coloring. By extension, this also applies to the  $\epsilon$ -edge and  $\epsilon$ -disjoint measures.*

*Proof.* Let the nodes follow the following simple algorithm RANDOM: in the first round, each node informs its neighbors of its color and identifier, and in  $k-2$  subsequent rounds, for each link a node has, it picks uniformly at random one of the (color, ID) pair it received from its neighbors in the previous round and sends it on this link. Any node that receives the same color twice but with a different ID, and such that it received the pairs in two (not necessarily distinct) rounds  $i$  and  $j$  such that  $i + j \leq k$ , flags the coloring as invalid. This  $(k-1)$ -rounds protocol is repeated  $T$  times. Let us analyze the probability of success of this protocol when  $k = 3$ .

Let  $\sigma = 4\sqrt{m/\epsilon}$ . We say that an edge  $uv$  is *good* if  $\min(d(u), d(v)) \leq \sigma$ , and *bad* otherwise. In an  $\epsilon$ -middle graph, there is a set  $\Pi$  of edges, each

of which is on a 3-path between same-colored nodes, with  $|II| \geq \epsilon m/3$ . Observe that if  $a, u, v, b$  is a path where  $a$  and  $b$  have the same color, then this will be detected if either  $u$  forwards the ID of  $a$  to  $v$  or  $v$  forwards the ID of  $b$  to  $u$ . The probability  $\bar{p}_{P_e}$  of non-detection along a path  $P_e$  with middle edge  $e = (u, v) \in II$  is therefore  $(1 - 1/d(u))^T \cdot (1 - 1/d(v))^T \leq e^{-T/\min(d(u), d(v))}$ . We say that a path  $P_e \in II$  is *good* if  $e$  is good. Let  $II' \subseteq II$  be the set of good paths. Let  $B$  be the set of nodes with degree at least  $\sigma$ . There are at most  $\sqrt{\epsilon m}/2$  nodes in  $B$ , as otherwise the total number incidences on nodes in  $B$  would exceed  $2m$ . Thus, there are at most  $\binom{|B|}{2} \leq \epsilon m/8$  edges with both endpoints in  $B$ . Hence, there are at least  $5\epsilon m/24$  good paths in  $II'$ .

The probability that none of those good paths detect a conflict in the color assignment is:

$$\prod_{P \in II'} \bar{p}_P \leq \exp\left(\frac{-T \cdot |II'|}{\sigma}\right) \leq \exp\left(-\frac{5}{96}T \cdot \epsilon^{3/2}m^{1/2}\right)$$

Therefore, running the protocol for  $T \in O(\epsilon^{-3/2}m^{-1/2})$  is enough to solve the problem with probability at least  $2/3$ .  $\square$

In particular, this protocol runs in constant time when  $\epsilon \in \Omega(m^{-1/3})$ . We give a matching lower bound later in the paper (Theorem 5). This algorithm is also able to detect  $\epsilon$ -edge and  $\epsilon$ -disjoint graphs with the same running time because of the relationships that exist between the measures, however the lower bounds we have for these measures are weaker (Theorem 3) and do not match our upper bound.

### 4.3 Lower bounds for $k \geq 3$

In this section, we prove lower bounds for the detection of  $\epsilon$ -disjoint colored graphs (Theorem 3),  $\epsilon$ -node colored graphs (Theorem 4) and  $\epsilon$ -middle colored graphs (Theorems 5 and 6) in the CONGEST model. By the relationships that exist between the separation measures of Definition 1, the lower bound on detecting an  $\epsilon$ -disjoint coloring also holds for  $\epsilon$ -edge colorings. Similarly, the lower bound on the detection of  $\epsilon$ -node colorings also holds for  $\epsilon$ -conflict colorings.


All lower bounds use the same following classical proof architecture: we take a two-party communication complexity problem  $f$  of communication complexity  $R^{cc}(f)$ , and show that the players can solve an instance  $f(x, y)$  of this problem by simulating a CONGEST algorithm for our testing task on a graph  $G_{x, y}$  with color assignment  $(c_v^{x, y})_{v \in V}$ . The vertices of

$G_{x,y}$  are partitioned into two sets  $V_A$  and  $V_B$ , and the edges are such that the colors and intraconnexions of  $V_A$ 's vertices only depend on  $x$ , and similarly with  $V_B$ 's vertices and  $y$ , while the interconnexions between a vertices of  $V_A$  and  $V_B$  are fixed and therefore independent of  $x$  and  $y$ . Let  $T$  be the number of rounds of a CONGEST algorithm for the CONGEST task, and  $C$  the number of edges between vertices of  $V_A$  and  $V_B$ . Simulating the CONGEST algorithm in the two-party communication complexity model can be done in  $T \cdot C \cdot \log(n)$  bits of communication. This last quantity has to exceed  $R^{cc}(f)$ , which yields that any CONGEST algorithm for our testing task requires at least  $T \geq \frac{R^{cc}(f)}{C \cdot \log(n)}$  rounds.

**Theorem 3.** *For  $k \geq 3$ , testing whether a distance- $k$  coloring is  $\epsilon$ -disjoint requires  $\tilde{\Omega}\left(\frac{1}{\epsilon \cdot (\epsilon m)}\right)$  rounds in the CONGEST model.*

Note that this lower bound matches neither our general upper bound (Theorem 1) nor our upper bound for  $k = 3$  (Theorem 2), leaving open the possibility of more efficient algorithms or stronger lower bounds.

For this lower bound, we consider graphs of the form presented in Figure 3. We conjecture that our analysis is not tight, and that detecting whether such graphs are  $\epsilon$ -disjoint actually requires  $\tilde{\Omega}(\epsilon^{-3/2} m^{-1/2})$ .



**Fig. 3.** The graph we use for our lower bound. It consists of 4 layers, with the outer layers having  $(\frac{1}{2} - \epsilon)m$  vertices and the inner layers  $\sqrt{2\epsilon m}$ . There only exist edges between adjacent layers, and vertices in layers 1 and 4 have degree 1, while layers 2 and 3 form a biclique. Layers 1 and 2 are randomly connected together by  $(\frac{1}{2} - \epsilon)m$  edges, as are layers 3 and 4.

*Proof.* Let  $m$  be an integer and  $\epsilon \in [\frac{1}{m}, \frac{1}{2})$ . Set  $N = m$ ,  $s = (\frac{1}{2} - \epsilon)m$ ,  $t = 2\epsilon m + 3$ , and consider an instance of  $\text{DISJ}_{s,t}^N$ : a pair of sets  $(X, Y)$ ,  $X, Y \subseteq [m]$ . Let us consider the graph  $G_{x,y} = (V, E)$  of Figure 3. Its vertices  $V$  are partitioned into four layers  $(V_i)_{i \in [4]}$ . Let Alice possess the two leftmost layers ( $V_A = V_1 \cup V_2$ ) and Bob possess the two rightmost layers ( $V_B = V_3 \cup V_4$ ). The inner layers are of size  $|V_2| = |V_3| = \sqrt{2\epsilon m}$

and form a biclique (complete bipartite graph) of  $2\epsilon m$  edges, while the outer layers are of size  $|V_1| = |V_4| = (\frac{1}{2} - \epsilon)m$ . Let us first describe how the players color their vertices, before describing how they connect the outer layers to the inner layers.

Let all of Alice's and Bob's vertices be initially uncolored. For each element  $x \in X \subseteq [m]$ , Alice picks an arbitrary uncolored vertex of  $V_1$  and colors it with  $x$ . Bob does the same with his input set  $Y$  and the layer  $V_4$ . Alice then colors her remaining uncolored vertices with distinct even numbers from  $[m+1, 2m]$ , while Bob colors his remaining uncolored vertices with distinct odd numbers from  $[m+1, 2m]$ .

Then, for each vertex  $u \in V_1$ , Alice connects it to a single vertex of  $V_2$  picked uniformly at random. Bob similarly connects vertices of  $V_4$  to vertices of  $V_3$ .

Let us now analyze the graph we constructed with respect to the Disjointness instance we started with. If  $X \cap Y = \emptyset$ , the way the players assigned colors ensures that the graph received a valid distance-3 coloring. If  $|X \cap Y| \geq t$ , however, there are  $t$  pairs  $(u, v) \in V_1 \times V_4$  of distinct vertices that are at distance 3 and received the same color. For each pair, there is a single length-3 path connecting them, and the only way those paths can share an edge is by sharing an edge in  $V_2 \times V_3$ . Let us prove that with high probability, more than  $\epsilon m$  of those paths are edge-disjoint, and therefore the graph is  $\epsilon$ -disjoint (see Definition 1).

Let  $S$  be an  $\epsilon m$ -sized subset of the  $2\epsilon m$  edges between  $V_2$  and  $V_3$ . The probability that none of those edges are directly connected to two vertices in layers  $V_1$  and  $V_4$  that received the same color is at most  $\left(1 - \frac{|S|}{2\epsilon m}\right)^t = 2^{-t}$ . As there are  $\binom{2\epsilon m}{\epsilon m} \leq 2^{2\epsilon m}$  such subsets  $S$ , the probability that less than  $\epsilon m$  edges of  $V_2 \times V_3$  are part of a length-3 path between similarly colored vertices of the outer layers is at most  $2^{2\epsilon m - t} \leq \frac{1}{8}$  for our choice of  $t$ .


Since the graph the players constructed is well-colored when they received disjoint sets, and  $\epsilon$ -disjoint with probability  $\geq 7/8$  when they received intersecting sets, the players can solve the Set Disjointness problem with error at most  $1/4$  by simulating a CONGEST algorithm to detect an  $\epsilon$ -disjoint distance-3 coloring that makes an error at most  $1/8$ . Since there are  $2\epsilon m$  edges between Alice's and Bob's vertices, the number of rounds  $T$  of a CONGEST algorithm detecting an  $\epsilon$ -disjoint coloring with probability  $\geq 7/8$  satisfies:

$$T \geq \frac{R_{1/4}^{cc}(\mathbf{DISJ}_{m/2, 2\epsilon m}^m)}{2\epsilon m \log(m)} \in \tilde{\Omega}\left(\frac{1}{\epsilon \cdot (\epsilon m)}\right)$$

□

Note that since a graph that is  $\epsilon$ -disjoint from a valid solution is also  $\epsilon$ -edge, the lower bound also applies to testing being  $\epsilon$ -edge. As corollary, we have that no constant-round algorithm can detect an  $\epsilon$ -disjoint coloring when  $\epsilon \in o(m^{-1/2})$ .

**Theorem 4.** *For  $k \geq 3$ , testing whether a distance- $k$  coloring is  $\epsilon$ -node requires  $\tilde{\Omega}\left(\frac{1}{\epsilon}\right)$  rounds in the CONGEST model.*



**Fig. 4.** The graph we use for our lower bound for the  $\epsilon$ -node measure. It consists of two stars of degree  $(\frac{n}{2} - 1)$ , linked by their roots.


The proof of this theorem and a later theorem are deferred to Appendix B. Note that contrary to our previous theorem for detecting  $\epsilon$ -disjoint colored graphs (Theorem 3), this lower bound is tight with respect to our first algorithm (Theorem 1).

**Theorem 5.** *For  $k = 3$ , testing whether a distance- $k$  coloring is  $\epsilon$ -middle requires  $\tilde{\Omega}\left(\frac{1}{\epsilon\sqrt{\epsilon m}}\right)$  rounds in the CONGEST model.*

Note that this lower bound matches our upper bound for  $k = 3$  (Theorem 2). For this lower bound, we consider graphs of the form presented in Figure 5.

*Proof.* Let  $m$  be an integer and  $\epsilon \in [\frac{1}{m}, \frac{1}{2}]$ . Set  $N = s = \frac{1-\epsilon}{2}\sqrt{m/\epsilon}$ ,  $t = 1$ , and consider an instance of  $\mathbf{DISJ}_{s,t}^N$ : a pair of sets  $(X, Y)$ ,  $X, Y \subseteq [m]$ . Consider the four layer graph  $G_{x,y} = (V, E)$  of Figure 5. The vertices  $V_2$  and  $V_3$  of layers 2 and 3 form a biclique. Every vertex of layer 2 is connected to  $s$  degree-1 vertices in layer 1, and layers 3 and 4 are similarly connected.

Let Alice possess as  $V_A$  the vertices of layers 1 and 2 and Bob possess the rest. For any vertex  $v \in V_2$ , let  $N_1(v) \subseteq V_1$  be the vertices of layer 1



**Fig. 5.** The graph we use for our lower bound on testing for  $\epsilon$ -middle colorings (Theorem 5), showing that the algorithm RANDOM is tight for this measure and  $k = 3$ . It consists of a central biclique between two layers of size  $\sqrt{\epsilon m}$ , and each vertex of these layers is connected to  $\sqrt{m/\epsilon}$  leaves in the outer layers.

connected to  $v$ , and similarly for any vertex  $v \in V_3$ , consider  $N_4(v)$  the vertices of layer 4 connected to  $v$ .

For each  $v \in V_2$ , Alice colors the nodes of  $N_1(v)$  with the elements of  $X$  as colors (without repetition, leaving uncolored nodes if necessary). Bob does the same with nodes of  $N_4(v)$  for each  $v \in V_3$ . Alice then colors her remaining uncolored colors with even numbers (avoiding distance-3 conflicts between her nodes); Bob does the same with even numbers on his side.

If the players received disjoint sets, the resulting graph  $G_{X,Y}$  is well-colored. If the sets' intersect, however, the coloring is  $\epsilon$ -middle, because for each pair of vertices  $(u, v) \in V_2 \times V_3$ , there exists a pair of vertices  $(u', v') \in N_1(u) \times N_4(v)$  that have the same color. Therefore, the players can solve their set disjointness instance by simulating a CONGEST algorithm for detection of  $\epsilon$ -middle colored graphs. Since there are  $\epsilon m$  edges between Alice's and Bob's parts of the graph, the number of rounds  $T$  of a CONGEST algorithm detecting an  $\epsilon$ -middle coloring with probability  $\geq 1/3$  satisfies:

$$T \geq \frac{R_{1/3}^{cc}(\text{DISJ}_{\frac{1-\epsilon}{2}\sqrt{m/\epsilon}, 1}^N)}{\epsilon m \cdot \log(m)} \in \tilde{\Omega}\left(\frac{1}{\epsilon\sqrt{\epsilon m}}\right)$$

□

Finally, we prove a lower bound on testing a distance-4 coloring in the  $\epsilon$ -middle model. The lower bound we obtain is strictly higher than our upper bound on the same task with distance-3, which shows that there is a clear gap between distance-3 and distance-4 colorings.

**Theorem 6.** *Testing whether a distance-4 coloring is  $\epsilon$ -middle requires  $\tilde{\Omega}\left(\frac{1}{\epsilon}\right)$  rounds in the CONGEST model.*

## 5 Conclusion

In this work, we studied the testing and verification of distance- $k$  colorings in the CONGEST model for  $k \geq 3$  and several notions of distance from a valid solution. We showed that the testing of distance-3 colorings admits a significantly more efficient algorithm than distance-4 for one of our measures ( $\epsilon$ -middle), and gave indications that it might also be the case for the other edge- and path-based measures. The node-based measures show no such gap. Our work does not give a full picture how the complexity of the problem evolves as  $k$  increases in the edge- and path-based models. A first open question is finding the exact complexity of testing in the  $\epsilon$ -disjoint and  $\epsilon$ -edge model: we conjecture that this complexity matches that of our algorithm for these models, rather than that of our lower bound or something intermediate.

Another open question is what algorithm we can design in the  $\epsilon$ -middle model for arbitrary  $k$ , as the BFS algorithm does not function in it. Even tackling the case  $k = 4$  is of interest, potentially to match our lower bound. Finally, the several measures we introduced to study this problem might be of independent interest. Are there other problems for which the same measures would make sense? A natural candidate here is testing edge-colorings.

## References

1. Abboud, A., Censor-Hillel, K., Khoury, S.: Near-linear lower bounds for distributed distance computations, even in sparse networks. In: DISC, pp. 29–42 (2016)
2. Afek, Y., Kutten, S., Yung, M.: The local detection paradigm and its application to self-stabilization. *Theor. Comput. Sci.* **186**(1-2), 199–229 (1997)
3. Awerbuch, B., Patt-Shamir, B., Varghese, G.: Self-stabilization by local checking and correction (extended abstract). In: FOCS, pp. 268–277 (1991)
4. Bamberger, P., Kuhn, F., Maus, Y.: Efficient deterministic distributed coloring with small bandwidth. *CoRR* **abs/1912.02814** (2019)
5. Bar-Yossef, Z., Jayram, T.S., Kumar, R., Sivakumar, D.: An information statistics approach to data stream and communication complexity. *J. Comput. System Sci.* **68**(4), 702–732 (2004)
6. Brakerski, Z., Patt-Shamir, B.: Distributed discovery of large near-cliques. *Distributed Computing* **24**(2), 79–89 (2011)
7. Brody, J., Chakrabarti, A., Kondapally, R., Woodruff, D.P., Yaroslavtsev, G.: Beyond set disjointness: the communication complexity of finding the intersection. In: PODC, pp. 106–113 (2014)
8. Censor-Hillel, K., Fischer, E., Schwartzman, G., Vasudev, Y.: Fast distributed algorithms for testing graph properties. *Distributed Computing* **32**(1), 41–57 (2019)
9. Censor-Hillel, K., Khoury, S., Paz, A.: Quadratic and near-quadratic lower bounds for the CONGEST model. In: DISC, pp. 10:1–10:16 (2017)

10. Drucker, A., Kuhn, F., Oshman, R.: On the power of the congested clique model. In: PODC, pp. 367–376 (2014)
11. Even, G., Fischer, O., Fraigniaud, P., Gonen, T., Levi, R., Medina, M., Montealegre, P., Olivetti, D., Oshman, R., Rapaport, I., Todinca, I.: Three notes on distributed property testing. In: DISC, pp. 15:1–15:30 (2017)
12. Feuilloley, L., Fraigniaud, P.: Survey of distributed decision. arXiv preprint arXiv:1606.04434 (2016)
13. Fraigniaud, P., Göös, M., Korman, A., Parter, M., Peleg, D.: Randomized distributed decision. *Distributed Computing* **27**(6), 419–434 (2014)
14. Fraigniaud, P., Korman, A., Peleg, D.: Towards a complexity theory for local distributed computing. *J. ACM* **60**(5), 35:1–35:26 (2013)
15. Fraigniaud, P., Olivetti, D.: Distributed detection of cycles. *ACM Trans. Parallel Comput.* **6**(3) (2019)
16. Fraigniaud, P., Patt-Shamir, B., Perry, M.: Randomized proof-labeling schemes. *Distributed Computing* **32**(3), 217–234 (2019)
17. Fraigniaud, P., Rapaport, I., Salo, V., Todinca, I.: Distributed testing of excluded subgraphs. In: DISC, pp. 342–356 (2016)
18. Frischknecht, S., Holzer, S., Wattenhofer, R.: Networks cannot compute their diameter in sublinear time. In: SODA, pp. 1150–1162 (2012)
19. Ghaffari, M., Harris, D.G., Kuhn, F.: On derandomizing local distributed algorithms. In: FOCS, pp. 662–673 (2018)
20. Goldreich, O. (ed.): Property Testing - Current Research and Surveys, *Lecture Notes in Computer Science*, vol. 6390. Springer (2010)
21. Göös, M., Suomela, J.: Locally checkable proofs in distributed computing. *Theory of Computing* **12**(1), 1–33 (2016)
22. Halldórsson, M.M., Kuhn, F., Maus, Y.: Distance-2 coloring in the CONGEST model (2020). Unpublished manuscript, available at <https://bit.ly/2TRfQ3b>
23. Hästad, J., Wigderson, A.: The randomized communication complexity of set disjointness. *Theory of Computing* **3**(11), 211–219 (2007)
24. Holzer, S., Wattenhofer, R.: Optimal distributed all pairs shortest paths and applications. In: PODC, pp. 355–364 (2012)
25. Itkis, G., Levin, L.A.: Fast and lean self-stabilizing asynchronous protocols. In: FOCS, pp. 226–239 (1994)
26. Kalyanasundaram, B., Schnitger, G.: The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math* **5**(4), 545–557 (1992)
27. Kol, G., Oshman, R., Saxena, R.R.: Interactive distributed proofs. In: PODC, pp. 255–264 (2018)
28. Korman, A., Kutten, S., Peleg, D.: Proof labeling schemes. *Distributed Computing* **22**(4), 215–233 (2010)
29. Naor, M., Parter, M., Yogev, E.: The power of distributed verifiers in interactive proofs. In: SODA, pp. 1096–115 (2020)
30. Razborov, A.A.: On the distributional complexity of disjointness. *Theor. Comp. Sci.* **106**, 385–390 (1992)
31. Rozhoň, V., Ghaffari, M.: Polylogarithmic-time deterministic network decomposition and distributed derandomization. *CoRR* **abs/1907.10937** (2019)
32. Sarma, A.D., Holzer, S., Kor, L., Korman, A., Nanongkai, D., Pandurangan, G., Peleg, D., Wattenhofer, R.: Distributed verification and hardness of distributed approximation. *SIAM J. Comput.* **41**(5), 1235–1265 (2012)




## A Verifying Distance- $k$ Colorings in Bounded-Degree Graphs

### A.1 A matching lower bound for the natural algorithm

In a graph of maximum degree  $\Delta$ , the nodes can learn their distance- $\lceil \frac{k}{2} \rceil$  neighborhood in  $O(\Delta^{\lceil k/2 \rceil - 1})$  rounds in CONGEST. In particular, an invalid distance- $k$  coloring can be detected with this number of rounds in CONGEST, since two nodes of distance at most  $k$  are both within a distance  $\lceil \frac{k}{2} \rceil$  of some node. This protocol is actually close to optimal, as our next theorem shows.

**Theorem 7.** *For  $k \geq 3$ , checking a distance- $k$  coloring requires  $\tilde{\Omega}(\Delta^{\lceil k/2 \rceil - 1})$  rounds in the CONGEST model.*



**Fig. 6.** The graph we use for our lower bound. It consists of 2 complete  $(\Delta - 1)$ -ary trees of depth  $\lfloor k/2 \rfloor - 1$  linked at their roots.

*Proof.* Set  $N = 2\Delta^{\lceil k/2 \rceil}$ ,  $s = (\Delta - 1)^{\lceil k/2 \rceil - 1}$ ,  $t = 1$ , and let us consider an instance  $(X, Y)$  of the Set Disjointness problem  $\text{DISJ}_{s,t}^N$  and show that Alice and Bob can simulate a CONGEST protocol for verifying a distance- $k$  coloring to solve their Disjointness instance.

Let us consider the graph of Figure A.1 and let Alice own the vertices of the left tree while Bob owns the vertices of the right tree. Let Alice color her vertices with the elements of her set  $X$  as distinct colors, and with even numbers greater than  $N$  if she runs out of colors. Bob does the same on his side with his set  $Y$  and odd numbers greater than  $N$ . The resulting coloring of the graph is a valid distance- $k$  coloring if and only if the sets  $X$  and  $Y$  do not intersect, so the players may simulate a CONGEST protocol on this graph to solve their set disjointness instance. Since there is only one edge between Alice's and Bob's vertices, simulating a round of the CONGEST protocol here only requires  $\log(n)$  bits of communication,

therefore the number of rounds  $T$  that a CONGEST protocol needs for verifying a distance- $k$  coloring in a graph of maximum degree  $\Delta$  satisfies:

$$T \geq \frac{R_{1/3}^{cc}(\mathbf{DISJ}_{s,t}^N)}{\log(n)} \in \Omega\left(\frac{\Delta^{\lceil k/2 \rceil - 1}}{k \cdot \log(\Delta)}\right)$$

□

## B Missing Proofs

**Theorem 4.** *For  $k \geq 3$ , testing whether a distance- $k$  coloring is  $\epsilon$ -node requires  $\tilde{\Omega}\left(\frac{1}{\epsilon}\right)$  rounds in the CONGEST model.*

*Proof.* Let  $n$  be an integer and  $\epsilon \in \left[\frac{1}{n}, \frac{1}{2}\right)$ . Set  $N = n$ ,  $s = n/2$ ,  $t = \epsilon n$ , and consider an instance of  $\mathbf{DISJ}_{s,t}^N$ : a pair of sets  $(X, Y)$ ,  $X, Y \subseteq [n]$ . Consider the graph  $G_{x,y} = (V, E)$  of Figure 4, constructed by taking two stars of degree  $\frac{n}{2} - 1$  and connecting their roots. Let Alice possess the  $n/2$  vertices  $V_A$  of one of the stars and Bob possess the vertices of the other star.

Let all of Alice's and Bob's vertices be initially uncolored. For each element  $x \in X \subseteq [n]$ , Alice picks an arbitrary uncolored vertex of  $V_A$  and colors it with  $x$ . Bob does the same with his input set  $Y$  and  $V_B$ . Alice then colors her remaining uncolored vertices with distinct even numbers from  $[n + 1, 2n]$ , while Bob colors his remaining uncolored vertices with distinct odd numbers from  $[n + 1, 2n]$ .


If the players received disjoint sets, the resulting graph  $G_{X,Y}$  is well-colored. If the sets' intersection is of size at least  $\epsilon n$ , however, the coloring is  $\epsilon$ -node. Therefore, the players can solve their set disjointness instance by simulating a CONGEST algorithm for detection of  $\epsilon$ -node colored graphs. Since there is only a single edge between Alice's and Bob's parts of the graph, the number of rounds  $T$  of a CONGEST algorithm detecting an  $\epsilon$ -node coloring with probability  $\geq 1/3$  satisfies:

$$T \geq \frac{R_{1/3}^{cc}(\mathbf{DISJ}_{n/2, \epsilon n}^n)}{\log(n)} \in \tilde{\Omega}\left(\frac{1}{\epsilon}\right)$$

□

**Theorem 6.** *Testing whether a distance-4 coloring is  $\epsilon$ -middle requires  $\tilde{\Omega}\left(\frac{1}{\epsilon}\right)$  rounds in the CONGEST model.*

For this lower bound, we consider graphs of the form presented in Figure 7.



**Fig. 7.** The graph we use for our lower bound on testing for  $\epsilon$ -middle distance-4 colorings (Theorem 6).

*Proof.* Let  $m$  be an integer and  $\epsilon \in [\frac{1}{m}, \frac{1}{2}]$ . Set  $N = s = \frac{1-\epsilon}{2}m$ ,  $t = 1$ , and consider an instance of  $\mathbf{DISJ}_{s,t}^N$ : a pair of sets  $(X, Y)$ ,  $X, Y \subseteq [m]$ . Consider the five layer graph  $G_{x,y} = (V, E)$  of Figure 7. Layers 2 and 4 each contain a single vertex ( $v_2$  and  $v_4$ ). Layer 3 contains  $\epsilon m$  vertices, all connected to  $v_2$  and  $v_4$ . Layer 1 contains  $s = \frac{1-\epsilon}{2}m$  vertices, all of them connected to  $v_2$ . Layer 5 is of the same size and connected to  $v_4$ .

Let Alice possess the vertices of layer 1 and 2, while Bob possesses layers 3, 4 and 5. As in our previous lower bounds, they use their inputs of the set disjointness problem to color vertices in layers 1 and 5. Vertices left uncolored after mapping the players' inputs on the graph are colored so as to never generate a conflict. The resulting colored graph is  $\epsilon$ -middle if and only if  $X$  and  $Y$  intersect, and thus the players can solve their disjointness instance by simulating a CONGEST protocol for the  $\epsilon$ -middle measure. Since there are  $\epsilon m$  edges between Alice's and Bob's parts of the graph, the number of rounds  $T$  of a CONGEST algorithm detecting an  $\epsilon$ -middle coloring with probability  $\geq 1/3$  satisfies:

$$T \geq \frac{R_{1/3}^{cc} \left( \mathbf{DISJ}_{\frac{1-\epsilon}{2}m, 1}^N \right)}{\epsilon m \cdot \log(m)} \in \tilde{\Omega} \left( \frac{1}{\epsilon} \right)$$

□