Improved Hardness of Approximation of Diameter in the CONGEST Model

Ofer Grossman
MIT, Cambridge, MA, USA
ofer.grossman@gmail.com

Seri Khoury
University of California, Berkeley, CA, USA
seri_khoury@berkeley.edu

Ami Paz
Faculty of Computer Science, Universität Wien, Austria
ami.paz@univie.ac.at

Abstract

We study the problem of approximating the diameter D of an unweighted and undirected n-node graph in the CONGEST model. Through a connection to extremal combinatorics, we show that a $(6/11+\epsilon)$-approximation requires $\Omega(n^{1/6}/\log n)$ rounds, a $(4/7+\epsilon)$-approximation requires $\Omega(n^{1/4}/\log n)$ rounds, and a $(3/5+\epsilon)$-approximation requires $\Omega(n^{1/3}/\log n)$ rounds. These lower bounds are robust in the sense that they hold even against algorithms that are allowed to return an additional small additive error. Prior to our work, only lower bounds for $(2/3+\epsilon)$-approximation were known [Frischknecht et al. SODA 2012, Abboud et al. DISC 2016].

Furthermore, we prove that distinguishing graphs of diameter 3 from graphs of diameter 5 requires $\Omega(n/\log n)$ rounds. This stands in sharp contrast to previous work: while there is an algorithm that returns an estimate $\lfloor 2/3D \rfloor \le \tilde{D} \le D$ in $\tilde{O}(\sqrt{n}+D)$ rounds [Holzer et al. DISC 2014], our lower bound implies that any algorithm for returning an estimate $2/3D \le \tilde{D} \le D$ requires $\tilde{\Omega}(n)$ rounds.

2012 ACM Subject Classification Mathematics of computing → Approximation algorithms; Theory of computation → Distributed computing models

Keywords and phrases Distributed graph algorithms, Approximation algorithms, Lower bounds

Digital Object Identifier 10.4230/LIPIcs.DISC.2020.19

Funding *Ofer Grossman*: Supported by a Fannie and John Hertz Foundation Fellowship.
Ami Paz: Supported by the Austrian Science Fund (FWF): P 33775-N, Fast Algorithms for a Reactive Network Layer.

Acknowledgements We thank Noga Alon for useful pointers regarding generalized polygons. This work was done while the fist two authors were visiting the Simon's Institute for the Theory of Computing.

1 Introduction and Related Work

The diameter D of a graph is one of the most fundamental parameters in graph theory. In distributed computing, the diameter is of utmost importance, as it captures the minimal number of rounds needed for a message to traverse all the nodes in the network. The complexity of computing the exact or approximate value of the diameter has been extensively studied in the distributed setting [1, 6, 7, 14, 16–19, 21, 23].

In the standard CONGEST model, the complexity of computing the exact diameter is $\Theta(n/\log n + D)$ rounds [14, 19]. On the other hand, there is a folklore algorithm yielding a 1/2-approximation for the diameter in $O(D)$ rounds: running a BFS (from an arbitrary node) and returning its depth.

34th International Symposium on Distributed Computing (DISC 2020).
Editor: Hagit Attiya; Article No. 19; pp. 19:1–19:16
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Table 1 A summary of the state of the art results for diameter approximation.

Approx.	Bound	Ref. and Comments
Exact	$\tilde{\Theta}(n)$	[14, 19]
2 vs. 3	$\tilde{\Omega}(n)$	[18]
$\lfloor 2/3D \rfloor < \tilde{D} \leq D$	$O(n^{1/2} + D)$	[16]
$2/3 + \epsilon$	$\tilde{\Omega}(n)$	[1]
3 vs. 5	$\tilde{\Omega}(n)$	**This paper** (Theorem 4)
$3/5 + \epsilon$	$\tilde{\Omega}(n^{1/3})$	**This paper** (Theorem 3)
$4/7$	$O(n^{1/3} + D)$	[3]
$4/7 + \epsilon$	$\tilde{\Omega}(n^{1/4})$	**This paper** (Theorem 2)
$6/11 + \epsilon$	$\tilde{\Omega}(n^{1/6})$	**This paper** (Theorem 1)
$1/2$	$O(D)$	Folklore

This raises the following natural question: for values of $1/2 < \alpha < 1$, how hard is it to α-approximate the diameter? Conversely, one may wonder how good of an approximation α is achievable in sub-linear time, and even in sub-polynomial time, as stated next.

- For which values of α does there exist a sub-polynomial time α-approximation algorithm for the diameter?
- For which values of α does there exist a truly sub-linear time α-approximation algorithm for the diameter?

We make progress on both these questions. For the first, we show that α must be at most $6/11$. For the second, we show that α must be at most $3/5$. The previous best known upper bound on α, for both cases, was $2/3$ [1]. All the results that are presented in this work, as well as the ones we compare with, are for unweighted and undirected graphs.

Our proofs use the well-established technique of reductions from communication complexity to distributed computing. Our main technical novelty is an interesting connection between extremal combinatorics, and specifically the existence of *generalized polygons* [15], and diameter approximation in the distributed setting. This extends prior work connecting extremal combinatorics and distributed computing [2, 8, 10, 12].

1.1 Our Contribution

Polynomial lower bound for .546-approximation. Our main result is that no sub-polynomial time algorithm can get better than a $6/11$-approximation.

▶ **Theorem 1.** *For any constant $0 < \epsilon < 5/11$, any algorithm for finding a $(6/11 + \epsilon)$-approximation for the diameter in the* CONGEST *model requires $\Omega(n^{1/6}/\log n)$ rounds.*

We prove analogous theorems for $(4/7 + \epsilon)$-approximation and $(3/5 + \epsilon)$-approximation, with lower bounds of $\Omega(n^{1/4}/\log n)$ and $\Omega(n^{1/3}/\log n)$, respectively.

▶ **Theorem 2.** *For any constant $0 < \epsilon < 3/7$, any algorithm for finding a $(4/7 + \epsilon)$-approximation for the diameter in the* CONGEST *model requires $\Omega(n^{1/4}/\log n)$ rounds.*

▶ **Theorem 3.** *For any constant $0 < \epsilon < 2/5$, any algorithm for finding a $(3/5 + \epsilon)$-approximation for the diameter in the* CONGEST *model requires $\Omega(n^{1/3}/\log n)$ rounds.*

These resutls hold even against constant diameter graphs and even against randomized algorithms that succeed with probability at least 2/3. Prior to our work, besides the near-linear lower bound of [14] for exact diameter, only a lower bound for $(2/3 + \epsilon)$-approximation was known: Abboud et al. [1] showed an $\Omega(n/\log^3 n)$ lower bound for this approximation factor. We note that Theorems 1, 2, and 3, as well as the aforementioned lower bound, also apply for algorithms that allow a constant *additive* error, in addition to the multiplicative one, as we explain in Section 1.2.

Near-linear lower bound for distinguishing diameter 3 vs 5. Next, we prove that distinguishing graphs of diameter 3 from graphs of diameter 5 requires a near-linear number of rounds.

▶ **Theorem 4.** *Any algorithm for distinguishing graphs of diameter 3 from graphs of diameter 5 in the* CONGEST *model requires* $\Omega(n/\log n)$ *rounds.*

We find this result rather surprising. There exist an algorithm [16] running in $O(\sqrt{n \log n} + D)$ rounds and returning an estimate $\lfloor \frac{2D}{3} \rfloor \le \tilde{D} \le D$. While the rounding in this equation might seem like an artifact of the proof, Theorem 4 shows that it is actually necessary.

That is, an algorithm for finding an estimate $\frac{2}{3}D \le \tilde{D} \le D$ can be used to distinguish diameter 3 from diameter 5, and we show that such a distinction must require $\Omega(n/\log n)$ rounds – much more than the $O(\sqrt{n \log n} + D)$ running time of the algorithm.

1.2 Robust Approximation

When dealing with diameter approximation, an important distinction to make is between *robust* and *non-robust* lower bounds. For example, as discussed above, an algorithm that finds an approximation $\tilde{D}$ of the diameter satisfying $\lfloor \frac{2D}{3} \rfloor \le \tilde{D} \le D$ does not in general imply a $\frac{2}{3}$-approximation.

However, as the diameter gets larger, the approximation ratio does approach 2/3. One way to view this is by saying that our 3 vs 5 lower bound is not a "robust" lower bound for $(3/5 + \epsilon)$-approximation. To show a "robust" lower bound for $(3/5 + \epsilon)$-approximation, we need a stronger result, i.e., that for any constant β, it is hard to distinguish between graphs of diameter $(3/5)D - \beta$ and graphs of diameter D. This would show that finding a $(3/5 + \epsilon)$-approximation of the diameter is hard not only in some low-diameter graphs, but also more generally. We formally define the notions of α-approximation for diameter and robust diameter lower bound.

▶ **Definition 5** (α-approximation for diameter)**.** *We say that an estimate $\tilde{D}$ is an α-approximation for diameter if*

$$\alpha D \le \tilde{D} \le D.$$

▶ **Definition 6** (Robust diameter lower bound)**.** *We say that α-approximating the diameter is robustly $T(n)$-hard if for any constant β, there is no algorithm which returns a value $\tilde{D}$ satisfying*

$$\alpha D - \beta \le \tilde{D} \le D$$

in $o(T(n))$ rounds in the CONGEST *model.*

In this paper, we prove both robust and non-robust lower bounds. The lower bounds presented in Theorems 1, 2, and 3 are robust, while the lower bound that is presented in Theorem 4 is not robust.

Notably, for a $(3/5+\epsilon)$-approximation, we give a robust lower bound of $\Omega(n^{1/3}/\log n)$, and a non-robust lower bound of $\Omega(n/\log n)$. The work of [16] rules out a robust lower bound better than $\Omega(\sqrt{n \log n})$, even for 2/3-approximation. This shows that there is an inherent, and large, gap between the robust and non-robust lower bounds.

This distinction between robust and non-robust approximation has been noted before, though not using this terminology. Holzer and Wattenhofer [18] showed that distinguishing diameter 2 from diameter 3 requires $\Omega(n/\log n)$ rounds, a result that can be viewed as a non-robust $(2/3+\epsilon)$-approximation lower bound. A robust lower bound of $\Omega(n/\log^3 n)$ for the same approximation ratio was later proven by Abboud et al. [1].

1.3 Further Related Work

The lower bound of Abboud et al. [1] for $(2/3+\epsilon)$-approximation follows from a lower bound for distinguishing between diameter $4\ell+2$ and $6\ell+1$, for some constant $\ell > 1$. Bringmann and Forster improved this result by showing the same hardness for distinguishing diameter $2\ell+1$ and $3\ell+1$ [5].

Very recently, in a concurrent and independent work [3], the authors show an upper bound of $O(n^{1/3}+D)$ for computing a 4/7-approximation for diameter.

All the results that are presented in this work are for unweighted graphs. For weighted graphs, Holzer and Pinsker [17] showed that $(1/2+\epsilon)$-approximation requires $\Omega(n/\log n)$ rounds. For (1/2)-approximation in the weighted case, one can compute single source shortest paths. The state of the art algorithm for single source shortest paths in the CONGEST model is by Forster and Nanongkai [13], who showed two algorithms for the problem. The first running in $\tilde{O}(\sqrt{nD})$ rounds and the second running in $\tilde{O}(\sqrt{n}D^{1/4}+n^{3/5}+D)$ rounds.

Road-map. In Section 2 we start with some basic definitions. The technical heart of the paper is in Sections 3 and 4. Theorem 4 is proved in Section 3, and Theorems 1, 2, and 3 are proved in Section 4.

2 Preliminaries

2.1 The Model

In the CONGEST model [22], a synchronized communication network of n computationally unbounded nodes is modeled by its communication graph $G=(V,E)$. Each of the nodes has a unique $O(\log n)$-bit identifier. The computation is split into rounds, and in each round each node can send a (possibly different) $O(\log n)$-bit message to each of its neighbors. The goal of the nodes is to compute some function of the network (e.g., its diameter, the value of the minimum vertex cover, etc.) while minimizing the number of communication rounds.

For a graph H that is not the input graph, we denote its set of nodes and edges by V_H and E_H, respectively. The distance between two nodes u,v in a graph G is denoted by $\mathrm{d}_G(u,v)$, and is the minimum number of hops in a path between them in G. The diameter D of the graph is the maximum distance between two nodes in it. The girth of the graph g is the minimum length of a cycle in it.

2.2 Communication Complexity

In the two-party communication setting [20, 26], two players, Alice and Bob, are given two input strings, $x,y \in \{0,1\}^K$, respectively, and need to jointly compute a function $f : \{0,1\}^K \times \{0,1\}^K \to \{\text{TRUE}, \text{FALSE}\}$ of their inputs, using a predefined communication protocol. The *communication complexity* of a function f is defined as follows. Definition 7 is a special case of [11, Definition 1].

▶ **Definition 7** (Communication Complexity). *Let $K \geq 1$ be an integer, f be a Boolean function $f : \{0,1\}^K \times \{0,1\}^K \to \{\text{TRUE}, \text{FALSE}\}$, and $\mathcal{Q}$ be the family of protocols that compute f correctly with probability at least $2/3$. Given 2 inputs $x, y \in \{0,1\}^K$, denote by $\Pi_Q(x,y)$ the transcript of a protocol Q on the inputs x, y, i.e., the sequence of bits that are exchanged between Alice and Bob. The cost of a protocol Q is $\mathrm{Cost}(Q) = \max_{x,y \in \{0,1\}^k} |\Pi_Q(x,y)|$.*

The communication complexity of f, denoted by $CC_f(K)$, is defined to be the minimum cost over all the possible protocols that compute f correctly with probability at least $2/3$: $CC_f(K) = \min_{Q \in \mathcal{Q}} \mathrm{Cost}(Q)$.

The set-disjointness function is defined as follows. For two strings $x, y \in \{0,1\}^K$, we say that x and y are not disjoint if and only if there is some index $i \in [K]$ such that $x_i = y_i = 1$. Otherwise we say that the strings are disjoint. It is well known that the communication complexity of set-disjointness is $\Omega(K)$ [24].

▶ Remark 8. Adding 0 bits to both input strings in matching locations does not change the output. Thus, we can assume a constant fraction of both input strings is 0 without affecting the asymptotic communication complexity. We use this fact in Section 4.

2.3 Lower Bound Graphs

Our lower bounds use the standard notion of family of lower bound graphs (see, e.g., [9]).

▶ **Definition 9** (Family of Lower Bound Graphs). *Let $K > 1$ be an integer, $f : \{0,1\}^K \times \{0,1\}^K \to \{\text{TRUE}, \text{FALSE}\}$ be a boolean function, and P be a graph predicate. A family of graphs $\left\{G_{(x,y)} = (V_{(x,y)}, E_{(x,y)}) \mid x, y \in \{0,1\}^K\right\}$ where each $G_{(x,y)}$ has a partition of the set of nodes $V_{(x,y)} = V_A \dot\cup V_B$ is said to be a family of* lower bound graphs for the CONGEST model w.r.t. f and P *if the following properties hold:*

1. *Only the existence of nodes in V_A or edges in $V_A \times V_A$ may depend on x;*
2. *Only the existence of nodes in V_B or edges in $V_B \times V_B$ may depend on y;*
3. *$G_{(x,y)}$ satisfies the predicate P iff $f(x,y) = \text{TRUE}$.*

For such a family, we denote by $C = E(V_A, V_B)$ the *cut*, i.e., the set of edges between V_A and V_B.

We use the following theorem, which is standard in the context of reductions to communication complexity (see, for example [1, 9, 10, 14, 17]). Its proof is by a standard simulation argument and appears in [9].

▶ **Theorem 10.** *Fix a function $f : \{0,1\}^K \times \{0,1\}^K \to \{\text{TRUE}, \text{FALSE}\}$ and a predicate P. If there is a family of lower bound graphs for the CONGEST model w.r.t. f and P then any algorithm for deciding P in the CONGEST model requires $\Omega\left(\frac{\mathrm{CC}_f(K)}{|C| \log n}\right)$ rounds.*

2.4 Generalized Polygons

Our proofs in Section 4 use the existence of generalized polygons [15]. A generalized polygon is an incidence relation whose incidence graph has several nice properties. In our context, we use the following key property of a a generalized polygon's incidence graph: its girth is twice its diameter. For the sake of simplifying the presentation, we also use the fact that the incidence graphs are balanced.

We use the notation $H = (L, R, E_H)$ to denote a bipartite graph H, where the bi-partition of the vertex set of H is L and R, and the set of edges of H is E_H. When $|L| = |R| = p$, we say that H is a *balanced* bipartite graph of size $2p$.

▶ **Definition 11.** *For two integers $p \geq t \geq 3$, we denote by* $\mathrm{Ex}(p,t)$ *the maximum number of edges in a balanced bipartite graph of size $2p$, diameter t, and girth $2t$.*

For $t \in \{3,4,6\}$, there are generalized polygons whose incidence graph has $2p$ nodes, diameter t, girth $2t$, and $\Theta(p^{1+\frac{1}{t-1}})$ edges. The cases of $t=3$ and $t=4$ were shown by Singelton [25], and Benson [4] gave a simplified proof and extended the result for $t=6$. This is summarized in the following theorem, which will be used later without explicitly re-mentioning generalized polygons.

▶ **Theorem 12** ([4, 25]). *For $t \in \{3,4,6\}$, it holds that* $\mathrm{Ex}(p,t) = \Omega(p^{1+\frac{1}{t-1}})$.

3 Diameter 3 vs 5

In this section we prove the following theorem.

▶ **Theorem 4.** *Any algorithm for distinguishing graphs of diameter 3 from graphs of diameter 5 in the* CONGEST *model requires $\Omega(n/\log n)$ rounds.*

To prove Theorem 4, we show a family of lower bound graphs $\{G_{(x,y)} \mid x,y \in \{0,1\}^K\}$ with respect to the set-disjointness function and the graph predicate that distinguishes between graphs of diameter 3 and graphs of diameter 5. That is, the predicate is defined only on a graph G with either 3 or 5, and is TRUE if and only if G has diameter 5. We start with the fixed graph construction G and then we show how to get the graph $G_{(x,y)}$ given two strings $x,y \in \{0,1\}^K$.

The Fixed Graph Construction G. The fixed graph construction is defined as follows. There are 8 sets of nodes $S, C^1, A^1, B^1, C^2, A^2, B^2, T$, each of size $p = n/8$. Each of the sets S, A^1, B^1, A^2, B^2, T is an independent set, and C^1 and C^2 are cliques. The nodes in the sets are denoted $S = \{s_i \mid i \in [p]\}$, $T = \{t_i \mid i \in [p]\}$, and for $h \in \{1,2\}$, $A^h = \{a_i^h \mid i \in [p]\}$, $B^h = \{b_i^h \mid i \in [p]\}$, and $C^h = \{C_i^h \mid i \in [p]\}$.

The connections between the sets are defined as follows. Each pair of sets $H_1 \neq H_2 \in \{S, C^1, A^1, B^1\}$ is connected by a perfect matching, where we connect the i'th node in H_1 to the i'th node in H_2. For example, the sets S and C^1 are connected by the perfect matching $\{(s_i, c_i^1) \mid i \in [p]\}$. Similarly, each pair of sets in $\{T, C^2, A^2, B^2\}$ is connected by a perfect matching. This concludes the fixed graph construction G. Let $K = p^2$. We define the graph $G_{(x,y)}$, given two strings $x,y \in \{0,1\}^K$, as follows.

Obtaining $G_{(x,y)}$ from G and $x,y \in \{0,1\}^K$. For each of the strings x and y, we index the $K = p^2$ positions by $x_{(i,j)}$ and $y_{(i,j)}$ for $i,j \in [p]$. The set of nodes of $G_{(x,y)}$ is exactly as in G. The set of edges of $G_{(x,y)}$ contains all the edges in G, and the following edges between pairs of nodes in $A^1 \times A^2$ and between pairs of nodes in $B^1 \times B^2$.

$$\{(a_i^1, a_j^2) \mid x_{(i,j)} = 0\}; \quad \{(b_i^1, b_j^2) \mid y_{(i,j)} = 0\}.$$

That is, if $x_{(i,j)} = 0$, we add an edge between a_i^1 and a_j^2, and if $y_{(i,j)} = 0$, we add an edge between b_i^1 and b_j^2. This concludes the definition of $G_{(x,y)}$ (See also Figure 1, for an illustration). Next, we prove that $G_{(x,y)}$ has diameter 3 if the strings x and y are disjoint, and otherwise it has diameter at least 5. We prove this in Lemmas 13 and 14.

▶ **Lemma 13.** *If the strings x and y are disjoint, then the diameter of $G_{(x,y)}$ is 3.*

Proof. We show that for any two nodes u, v, $\mathrm{d}_{G_{(x,y)}}(u,v) \le 3$. Let $L = S \cup C^1 \cup A^1 \cup B^1$, and let $R = T \cup C^2 \cup A^2 \cup B^2$ The proof is by the following case analysis.

1. $u, v \in L$ or $u, v \in R$: We prove the claim for the case in which $u, v \in L$. The case for which $u, v \in R$ is similar. Observe that any node in L is connected by an edge to some node in C^1. Hence, since C^1 is a clique, this implies that $\mathrm{d}_{G_{(x,y)}}(u,v) \le 3$.
2. $u \in L$ and $v \in R$: Hence, u belongs to one of the sets in $\{S, C^1, A^1, B^1\}$ and v belongs to one of the sets in $\{T, C^2, A^2, B^2\}$. We assume that $u \in S$ and $v \in T$; the proof for the other cases is similar. Let i be such that $u = s_i$, and j such that $v = t_j$. Since the sets are disjoint, it holds that either $x_{(i,j)} = 0$, or $y_{(i,j)} = 0$ (or both). Hence, either there is an edge between a_i^1 and a_j^2, or there is an edge between b_i^1 and b_j^2 (or both), and assume the former without loss of generality. Since s_i is connected to a_i^1 and t_j is connected to a_j^2, we have $\mathrm{d}_{G_{(x,y)}}(s_i, t_j) \le 3$. Furthermore, one can verify that the distance between s_i and t_j cannot be smaller than 3, which implies that the diameter of the graph is 3. ◀

▶ **Lemma 14.** *If the strings x and y are not disjoint, then the diameter of $G_{(x,y)}$ is at least 5.*

Proof. As the sets are not disjoint, there are $i, j \in [p]$ for which it holds that $x_{(i,j)} = y_{(i,j)} = 1$. We show that in this case, any path P from s_i to t_j is of length at least 5, i.e., $\mathrm{d}_{G_{(s_i,t_j)}}(u,v) \ge 5$. Observe that any path P from s_i to t_j must either pass from a node in A^1 to a node in A^2, or from a node in B^1 to a node in B^2. We assume that former case; the latter is similar. The proof is by the following case analysis.

1. The path P visits a node $a_{j'}^2 \in A^2$ for which $j' \ne j$: Observe that $\mathrm{d}_{G_{(x,y)}}(s_i, a_{j'}^2) \ge 2$, and that $\mathrm{d}_{G_{(x,y)}}(a_{j'}^2, t_j) = 3$. Hence, $\mathrm{d}_{G_{(x,y)}}(s_i, t_j) \ge 5$.
2. The path P visits a_j^2. Since $x_{(i,j)} = 1$, there is no edge between a_i^1 and a_j^2. This implies that $\mathrm{d}_{G_{(x,y)}}(s_i, a_j^2) \ge 4$, and hence $\mathrm{d}_{G_{(x,y)}}(s_i, t_j) \ge 5$. ◀

Proof of Theorem 4. First, we define $V_A = S \cup T \cup A^1 \cup A^2 \cup C^1 \cup C^2$, and $V_B = B^1 \cup B^2$. Lemmas 13 and 14 imply that $\{G_{(x,y)} \mid x, y \in \{0,1\}^K\}$ is a family of lower bound graphs with respect to the set-disjointness problem and the graph predicate that distinguishes between graphs of diameter 3 and graphs of diameter 5. Observe that the cut size is $E(V_A, V_B) = \Theta(p)$, and $p = \Theta(n)$. Hence, since the length of the input strings is $K = p^2$, and since the communication complexity of set-disjointness is $\Omega(K) = \Omega(p^2)$, Theorem 10, implies that any algorithm for deciding whether a graph has diameter 3 or 5 in the CONGEST model requires $\Omega(p^2/p \log p) = \Omega(p/\log p) = \Omega(n/\log n)$ rounds. ◀

The connectivity of $G_{(x,y)}$. One may wonder about the connectivity of $G_{(x,y)}$. If the graph $G_{(x,y)}$ is not connected, then the construction wouldn't be meaningful as there is a trivial lower bound of $\Omega(D)$, where D is the diameter of the graph, which is ∞ in graphs that are not connected. Observe that the only case in which $G_{(x,y)}$ is not connected is when $x = y = 1^K$. To ensure connectivity (and in fact constant diameter, due to the cliques C^1 and C^2), we can assume that at least one of the strings x or y has a zero bit. Clearly, the communication complexity of set-disjointness doesn't change under this assumption. In fact, Remark 8 allows to make an even stronger assumption, which we only need in the next section.

4 Robust Lower Bounds

In this section we prove robust lower bounds for $(6/11 + \epsilon)$-approximation, $(4/7 + \epsilon)$-approximation, and $(3/5 + \epsilon)$-approximation of the diameter. Our lower bounds follow from the following theorem.

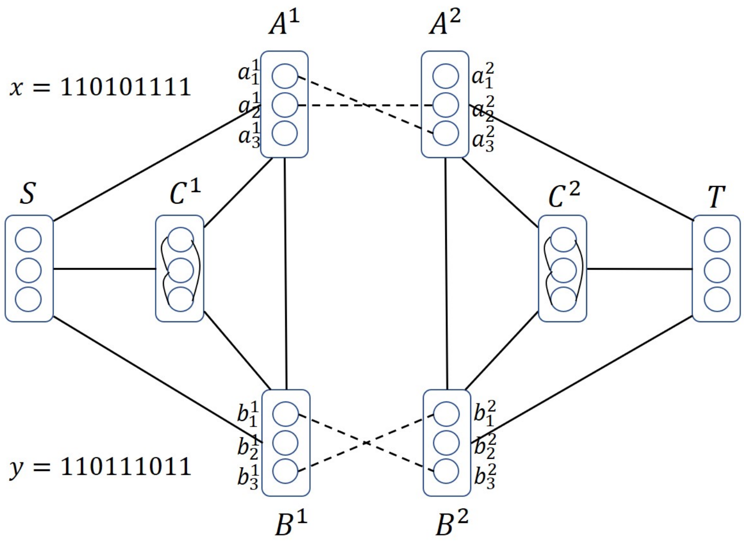

Figure 1 Diameter 3 vs 5: An example for the graph construction $G_{(x,y)}$ for $p = 3$: There are 8 sets of nodes $S, C^1, A^1, A^2, C^2, B^1, B^2, T$, each of size $p = 3$. Each of the sets S, A^1, A^2, B^1, B^2, T forms an independent set, and the sets C^1 and C^2 are cliques. In this diagram, an edge between two sets represents a perfect matching connecting them. For example, the edge between S and C^1 represents all the edges in $\{(s_i, c_i^1) \mid i \in [p]\}$. The dashed edges between A^1 and A^2 are the input edges which depend on the input sting x. Recall that we index the $p^2 = 9$ positions of x by pairs of indices $(i, j) \in [p] \times [p]$. In this example, we have that $x_{(1,3)} = x_{(2,2)} = 0$, and all the other bits of x are 1's. Hence, the only edges between A^1 and A^2 are (a_1^1, a_3^2) and (a_2^1, a_2^2). Similarly, the dashed edges between B^1 and B^2 represent the input edges which depend on the string y. Since in this example we have $y_{(1,3)} = y_{(3,1)} = 0$, and all the other bits of y are 1's, the only edges between B^1 and B^2 are (b_1^1, b_3^2) and (b_3^1, b_1^2).

▶ **Theorem 15.** *Let $t \in \{3, 4, 6\}$. For any constant $0 < \epsilon < 1 - \frac{t}{2t-1}$, any algorithm for computing a $(\frac{t}{2t-1} + \epsilon)$-approximation to the diameter in the* CONGEST *model requires $\Omega\left(n^{1/t}/\log n\right)$ rounds, where n is the number of nodes in the input graph.*

The theorem has the following consequences, when plugging in $t = 6$, $t = 4$ and $t = 3$, in this order.

▶ **Theorem 1.** *For any constant $0 < \epsilon < 5/11$, any algorithm for finding a $(6/11 + \epsilon)$-approximation for the diameter in the* CONGEST *model requires $\Omega(n^{1/6}/\log n)$ rounds.*

▶ **Theorem 2.** *For any constant $0 < \epsilon < 3/7$, any algorithm for finding a $(4/7 + \epsilon)$-approximation for the diameter in the* CONGEST *model requires $\Omega(n^{1/4}/\log n)$ rounds.*

▶ **Theorem 3.** *For any constant $0 < \epsilon < 2/5$, any algorithm for finding a $(3/5 + \epsilon)$-approximation for the diameter in the* CONGEST *model requires $\Omega(n^{1/3}/\log n)$ rounds.*

Recall that $\mathrm{Ex}(p, t)$ is the maximum number of edges of a balanced bipartite graph of size $2p$, diameter t, and girth $2t$ (see Definition 11). Let $t \in \{3, 4, 6\}$, $p \geq t$ and let $K = \mathrm{Ex}(p, t)$. To prove Theorem 15, we show a family of lower bound graphs $\left\{G_{(x,y)} \mid x, y \in \{0, 1\}^K\right\}$ with respect to the set-disjointness function and the graph predicate that distinguishes between graphs of diameter $t(b + 1) + 1$ and graphs of diameter $(2t - 1)b$, for some integer $b = \Theta(1/\epsilon)$ that will be chosen later.

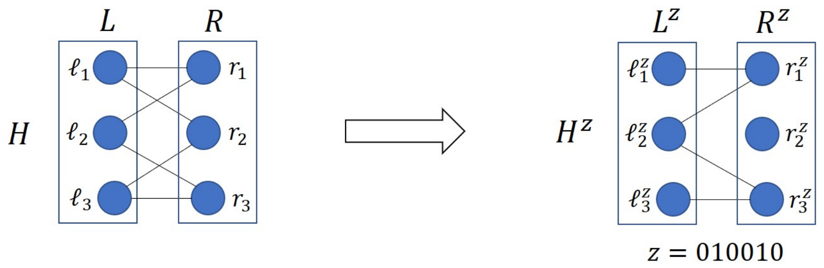

Figure 2 An example for H and H^z. In this example $t = p = 3$ and therefore H is a bipartite graph of diameter $t = 3$ and girth $2t = 6$. For these parameters, we have $K = \text{Ex}(3, 3) = 6$. Recall that $\pi : E_H \to [K]$ is an arbitrary $1 : 1$ mapping from the set of pairs $(\ell_i, r_j) \in E_H$ to $[K]$. In this example, we choose $\pi(\ell_1, r_1) = 1, \pi(\ell_1, r_2) = 2, \pi(\ell_2, r_1) = 3, \pi(\ell_2, r_3) = 4, \pi(\ell_3, r_2) = 5$ and $\pi(\ell_3, r_3) = 6$. Furthermore, in this example we have $z = 010010$. Hence, since H^z is obtained from H by keeping only the edges that correspond to the 0 bits in z, we have that the only edges in H^z are $(\ell_1^z, r_1^z), (\ell_2^z, r_1^z), (\ell_2^z, r_3^z)$ and (ℓ_3^z, r_3^z).

The rest of this section is organized as follows. In section 4.1, we start with the description of $G_{(x,y)}$ given two strings $x, y \in \{0,1\}^K$. In section 4.2, we show that $\left\{G_{(x,y)} \mid x, y \in \{0,1\}^K\right\}$ is a family of lower bound graphs with the required properties. In Section 4.3, we deduce Theorem 15. While the graphs $G_{(x,y)}$ need to be connected, we ignore this fact in Section 4.1; in Section 4.4 we show how to slightly modify the construction so that the graphs become connected (and even of constant diameter) for any x and y.

4.1 Description of $G_{(x,y)}$

Given two strings $x, y \in \{0,1\}^K$, we describe the graph $G_{(x,y)}$ in three steps. In the first step, given a string $z \in \{0,1\}^K$, we define a bipartite graph H^z. Roughly speaking, H^z is obtained from a densest possible balanced bipartite graph H of size $2p$, diameter t and girth $2t$, where we keep only some of the edges of H in H^z according to the string z. In the second step, we define a graph $\widetilde{H}^z$, which is obtained form H^z by stretching each edge to a path of length b. In the third step, we describe how to get $G_{(x,y)}$ from $\widetilde{H}^x$ and $\widetilde{H}^y$.

Description of H^z. Let $H = (L, R, E_H)$ be a balanced bipartite graph of size $2p$, diameter t, girth $2t$, and a maximum number of edges. That is, the number of edges of H is $|E_H| = \text{Ex}(p, t) = K$. We denote the nodes of H by $L = \{\ell_1, \cdots, \ell_p\}$ and $R = \{r_1, \cdots, r_p\}$. Furthermore, let $\pi : E_H \to [K]$ be an enumeration of E_H, that is, π is an arbitrary ordering over the set of pairs $E_H \subseteq L \times R$. By this mapping, each bit of a string $z \in \{0,1\}^K$ corresponds to a unique edge in E_H.

Given a string $z \in \{0,1\}^K$, the graph H^z is defined as follows. H^z is a version of H where we keep only the edges for which the corresponding bits in z are 0. More formally, $H^z = (L^z, R^z, E_{H^z})$ is a balanced bipartite graph with $|L^z| = |R^z| = p$, where $L^z = \{\ell_1^z, \cdots, \ell_p^z\}$ and $R^z = \{r_1^z, \cdots, r_p^z\}$. A pair of nodes $(\ell_i^z, r_j^z) \in L^z \times R^z$ is connected by an edge in H^z if (ℓ_i, r_i) is an edge of H and $z_{\pi(\ell_i, r_i)} = 0$, that is, $E_{H^z} = \{(\ell_i^z, r_j^z) \mid (\ell_i, r_j) \in E_H \wedge z_{\pi(\ell_i, r_j)} = 0\}$. See Figure 2 for an illustration of obtaining H^z from H and an input string $z \in \{0,1\}^K$.

Description of $\widetilde{H}^z$. $\widetilde{H}^z$ is obtained from H^z by replacing each edge $(\ell_i^z, r_j^z) \in E_{H^z}$ with a path of $b + 1$ nodes and b edges, starting at ℓ_i^z and ending at r_j^z, where b is some positive integer to be chosen later. We denote this path by $P^z_{(\ell_i^z, r_j^z)}$. We slightly abuse notation and denote the set of nodes on this path also by $P^z_{(\ell^z_i, r^z_j)}$. We sometimes treat $P^z_{(\ell^z_i, r^z_j)}$ as a set of nodes, and sometimes as a path, but this will be clear from the context. Hence, the set of nodes of $\widetilde{H}^z$ is

$$V_{\widetilde{H}^z} = L^z \cup R^z \cup \bigcup_{(\ell^z_i, r^z_j) \in E_{H^z}} P^z_{(\ell^z_i, r^z_j)}$$

and the edges of $\widetilde{H}^z$ are only the ones on the paths in $\{P^z_{(\ell^z_i, r^z_j)} \mid (\ell^z_i, r^z_j) \in E_{H^z}\}$. Observe that $\widetilde{H}^z$ is not necessarily bipartite.

Obtaining $G_{(x,y)}$ from $\widetilde{H}^x$ and $\widetilde{H}^y$. Given two input strings $x, y \in \{0,1\}^K$, $G_{(x,y)}$ is composed of $\widetilde{H}^x$ and $\widetilde{H}^y$ where we add a perfect matching between L^x and L^y, $\{(\ell^x_i, \ell^y_i) \mid i \in [p]\}$, and a perfect matching between R^x and R^y, $\{(r^x_i, r^y_i) \mid i \in [p]\}$. This concludes our construction. See also figures 3 and 4 for illustrations of $G_{(x,y)}$. In these figures, we also illustrate $\widetilde{H}^z$ for $z = x \wedge y$, where the string $x \wedge y \in \{0,1\}^K$ is defined by $(x \wedge y)_h = x_h \cdot y_h$ for any $h \in [K]$. That is, $(x \wedge y)_h = 1$ if and only if $x_h = y_h = 1$. The reason that we illustrate $\widetilde{H}^z$ in the same figures is that our proof heavily relies on comparing distances in $G_{(x,y)}$ to distances in $\widetilde{H}^z$ for $z = x \wedge y$. Figure 3 is an illustration of the two graphs when the strings are not disjoint, while Figure 4 is an illustration of the two graphs when the strings are disjoint. Before we prove that $\left\{G_{(x,y)} \mid x, y \in \{0,1\}^K\right\}$ is a family of lower bound graphs, we show the following two useful properties of the balanced bipartite graph $H = (L, R, E_H)$ that was described above.

▶ **Property 1.** *If t is odd, then the distance between any two nodes $u, v \in L$ in H is at most $t-1$. Similarly, the distance between any two nodes $u, v \in R$ in H is at most $t-1$.*

Proof. The distance between every two nodes in H is at most its diameter t, but the distance between every two nodes in the same side of the bi-partition is even, so it is at most $t-1$. ◀

▶ **Property 2.** *If t is even, then the distance between any pair of nodes $u \in L$ and $v \in R$ in H is at most $t-1$.*

Proof. The distance between every two nodes in H is at most its diameter t, and the distance between two nodes in different sides of the bi-partition is odd, so it is at most $t-1$. ◀

4.2 $G_{(x,y)}$ is a family of lower bound graphs

Our goal in this section is to prove that $\{G_{(x,y)} \mid x, y \in \{0,1\}^K\}$ is a family of lower bound graphs with respect to the set-disjointness function and the graph predicate that distinguishes graphs of diameter $t(b+1)+1$ from graphs of diameter $(2t-1)b$. For the rest of the paper, $z = x \wedge y$. Our proof relies on comparing distances between nodes in $G_{(x,y)}$ to distances between nodes in $\widetilde{H}^z$. While the proof contains many technical details that require some care, it follows from the following simple intuition.

Intuition and overview of the proof. First, it is not very hard to see that the diameter of $G_{(x,y)}$ is roughly equal to the diameter of $\widetilde{H}^z$ (up to an additive $t+1$). Hence, it suffices to argue that if the strings x and y are disjoint, then the diameter of $\widetilde{H}^z$ is at most tb, and otherwise the diameter of $\widetilde{H}^z$ is at least $(2t-1)b$. The main idea is to note that H^z is isomorphic to H if and only if the strings x and y are disjoint. Hence, if the strings are disjoint, the diameter of H^z is equal to the diameter of H which is t, and since $\widetilde{H}^z$ is obtained from H^z by stretching each edge to a path of length b, the diameter of $\widetilde{H}^z$ is tb.

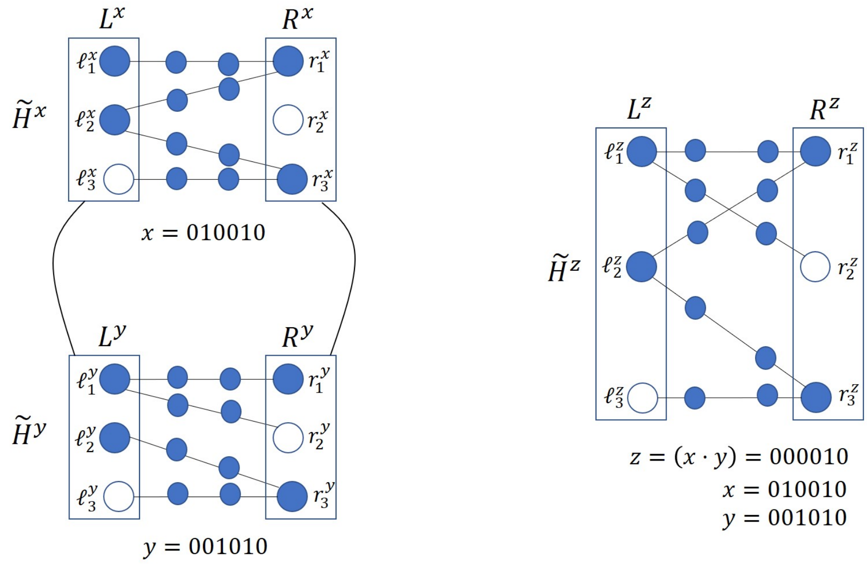

Figure 3 An Illustration of $G_{(x,y)}$ (on the left) and $\widetilde{H}^z$ for $z = x \wedge y$ (on the right). $G_{(x,y)}$ is composed of $\widetilde{H}^x$ and $\widetilde{H}^y$, where we add a perfect matching between L^x and L^y, $\{(\ell_i^x, \ell_i^y) \mid i \in [p]\}$, and a perfect matching between R^x and R^y, $\{(r_i^x, r_i^y) \mid i \in [p]\}$. The edge between L^x and L^y in this figure represents the matching between them. Similarly, the edge between R^x and R^y in this figure represents the matching between them. The parameters in this example are exactly as the ones chosen for Figure 2. That is, $t = p = 3$, $b = 3$, $\pi(\ell_1, r_1) = 1, \pi(\ell_1, r_2) = 2, \pi(\ell_2, r_1) = 3, \pi(\ell_2, r_3) = 4, \pi(\ell_3, r_2) = 5$ and $\pi(\ell_3, r_3) = 6$. Furthermore, we have $x = 010010, y = 001010$ and $z = x \wedge y = 000010$. Hence, the only paths of length b that we have in $\tilde{H}^x$ are $P^x_{(\ell_1^x, r_1^x)}, P^x_{(\ell_2^x, r_1^x)}, P^x_{(\ell_2^x, r_3^x)}$ and $P^x_{(\ell_3^x, r_3^x)}$. And the only paths of length b that we have in $\tilde{H}^y$ are $P^y_{(\ell_1^y, r_1^y)}, P^y_{(\ell_1^y, r_2^y)}, P^y_{(\ell_2^y, r_3^y)}$ and $P^y_{(\ell_3^y, r_3^y)}$. Observe that the path $P^x_{(\ell_3^x, r_2^x)}$ doesn't exist in $\widetilde{H}^x$, and that the path $P^y_{(\ell_3^y, r_2^y)}$ doesn't exist in $\widetilde{H}^y$. This is because $x_{\pi(\ell_3, r_2)} = x_5 = 1$ and $y_{\pi(\ell_3, r_2)} = y_5 = 1$. This implies that $z_{\pi(\ell_3, r_2)} = z_5 = 1$ as well, and therefore the path $P^z_{(\ell_3^z, r_2^z)}$ doesn't exist in $\widetilde{H}^z$. It is easy to see that the distance between ℓ_3^z and r_2^z in $\widetilde{H}^z$ is $5b = (2t-1)b$. One can verify that the distance between ℓ_3^x and r_2^x in $G_{(x,y)}$ is at least $5b = (2t-1)b$ as well. This illustrates that when the strings x and y are disjoint, the diameter of $G_{(x,y)}$ is at least $(2t-1)b$.

On the other hand, if H^z is not isomorphic to H, then there is an edge in H for which the corresponding edge in H^z doesn't exist. Since the girth of H is $2t$, it implies that there are two nodes in H^z at distance at least $(2t-1)$ from each other. Hence, the distance between the corresponding two nodes in $\widetilde{H}^z$ is at least $(2t-1)b$. Next, we formalize these ideas and give a more detailed proof. The non-disjointness case is proved in Lemma 19, which uses claims 16, 17 and 18. The disjointness case is proved in Lemma 23, which uses claims 20, 21, and 22. Recall that given a graph G and two nodes u, v in it, we denote by $\mathrm{d}_G(u,v)$ the distance between u and v in G.

Non-disjointness case

▷ Claim 16. If x and y are not disjoint, then the diameter of H^z is at least $2t-1$ and the diameter of $\widetilde{H}^z$ is at least $(2t-1)b$. In particular, there are $\ell_i^z \in L^z$ and $r_j^z \in R^z$ such that $\mathrm{d}_{\widetilde{H}^z}(\ell_i^z, r_j^z) \geq (2t-1)b$.

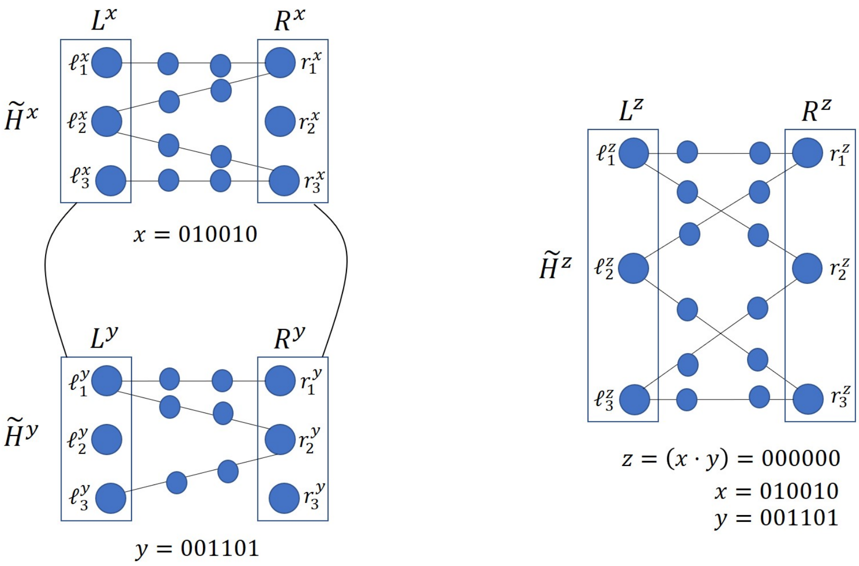

Figure 4 An illustration of $G_{(x,y)}$ and $\widetilde{H}^z$ for $z = x \wedge y$. The parameters in this example are exactly as the ones chosen for Figures 2 and 3. The only difference in this example compared to the one in Figure 3 is that the strings x and y are disjoint. Hence, $z = x \wedge y$ is an all zeros string. It is easy to see that in this case the diameter of $\widetilde{H}^z$ is tb. The key point of the proof is that the diameter of $G_{(x,y)}$ is not very much larger than the diameter of $\widetilde{H}^z$ (in fact, it is larger by at most $t+1$, which is negligible compared to tb for values of $b \gg t$). To illustrate this in this example, we show a path of length $tb + t = 3b + 3$ from ℓ_1^x and r_3^y. We start by moving from ℓ_1^x to ℓ_1^y using the edge (ℓ_1^x, ℓ_1^y) that is part of the matching between L^x and L^y. Then, we use the path of length b from ℓ_1^y to r_2^y, and the path of length b from r_2^y to ℓ_3^y. After that, we use the edge (ℓ_3^y, ℓ_3^x) to move to ℓ_3^x, and the path of length b from ℓ_3^x to r_3^x. Finally, we use the edge (r_3^x, r_3^y) to reach r_3^y. This example illustrates that the diameter of $G_{(x,y)}$ is relatively small if the strings are disjoint.

Proof. Observe that if the strings x and y are not disjoint, then there is an $h \in [K]$ for which it holds that $x_h = y_h = 1$. Hence, $z_h = 1$. Since H^z is obtained from H by keeping only the edges that correspond to the 0 bits in z, it follows that there is an edge $(\ell_i, r_j) \in H$ such that there is no edge between the corresponding pair (ℓ_i^z, r_j^z) in H^z. Hence, since H has girth $2t$, if follows that the distance between ℓ_i^z and r_j^z in H^z is at least $2t - 1$. Since $\widetilde{H}^z$ is obtained from H^z by replacing each edge with a path of length b, it follows that the diameter of $\widetilde{H}^z$ is at least $(2t-1)b$. ◁

▷ Claim 17. For any $(\ell_i, r_j) \in E_H$, if one of the paths $P^x_{(\ell_i^x, r_j^x)}$ and $P^y_{(\ell_i^y, r_j^y)}$ exists in $G_{(x,y)}$, then the path $P^z_{(\ell_i^z, r_j^z)}$ exists in $\widetilde{H}^z$. Similarly, if $P^z_{(\ell_i^z, r_j^z)}$ exists in $\widetilde{H}^z$, then either $P^x_{(\ell_i^x, r_j^x)}$ exists in $G_{(x,y)}$ or $P^y_{(\ell_i^y, r_j^y)}$ exists in $G_{(x,y)}$.

Proof. Let $h = \pi(\ell_i, r_j)$. Observe that if one of the paths $P^x_{(\ell_i^x, r_j^x)}$ and $P^y_{(\ell_i^y, r_j^y)}$ exists in $G_{(x,y)}$, then it must be the case that either $x_h = 0$ or $y_h = 0$. Hence, $z_h = 0$. Therefore there is an edge between ℓ_i^z and r_j^z in H^z, which is stretched to a path $P^z_{(\ell_i^z, r_j^z)}$ in $\widetilde{H}^z$. The other direction of the claim is proved similarly. ◁

▷ Claim 18. For any $\ell_i^x \in L^x$ and $r_j^x \in R^x$, it holds that $\mathrm{d}_{G_{(x,y)}}(\ell_i^x, r_j^x) \geq \mathrm{d}_{\widetilde{H}^z}(\ell_i^z, r_j^z)$.

Proof. Consider a shortest path between ℓ_i^x and r_j^x in $G_{(x,y)}$. Observe that this path is composed of edges crossing from $\widetilde{H}^x$ to $\widetilde{H}^y$ or vice versa (i.e., edges in $(L^x \times L^y) \cup (R^x \times R^y)$), and of paths of length b crossing from $L^x \cup L^y$ to $R^x \cup R^y$ or vice versa. Let q be the number of paths of length b crossing from $L^x \cup L^y$ to $R^x \cup R^y$ (or vice versa) that are used by the shortest path. And denote these paths by $P^1, P^2, \cdots, P^q$. Clearly, $\mathrm{d}_{G_{(x,y)}}(\ell_i^x, r_j^x) \geq qb$. Hence, it suffices to show that $qb \geq \mathrm{d}_{\widetilde{H}^z}(\ell_i^z, r_j^z)$.

For this, observe that for any $h \in [q]$, there are $ih, jh \in [p]$ and $w \in \{x, y\}$, for which $P^h = P^w_{(\ell^w_{ih}, r^w_{jh})}$ (That is, P^h is connecting either a pair $(\ell^x_{ih}, r^x_{jh}) \in L^x \times R^x$ or a pair $(\ell^w_{ih}, r^w_{jh}) \in L^y \times R^y$). Hence, by Claim 17, this implies that for any $h \in [q]$, the path $P^z_{(\ell^z_{ih}, r^z_{jh})}$ exists in $\widetilde{H}^z$. Therefore, by starting at ℓ_i^z and following these q paths of length b in $\widetilde{H}^z$ we reach r_j^z. Hence, $qb \geq \mathrm{d}_{\widetilde{H}^z}(\ell_i^z, r_j^z)$. ◁

▶ **Lemma 19.** *If x and y are not disjoint, then the diameter of $G_{(x,y)}$ is at least $(2t-1)b$.*

Proof. By Claim 16, if the strings are not disjoint, then there are $\ell_i^z \in L^z$ and $r_j^z \in R^z$ such that $\mathrm{d}_{\widetilde{H}^z}(\ell_i^z, r_j^z) \geq (2t-1)b$. Furthermore, by Claim 18, it holds that $\mathrm{d}_{G_{(x,y)}}(\ell_i^x, r_j^x) \geq \mathrm{d}_{\widetilde{H}^z}(\ell_i^z, r_j^z)$. Hence, there are two nodes in $G_{(x,y)}$ at distance at least $(2t-1)b$ from each other. ◀

Disjointness case

▷ Claim 20. If x and y are disjoint, then H^z is isomorphic to H. In particular, this implies:

1. H^z has diameter t.
2. for odd values of t, and for any two nodes $u, v \in V_{\widetilde{H}^z}$, if $u, v \in L^z$ or $u, v \in R^z$, then $\mathrm{d}_{\widetilde{H}^z}(u, v) = (t-1)b$.
3. for even values of t, and for any two nodes $u, v \in V_{\widetilde{H}^z}$ such that $u \in L^z$ and $v \in R^z$, it holds that $\mathrm{d}_{\widetilde{H}^z}(u, v) = (t-1)b$.

Proof. Observe that if the strings x and y are disjoint, then for any $h \in [K]$, either $x_h = 0$ or $y_h = 0$, so $z = x \wedge y$ is the all-zero string. Therefore, since H^z is obtained from H by keeping the edges that correspond to the 0 bits in z, H^z is isomorphic to H. Since H has diameter t, H^z also has diameter t.

Moreover, since H^z is isomorphic to H, by Property 1, it holds that for odd values of t, and for u, v that are on the same side (i.e., $u, v \in L^z$ or $u, v \in R^z$), it holds that $\mathrm{d}_{H^z}(u, v) = t - 1$, and therefore $\mathrm{d}_{\widetilde{H}^z}(u, v) = (t-1)b$.

Similarly, by Property 2, it holds that for even values of t, and for u, v that are not on the same side (i.e., $u \in L^z$ and $v \in R^z$), it holds that $\mathrm{d}_{H^z}(u, v) = t - 1$, and therefore $\mathrm{d}_{\widetilde{H}^z}(u, v) = (t-1)b$. ◁

For two nodes $u, v \in L^x \cup L^y \cup R^x \cup R^y$ in $G_{(x,y)}$, we say that u and v are on the same side if $u, v \in L^x \cup L^y$ or $u, v \in R^x \cup R^y$. Similarly, we say that u and v are on different sides if $u \in L^x \cup L^y$ and $v \in R^x \cup R^y$.

▷ Claim 21. For odd values of t, if x and y are disjoint then for any two nodes $u, v \in L^x \cup L^y \cup R^x \cup R^y$ that are on the same side (i.e., either $u, v \in L^x \cup L^y$ or $u, v \in R^x \cup R^y$) it holds that $\mathrm{d}_{G_{(x,y)}}(u, v) \leq (t-1)b + t$.

Due to space limitations, the proof of Claim 21 is deferred to the full version.

▷ Claim 22. For even values of t, if x and y are disjoint then for any two nodes $u, v \in L^x \cup L^y \cup R^x \cup R^y$ that are on different sides (i.e., $u \in L^x \cup L^y$ and $v \in R^x \cup R^y$) it holds that $\mathrm{d}_{G_{(x,y)}}(u, v) = (t-1)b + t$.

Due to space limitations, the proof of Claim 22 is deferred to the full version.

▶ **Lemma 23.** *If x and y are disjoint, then the diameter of $G_{(x,y)}$ is at most $tb+t+1$.*

Proof. Let $u, v \in V_{G_{(x,y)}}$ be two nodes in $G_{(x,y)}$. Let a_u^ℓ be the distance from u to the closest node in $L^x \cup L^y$, and let a_u^r be the distance from u to the closest node in $R^x \cup R^y$. a_v^ℓ and a_v^r are defined similarly for v. For example, if $u \in L^x \cup L^y$ then $a_u^\ell = 0$. The key point to note is that either $a_u^\ell + a_v^\ell \le b+1$, or $a_u^r + a_v^r \le b+1$. That is, either taking the two nodes to the left side of $G_{(x,y)}$ (i.e., to $L^x \cup L^y$) costs at most $b+1$, or taking the two nodes to the right side costs at most $b+1$. Similarly, it holds that either $a_u^\ell + a_v^r \le b+1$, or $a_u^r + a_v^\ell \le b+1$. The rest of the proof is by the following case analysis.

1. t is odd: By Claim 21, the distance between any two nodes in $L^x \cup L^y$ and the distance between any two nodes in $R^x \cup R^y$ is at most $(t-1)b+t$. Furthermore, we can either move u and v to $L^x \cup L^y$ by using at most $b+1$ steps (in total, for moving both u and v), or we can move u and v to $R^x \cup R^y$ by using at most $b+1$ steps (in total, for moving both u and v). After moving u and v to one of the sides, we can use Claim 21 and deduce that $\mathrm{d}_{G_{(x,y)}}(u,v) \le (t-1)b+t+b+1 = tb+t+1$.
2. t is even: By Claim 22, the distance between any $u' \in L^x \cup L^y$ and $v' \in R^x \cup R^y$ is at most $(t-1)b+t$. Furthermore, we can either move u to $L^x \cup L^y$ and v to $R^x \cup R^y$ by using at most $b+1$ steps (in total, for moving both u and v), or we can move u to $R^x \cup R^y$ and v to $L^x \cup L^y$ by using at most $b+1$ steps (in total, for moving both u and v). Hence, we can use Claim 22 and deduce that $\mathrm{d}_{G_{(x,y)}}(u,v) \le (t-1)b+t+b+1 = tb+t+1$. ◀

4.3 Proof of Theorem 15

First, we define the following partition $V = V_A \dot\cup V_B$ of the set of nodes of $G_{(x,y)}$. $V_A = V_{\widetilde{H}^x}$, and $V_B = V_{\widetilde{H}^y}$. Hence, the size of the cut $C = E(V_A, V_B)$ is $\Theta(p)$. This is because the only edges connecting between nodes in $\widetilde{H}^x$ and nodes in $\widetilde{H}^y$ in $G_{(x,y)}$ are the $2p$ edges of the matching between L^x and L^y, and the matching between R^x and R^y.

Since our goal is to show a lower bound as a function of the number of nodes n in the input graph $G_{(x,y)}$, we need to analyze the size of the cut and the size of the input strings with respect to n. By Theorem 12, we have that for $t \in \{3,4,6\}$, $K = \mathrm{Ex}(p,t) = \Omega(p^{1+\frac{1}{t-1}})$. Furthermore, by Remark 8 we can assume that a constant fraction of the bits in the strings x and y are 0. Hence, these 0 bits are translated to paths of length b in $G_{(x,y)}$. This implies that the number of nodes in $G_{(x,y)}$ is $n = \Theta(Kb) = \Omega(p^{1+\frac{1}{t-1}}b) = \Omega(p^{1+\frac{1}{t-1}})$ for constant values of b. Therefore, the size of the cut C is $\Theta(p) = O(n^{\frac{t-1}{t}})$.

Lemmas 19 and 23 imply that $\{G_{(x,y)} \mid x, y \in \{0,1\}^K\}$ is a family of lower bound graphs with respect to the set-disjointness function and the graph predicate that distinguishes between graphs of diameter $tb+t+1$ and graphs of diameter $(2t-1)b$. Hence, by Theorem 10 and the fact that the communication complexity of set-disjointness is $\Omega(K)$, we have that any algorithm for distinguishing between these two cases in the CONGEST model requires $\Omega\left(\frac{K}{|C|\log n}\right) = \Omega\left(\frac{n}{n^{(t-1)/t}\log n}\right) = \Omega\left(\frac{n^{1/t}}{\log n}\right)$ rounds. To get this lower bound for $(\frac{t}{2t-1}+\epsilon)$-approximation for diameter, we need that $\left(\frac{t}{2t-1}+\epsilon\right)(2t-1)b > tb+t+1$. Hence, we can pick $b = \Theta\left(\frac{t+1}{\epsilon(2t-1)}\right) = \Theta\left(\frac{1}{\epsilon}\right)$.

4.4 Handling the connectivity issue

One may wonder about the connectivity of the graph $G_{(x,y)}$. As the construction was described so far, there could be some values of x and y such that $G_{(x,y)}$ is not connected. In this section, we show how to slightly modify the construction of $G_{(x,y)}$ so that it is always connected, and in fact of constant diameter, without changing the analysis. Observe that it suffices to make $\widetilde{H}^x$ always connected. This is because any node in $\widetilde{H}^y$ has some path connecting it to a node in $\widetilde{H}^x$. Since $\widetilde{H}^x$ is obtained from H^x by stretching each edge in H^x to a path of length b, it suffices to make H^x always connected, regardless of the input string x.

Recall that given a string $x \in \{0,1\}^K$, we defined H^x to be the graph obtained from H where we keep only the edges that correspond to the 0 bits in x. Recall that H is a balanced bipartite graph of size $2p$, diameter t and girth $2t$. Of course, H is always connected. But since some of the edges of H may not exist in H^x, H^x may not be connected. To ensure that H^x is connected, let S be a shortest paths tree starting from an arbitrary node in H. Of course, the number of edges in S is $O(p)$, which is small with respect to the size of the input string x which is $K = \mathrm{Ex}(p,t)$.

We modify the definition of H^x such that the edges that correspond to the spanning tree S always exist in H^x. In particular, their existence in H^x doesn't depend on x. For this, we need to modify the size of the string x to $K - |S| = \Theta(K)$, so that only the edges that are not in S depend on x. The proof that $\{G_{(x,y)} \mid (x,y) \in \{0,1\}^{K-|S|} \times \{0,1\}^K\}$ is a family of lower bound graphs remains exactly the same as in Section 4.2. Furthermore, since the size of x didn't change asymptotically, the deduced lower bound from Section 4.3 doesn't change asymptotically. Observe that under the new definition of H^x, the diameter of H^x is at most $2t$. Hence, the diameter of $\widetilde{H}^x$ is at most $2tb$. It is not very hard to verify that the diameter of $G_{(x,y)}$ in this case is at most $2tb + 2b + 2$, which is constant for constant values of t and b.

References

1. Amir Abboud, Keren Censor-Hillel, and Seri Khoury. Near-linear lower bounds for distributed distance computations, even in sparse networks. In *Distributed Computing - 30th International Symposium, DISC*, volume 9888 of *Lecture Notes in Computer Science*, pages 29–42. Springer, 2016. `doi:10.1007/978-3-662-53426-7_3`.
2. Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Christoph Lenzen. Fooling views: A new lower bound technique for distributed computations under congestion. *Distributed Computing*, pages 1–15, 2020.
3. Bertie Ancona, Keren Censor-Hillel, Yuval Efron, Mina Dalirrooyfard, and Virginia Vassilevska Williams. Distributed distance approximation. In *submission*, 2020.
4. Clark T. Benson. Minimal regular graphs of girths eight and twelve. *Canadian Journal of Mathematics*, 18:1091–1094, 1966. `doi:10.4153/CJM-1966-109-8`.
5. Karl Bringmann and Sebastian Krinninger. Brief announcement: A note on hardness of diameter approximation. In *31st International Symposium on Distributed Computing, DISC*, volume 91 of *LIPIcs*, pages 44:1–44:3. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. `doi:10.4230/LIPIcs.DISC.2017.44`.
6. Keren Censor-Hillel, Michal Dory, Janne H. Korhonen, and Dean Leitersdorf. Fast approximate shortest paths in the congested clique. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC*, pages 74–83. ACM, 2019. `doi:10.1145/3293611.3331633`.
7. Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, and Jukka Suomela. Algebraic methods in the congested clique. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC*, pages 143–152. ACM, 2015. `doi:10.1145/2767386.2767414`.
8. Keren Censor-Hillel, Telikepalli Kavitha, Ami Paz, and Amir Yehudayoff. Distributed construction of purely additive spanners. In *Proceedings of the 30th International Symposium on Distributed Computing, DISC*, pages 129–142, 2016.
9. Keren Censor-Hillel, Seri Khoury, and Ami Paz. Quadratic and near-quadratic lower bounds for the CONGEST model. In *31st International Symposium on Distributed Computing, DISC*, pages 10:1–10:16, 2017.

10 Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing, PODC*, pages 367–376, 2014.

11 Yuval Efron, Ofer Grossman, and Seri Khoury. Beyond alice and bob: Improved inapproximability for maximum independent set in CONGEST. In Yuval Emek and Christian Cachin, editors, *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event*, pages 511–520. ACM, 2020. doi:10.1145/3382734.3405702.

12 Orr Fischer, Tzlil Gonen, Fabian Kuhn, and Rotem Oshman. Possibilities and impossibilities for distributed subgraph detection. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures, SPAA*, pages 153–162. ACM, 2018.

13 Sebastian Forster and Danupon Nanongkai. A faster distributed single-source shortest paths algorithm. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 686–697. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00071.

14 Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1150–1162, 2012.

15 Zoltán Füredi and Miklós Simonovits. *The History of Degenerate (Bipartite) Extremal Graph Problems*, pages 169–264. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. doi:10.1007/978-3-642-39286-3_7.

16 Stephan Holzer, David Peleg, Liam Roditty, and Roger Wattenhofer. Distributed 3/2-approximation of the diameter. In *Proceedings of the 28th International Symposium on Distributed Computing, DISC*, pages 562–564, 2014.

17 Stephan Holzer and Nathan Pinsker. Approximation of distances and shortest paths in the broadcast congest clique. In *19th International Conference on Principles of Distributed Systems, OPODIS*, volume 46 of *LIPIcs*, pages 6:1–6:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.OPODIS.2015.6.

18 Stephan Holzer and Roger Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC*, pages 355–364, 2012.

19 Qiang-Sheng Hua, Haoqiang Fan, Lixiang Qian, Ming Ai, Yangyang Li, Xuanhua Shi, and Hai Jin. Brief announcement: A tight distributed algorithm for all pairs shortest paths and applications. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA*, pages 439–441, 2016.

20 Eyal Kushilevitz and Noam Nisan. *Communication Complexity.* Cambridge University Press, New York, NY, USA, 1997.

21 Christoph Lenzen and David Peleg. Efficient distributed source detection with limited bandwidth. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC*, pages 375–382, 2013.

22 David Peleg. *Distributed Computing: A Locality-Sensitive Approach.* Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.

23 David Peleg, Liam Roditty, and Elad Tal. Distributed algorithms for network diameter and girth. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 660–672, 2012.

24 Alexander A. Razborov. On the distributional complexity of disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992.

25 R. Singleton. On minimal graphs of maximum even girth. *Journal of Combinatorial Theory*, 1:306–332, December 1966. doi:10.1016/S0021-9800(66)80054-6.

26 Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the 11h Annual ACM Symposium on Theory of Computing, STOC*, pages 209–213, 1979.