



**HAL**  
open science

## k-Center Clustering in Distributed Models

Leyla Biabani, Ami Paz

► **To cite this version:**

Leyla Biabani, Ami Paz. k-Center Clustering in Distributed Models. SIROCCO 2024, May 2024, Vietri sul Mare, Italy. pp.83-100, 10.1007/978-3-031-60603-8\_5 . hal-04799381

**HAL Id: hal-04799381**

**<https://hal.science/hal-04799381v1>**

Submitted on 22 Nov 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# $k$ -Center Clustering in Distributed Models

Leyla Biabani, Eindhoven University of Technology  
Ami Paz, LISN — CNRS & Paris-Saclay University

## Abstract

The  $k$ -center problem is a central optimization problem with numerous applications for machine learning, data mining, and communication networks. Despite extensive study in various scenarios, it surprisingly has not been thoroughly explored in the traditional distributed setting, where the communication graph of a network also defines the distance metric.

We initiate the study of the  $k$ -center problem in a setting where the underlying metric is the graph’s shortest path metric in three canonical distributed settings: the LOCAL, CONGEST, and CLIQUE models. Our results encompass constant-factor approximation algorithms and lower bounds in these models, as well as hardness results for the bi-criteria approximation setting.

**keywords:**  $k$ -Center clustering; Distributed graph algorithms; Shortest path metric

## 1 Introduction

### 1.1 Distributed $k$ -Center

The  $k$ -center problem is a key optimization problem, seeks to locate a small set of points in space (“centers”) that minimize the maximal distance from them to any other point. The problem was extensively studied in the centralized setting, where the points are taken from a metric space, with or without guarantees on the metric. Over the years, there was also some work on the  $k$ -center problem in parallel and distributed models, looking to solve the problem faster using multiple computational units. As in the centralized setting, the points and metric under consideration in these works are taken from some metric space, and it is usually assumed that distances are given to the computational units by an oracle.

A natural setting for the  $k$ -center problem is that of a network. We consider a graph representing a communication network in a natural way, and the goal is to make  $k$  of the nodes into centers while minimizing the maximal distance between these nodes and all others. Here, the metric space is not arbitrary, but is the *graph metric* on the graph itself. This is significant, for example, when determining server placement within a network to minimize the maximum delay, as seen in scenarios like content distribution over the Internet.

Two related, yet different problems are the metric *facility location* problem, where the goal is to minimize the average distance from the servers (or equivalently, the sum of distances), and the *online  $k$ -server* problem, where the points and servers can move in space.

In this work, we initiate the study of  $k$ -center in the distributed setting, where an undirected graph represents both the communication network and the problem’s metric. We address the problem in the popular distributed models of LOCAL, CONGEST, and CLIQUE, and derive upper and lower bounds. While the problem was studied earlier in some of these distributed models, it is important to note that all the prior work considered points in an arbitrary metric space, independent of the communication graph. As far as we know, our work is the first to consider the natural setting where the input graph represents both the network used for communication and the metric on which  $k$ -center is solved.

## 1.2 Our Results and Techniques

We present upper and lower bounds for the  $k$ -center problem in different distributed models. The metric considered is the shortest-path metric, and the graphs are unweighted unless otherwise specified.

### The LOCAL Model

A simple and natural model for communication networks is the LOCAL model [22], where the network's nodes communicate by synchronously exchanging messages of unbounded size. The local computation time is neglected, and the complexity measure in this case is the number of communication rounds. In this model, any problem is solvable in  $D$  rounds, where  $D$  is the diameter of the graph, i.e., the largest distance between two nodes in it. Hence, our upper bounds should be read as the minimum of the given value and  $D$ , and the lower bounds only apply for graphs with a diameter larger than the lower bound value.

We start Section 3 by giving a relatively simple algorithm that finds a  $(2k + \epsilon)$ -approximate solution to the  $k$ -center problem in  $O(k/\epsilon)$  rounds of the LOCAL model. This result relies on the assumption that the nodes have unique IDs in  $1, \dots, n$  — if we assumed the range to be larger, even ensuring a specific number of centers would require global communication. We then show that reducing the approximation ratio to be below  $k - 1$  requires  $\Omega(n)$  time. We stress that  $D \leq n$ , so in  $\Omega(n)$  time the nodes can aggregate the full structure of the network and find an exact solution.

Put differently, our results present two extremes. On the one hand, if a large approximation ratio of at least  $2k + \epsilon$  is allowed, the problem is simple—for  $\epsilon, k$  constants, it is solvable in constant time. On the other hand, if a lower approximation ratio is required, e.g., a constant that does not depend on  $k$ , then the running time is so large that the nodes may as well compute an optimal solution.

### The CONGEST Model

The CONGEST model is a restrictive variant of the LOCAL model, where the messages are limited to  $O(\log n)$  bits. Specifically, problems are no longer trivially solvable in  $O(D)$  time, and getting a constant approximation is solvable in non-trivial time, but is challenging.

In Section 4 we present a 2-approximation CONGEST algorithm for the  $k$ -center problem running in  $O(kD)$  rounds. It constructs different BFS trees, sometimes from multiple sources simultaneously, and simulates a centralized greedy approximation algorithm for the  $k$ -center problem. On the other hand, in Section 5 we prove that improving the approximation ratio to below  $4/3$ , and hence also finding an exact solution, requires a much longer time: any algorithm for this problem, even randomized, must take  $\tilde{\Omega}(n/k)$  rounds, and this is true even for graphs of diameter as small as 12.

This lower-bound proof is rather involved: it uses a reduction to communication complexity with a twist. Proving CONGEST lower bounds by reducing them to communication complexity is common in the literature [1, 10, 28], where usually the solution for the CONGEST problem directly implies an answer communication complexity problem. In our reduction, a new post-processing phase is added, where the players must do extra computation and also communicate more after getting the solution to the (approximate)  $k$ -center problem and before finding an answer to the communication problem.

### The CLIQUE Model

Finally, in Section 6 we consider the more recent congested clique model, denoted CLIQUE. It resembles the CONGEST, but with an all-to-all communication—the communication graph is a clique, and the input graph is a subgraph of it on the same set of nodes. One might think that the

$O(kD)$ -round CONGEST algorithm will directly translate to an  $O(k)$ -round CLIQUE algorithm, as the communication network has a diameter 1. However, this is not the case: the CONGEST algorithm builds BFS trees in  $O(D)$  time, and this step cannot be translated to an  $O(1)$ -round subroutine in the CLIQUE model.

At a high level, we show that finding a 2-approximate solution for the  $k$ -center problem in this model can be done deterministically in the same time as computing all-pairs-shortest paths, which is  $O(n^{1/3})$  for weighted graphs, and  $O(n^{0.158})$  for unweighted graphs [5]. Interestingly, if  $\omega$ , the exponent of the (centralized) matrix multiplication problem, will be discovered to be 2, the time for distributed matrix multiplication algorithms 2-approximation of  $k$ -center will also be improved, to  $O(n^{\epsilon'})$  for any  $\epsilon' > 0$ , or even to be poly-logarithmic. Previous work implies that a  $(2 + \epsilon)$ -approximate solution for  $k$ -center can be found in only  $O(\text{poly log } n)$  rounds, even in weighted graphs [2]. By allowing higher approximation ratios we can get even faster algorithms, such as a  $(4 + \epsilon)$ -approximation in  $O(\text{poly log log } n + k)$  rounds, and up to an  $O(\log n)$ -approximation in  $O(k)$  rounds.

En route, we prove a new result regarding approximate  $k$ -center. A simple greedy algorithm [15] finds a 2-approximation of  $k$ -center, assuming the distance between every two nodes is known. We extend this claim in Lemma 4, to show that if only a one-way, multiplicative  $\alpha$ -approximation of the distances is known, then a similar greedy algorithm gives a  $(2\alpha)$ -approximate solution to the  $k$ -center problem. Hence, by studying distributed  $k$ -center, we also provide new insights to the centralized case, which could be of independent interest.

Unfortunately, current techniques cannot establish lower bounds in the CLIQUE model. Any non-trivial lower bound in this model will imply circuit complexity lower bounds, solving a long-standing and notoriously hard open problem [13].

## 2 Preliminaries

### 2.1 The $k$ -Center Problem

Consider a graph  $G = (V, E)$  with nodes  $V$  and edges  $E$  where  $|V| = n$ . The edges in set  $E$  may have weights or be unweighted, and we mainly focus on the former case, and state it explicitly when this is not the case. The length of a path is determined by the number of edges in the path when  $E$  is unweighted, and by the total weight of the edges in the path when  $E$  is weighted. The diameter of the graph  $G$ , denoted  $D$ , is the maximum distance between any two nodes in  $V$ , i.e.,  $D = \max_{u,v \in V} d(u, v)$ . For any node  $u \in V$ , the eccentricity is  $\text{ecc}(u) = \max_{v \in V} d(u, v)$ , so we also have  $D = \max_{u \in V} \text{ecc}(u)$ .

Let  $k \in \mathbb{N}$  be a given parameter. In the  $k$ -center problem, we aim to find a set  $S^* \subseteq V$  with a size of at most  $k$  that minimizes the maximum distance of any node of  $V$  to its nearest center in  $S$ . More formally, we seek a set  $S^*$  with  $|S^*| \leq k$  that minimizes the value  $\max_{v \in V} \min_{s \in S^*} d(v, s)$ . This value is denoted as  $\text{OPT}_k(G)$  or simply  $\text{OPT}_k$  when the context is clear. For  $\alpha \geq 1$ , an algorithm is considered  $\alpha$ -approximation if it can compute a solution  $S$  such that the distance of any node  $v \in V$  to its nearest node in  $S$  is at most  $\alpha \cdot \text{OPT}_k$ .

### 2.2 Computational Models

We consider three common computational models for studying distributed graph algorithms, namely the LOCAL, CONGEST and CLIQUE models. We model a communication network using its graph, with nodes representing computational units and edges representing communication links. We use  $n$  for the number of nodes (computational units) and assume they have unique ids in  $\{1, \dots, n\}$

(specifically, there is always a node with id 1). The computation proceeds in synchronous rounds, where in each round each node sends messages to its neighbors, receives messages from them, and updates its local state accordingly. The input is local, in the sense that each node initially knows only its id, list of neighbors, and if there are inputs such as edge weights, then also the weights of its incident edges. The outputs are similarly local, e.g., at the end of the algorithm’s execution each node should know if it is a center or not.

In the LOCAL model [27], the message sizes are unbounded, and an  $r$ -round algorithm is equivalent to having each node deciding by its distance- $r$  neighborhood [22]. The CONGEST model [10, 28] is similarly defined, but each node is limited to sending  $O(\log n)$ -bit messages to each of its neighbors in each round, where messages to different neighbors might be different from one another. This model allows each node to send, e.g., its id or the ids of some of its neighbors in a single round, but not its full list of neighbors. A common primitive in this model is that of constructing a BFS tree from a node. This is sometimes extended to multiple BFS trees, or trees of bounded depth. See, e.g., [17, 21].

Finally, we model a network with a congested all-to-all overly network by the CLIQUE model [13, 23]. In this model, the input is a network as before, but the communication is less limited: in each round, each node can send  $O(\log n)$ -bit messages to each other node in the graph, and not only to its neighbors. This allows the nodes, e.g., to re-distribute the inputs in constant time [11, 20], and compute all-pairs-shortest-paths in sub-linear time [5].

### 2.3 Communication Complexity

We prove lower bounds for the  $k$ -center problem in the CONGEST model using a reduction to *communication complexity*, a well-studied topic in theoretical computer science [19, 30].

In the two-party *set disjointness* (henceforth: disjointness) communication complexity problem, two players referred to as Alice and Bob get two  $\ell$ -bit strings  $x$  (Alice’s string) and  $y$  (Bob’s string), and communicate by exchanging messages on a reliable asynchronous channel. Their goal is to decide if the sets represented by the indicator vectors  $x$  and  $y$  are disjoint or not, i.e., if there is an index  $i$  such that  $x[i] = y[i] = 1$ , in which case we say they are not disjoint and the players must output 0, or otherwise, the sets are disjoint and they should output 1.

Alice and Bob follow some protocol indicating who should send messages at each step and what message to send. The communication complexity of a protocol (as a function of  $\ell$  is the maximal number of bits they exchange when executing it, and the deterministic communication complexity of a problem the minimal communication complexity of a protocol solving the problem. The randomized communication complexity is similarly defined, but the players may also use random bits when executing the protocol, and the success probability is at least  $2/3$ , when taken on the choice of random bits.

For the disjointness problem, there is an  $\Omega(\ell)$  lower bound, which holds for deterministic and randomized algorithms alike [19, Example 3.22].

### 2.4 Related Work

To our surprise, the  $k$ -center problem was not studied in traditional distributed computing models. We survey below works in related computational models, and works on the related problem of metric facility location in distributed settings.

### 2.4.1 $k$ -Center in Related Computational Models

Bandyapadhyay, Inamdar, Pai, and Pemmaraju [2] studied the  $k$ -center problem, along with the related uncapacitated facility location and  $k$ -median problems. They consider the  $k$ -machine model (the parameter  $k$  here need not be the same as in  $k$ -center) which is closely related to the CLIQUE model. As mentioned, their work implies a randomized  $(2 + \epsilon)$ -approximation  $O(\text{poly } \log n)$ -time  $k$ -center algorithm in the CLIQUE model. Chiplunkar, Kale, and Ramamoorthy [8] studied approximate  $k$ -center in *streaming* models, and also in a parallel model where multiple processors perform local computations and then send the results to a central processor. Surprisingly, these seem to be the only works on the  $k$ -center problem in distributed settings. Cruciani, Forster, Goranci, Nazari, and Skarlatos [9] recently studied the  $k$ -center problem in a centralized *dynamic graphs* setting.

### 2.4.2 Metric Facility Location

The metric facility location problem has attracted more attention in distributed settings than the  $k$ -center problem. Works on this problem consider models that share similarities with the ones we consider here, although typically the models are not precisely identical.

**The CONGEST Model** In an unpublished manuscript, Briest, Degener, Kempkes, Kling, and Pietrzyk [4] studied metric facility location in the CONGEST model. They focused on a bipartite graph, with one side representing facilities and the other representing clients. Some of their results appeared in the thesis of Pietrzyk [29]. This work improved upon previous work that also consider metric facility location in bipartite setting in the CONGEST model [25, 26].

**The CLIQUE Model** Gehweiler, Lammersen, and Sohler [14] studied metric facility location in a model resembling the CLIQUE. They present a 3-round randomized algorithm that gives a constant approximation factor, based on the method of Mettu and Plaxton [24].

**Distributed Large-Scale Computational Models** Inamdar, Pai, and Pemmaraju [18] studied metric facility location in the CLIQUE model, MPC and  $k$ -machine, and gave an  $O(1)$ -approximation algorithm using a Mettu-Plaxton-style algorithm. When considering the CLIQUE model with what they call “implicit metric”, their model coincides with ours. They also consider another input regime (“explicit metric”), but not in the LOCAL and CONGEST models, and not for the  $k$ -center problem. Their work improves upon earlier works on metric facility location in the CLIQUE model [3, 16].

## 3 The $k$ -Center Problem in the LOCAL Model

In this section, we consider the  $k$ -center problem in the LOCAL model. We start with a fast and simple  $((2 + \epsilon)k)$ -approximation algorithm, Algorithm 1.

**Lemma 1.** *For any  $\epsilon > 0$  there is a deterministic  $O(k/\epsilon)$ -round algorithm in the LOCAL model that gives a  $((2 + \epsilon)k)$ -approximate solution for the  $k$ -center problem.*

*Proof.* Let  $t = 2 + \frac{4}{\epsilon}$ . The algorithm starts by having the node with id 1 (or any other arbitrary node) initiating the construction of a BFS tree of depth  $tk$ . After  $tk$  rounds, the nodes report back up the tree if the tree construction algorithm terminated in all its branches (i.e., reached all the nodes) or not. If it terminated, we have  $D \leq tk$ , and in another  $O(tk)$  rounds node 1 can aggregate all the graph structure and find an optimal solution. Otherwise, node 1 becomes the only center.

---

**Algorithm 1** A  $((2 + \epsilon)k)$ -approximation in the LOCAL model

---

- 1:  $t \leftarrow 2 + \frac{4}{\epsilon}$
  - 2: Perform a BFS from  $v_1$  for  $tk$  rounds
  - 3: Report back on the tree if terminated on all branches under you
  - 4: **if** All branches terminated **then**
  - 5:     Aggregate all the graph to  $v_1$
  - 6:     Locally compute an optimal solution
  - 7:     Disseminate the solution on the BFS tree
  - 8: **else** Only  $v_1$  marks itself as a center
- 

We claim that this simple algorithm gives a  $(2 + \epsilon)k$  approximate solution; clearly, we only need to prove it for the case  $D > tk$ . The eccentricity of node 1 is at most  $D$ , which gives an upper bound on ALG, the quality of the solution given by the algorithm.

Recall that  $\text{OPT}_k$  is the largest distance from a node to its center in an optimal solution. Let  $\pi$  be the shortest path between two nodes of the largest distance, i.e., two nodes of distance  $D$ . In an optimal solution, there are at least  $D + 1 - k$  non-center nodes in  $\pi$ . Therefore, there is a center that covers at least  $(D + 1 - k)/k$  nodes of  $\pi$ . Hence,  $\text{OPT}_k$  is at least  $(D + 1 - k)/(2k) = D/2k + 1/2k - 1/2 \geq D/2k - 1$ .

Note that  $D > tk$  implies  $D/(tk) > 1$ . The approximation ratio is thus

$$\frac{\text{ALG}}{\text{OPT}_k} \leq \frac{D}{D/2k - 1} \leq \frac{D}{D/2k - D/(tk)} = \frac{2tk}{t - 2} = (2 + \epsilon)k$$

by the choice of  $t$ , as required. □

This algorithm utilizes the fact that local computation is not limited in the LOCAL model, so  $v_1$  can find an optimal solution. However, it is not strictly necessary to use all this computational power: in case  $v_1$  has to locally solve  $k$ -center, it can instead compute a 2-approximation in a greedy manner. The approximation ratio is still as required, and the local computation now takes polynomial time.

Next, we move to the main result of the section: a lower bound for the  $k$ -center problem in the LOCAL model. Our lower bound also works for bi-criteria algorithms, which return at most  $\beta k$  centers instead of at most  $k$  centers, for some  $\beta \geq 1$ ; if the algorithm returns at most  $k$  centers, we can simply set  $\beta = 1$ . Our lower bound states that any  $t$ -round algorithm cannot achieve an approximation ratio better than  $k - \frac{k^2 + k(\beta k - 1)(2t + 1)}{n + k}$ . If  $k^2 + k(\beta k - 1)(2t + 1) < n + k$  holds, then our lower bound states that it is not possible to get an approximation ratio better than  $k - 1$ . The typical case is the non-bi-criteria one, with  $k$  being a constant, and there our lower bound translates to stating that getting an approximation ratio better than  $k - 1$  requires linear time.

**Theorem 1.** *Let  $t, k \in \mathbb{N}$ ,  $\beta \geq 1$ . Any  $t$ -round bi-criteria deterministic LOCAL algorithm that solves the  $k$ -center problem and reports at most  $\beta k$  centers as the solution, cannot have an approximation ratio better than  $k - \frac{k^2 + k(\beta k - 1)(2t + 1)}{n + k}$ , where  $n > 2\beta kt$  is the size of the graph and  $\beta \geq 1$ .*

At a high-level, the proof considers a communication graph which is a cycle, an algorithm in the LOCAL model that is faster than the lower bound, and the centers chosen by it. In  $t$  rounds, each node (and specifically, each chosen center) can gather information only from a segment of  $2t + 1$  nodes. We thus create a new cycle by concatenating all the  $(2t + 1)$ -node segments around the chosen centers. By the choice of the parameters, this leaves some “leftover nodes” that are not in any segment, which are concatenated after all the segments. When executing the same algorithm on the new cycle, the




Figure 1: Illustration for the proof of Theorem 1. Left: cycle  $C$ . Right: cycle  $C'$

same nodes as before will become centers, as they gather exactly the same information. The “leftover nodes” are now far from all the chosen centers, rendering a bad approximation ratio, as claimed.

*Proof.* Let  $\mathcal{A}$  be an algorithm that finds an (approximate, bi-criteria) solution for the  $k$ -center problem. Let  $k' \leq \beta k$  be the maximum number of centers that  $\mathcal{A}$  reports for any cycle of length  $n$ , and let  $C$  be such a cycle. For simplicity, we label the nodes of  $C$  with numbers  $1, \dots, n$  in a clockwise order. Assume that  $c_1 < \dots < c_{k'}$  are the centers returned by  $\mathcal{A}$  in clockwise order. Since  $n > 2kt$ , there exists at least two consecutive centers such that the distance between them is more than  $2t$ . We introduce a new cycle  $C'$  by re-arranging the nodes and show that the solution that  $\mathcal{A}$  finds for  $C'$  has an approximation ratio of at least  $k - \frac{k^2 + k(\beta k - 1)(2t + 1)}{n + k}$ .

To build the cycle  $C'$ , we first define segments  $S_1, \dots, S_{k'}$  of  $C$  as follows. Each segment  $S_i$  is of the form  $S_i = [b_i, e_i]$ , which refers to the nodes  $b_i$  to  $e_i$  in the clockwise order. Roughly speaking, each  $S_i$  is a segment of  $2t + 1$  nodes centered around  $c_i$ , but since such segments might overlap, the exact definition is a bit more subtle.

We define  $e_i := \min(c_i + t, c_{i+1} - 1)$  for any  $1 \leq i < k'$  and  $e_{k'} := \min(c_{k'} + t, n + c_1 - 1)$ . We also define  $b_1 := \max(c_1 - t, n - c_{k'} + 1)$  and  $b_i := \max(c_i - t, e_{i-1} + 1)$  for any  $1 < i \leq k'$ . Note that  $e_k$  may be larger than  $n$ , in which case we consider  $S_{k'}$  as the segment starting at  $b_i$  and ending at  $e_{k'} - n$  in the clockwise order, i.e.  $S_{k'} = [b_{k'}, n] \cup [1, e_{k'} - n]$ . Similarly,  $b_1$  might be smaller than 1, in which case we consider nodes  $b_1 + n$  to  $e_1$  in the clockwise order, i.e.  $S_1 = [b_1 + n, n] \cup [1, e_1]$ .

Since we assumed  $n > 2\beta kt$ , then there is an  $i^*$  such that the distance between  $c_{i^*}$  and  $c_{i^*+1}$  is more than  $2t$  (we consider  $c_1$  as  $c_{i^*+1}$  if  $i^* = k'$ ). To build  $C'$ , we concatenate the segments  $S_{i^*+1}, S_{i^*+2}, \dots, S_{k'}$  followed by  $S_1, S_2, \dots, S_{i^*}$ , and finally all the remaining nodes of  $C$ , in a clockwise order (see Figure 1).

Now, we claim that there is an execution of  $\mathcal{A}$  on  $C'$  that reports  $c_1, \dots, c_{k'}$  as the centers. Since the distance between  $c_{i^*}$  and  $c_{i^*+1}$  is more than  $2t$ , then the neighbourhood of length  $t$  for the nodes  $c_{i^*}$  and  $c_{i^*+1}$  is disjoint. Along with the definition of the segments, we can conclude that the distance- $t$  neighborhood for any node  $c_i$  is the same in both  $C$  and  $C'$ . Hence, any node  $c_i$  receives the same information for both graphs  $C$  and  $C'$ , and therefore, each node  $c_i$  makes the same decisions. Since  $c_1, \dots, c_k$  chose to be the centers in  $C$ , they also chose to be the centers in  $C'$ . Moreover, we assumed that  $k'$  is the maximum number of centers that  $\mathcal{A}$  reports for a cycle of



length  $n$ , and hence, no other node chooses to be a center.

We next discuss the approximation ratio. According to the definition of  $S_i$ , each  $S_i$  is of size at most  $2t+1$ . Therefore, there are at most  $(k'-2)(2t+1)$  nodes in  $S_{i^*+2}, S_{i^*+3}, \dots, S_k, S_1, S_2, \dots, S_{i^*-1}$ . Then, in the clockwise path from  $c_{i^*+1}$  to  $c_{i^*}$  in  $C'$  there are at most  $(t+1) + (k'-2)(2t+1) + (t+1) = (k'-2)(2t+1) + 2t + 2$  nodes. This means that the distance from  $c_{i^*+1}$  to  $c_{i^*}$  in  $C'$  in the anti-clockwise order is at least  $n - ((k'-2)(2t+1) + 2t + 2) + 1 = n - (k'-1)(2t+1)$ , and note that there is no center in this path. Hence, there is a node in the anti-clockwise path from  $c_{i^*+1}$  to  $c_{i^*}$  in  $C'$  such that its distance to the nearest center is at least  $(n - (k'-1)(2t+1))/2$ .

On the other hand, the optimal solution has  $\text{OPT}_k \leq \lceil (n-k)/(2k) \rceil$ . To see this, we can choose the  $k$  centers such that the distance between any two consecutive centers in the clockwise order is  $\lceil n/k \rceil$  or  $\lfloor n/k \rfloor$ . Thus,  $\text{OPT}_k \leq \lceil (n-k)/(2k) \rceil$  holds and the approximation ratio of Algorithm  $\mathcal{A}$  is at least

$$\begin{aligned} \frac{(n - (k' - 1)(2t + 1))/2}{\lceil (n - k)/(2k) \rceil} &\geq \frac{(n - (k' - 1)(2t + 1))/2}{(n - k)/(2k) + 1} = k \frac{n - (k' - 1)(2t + 1)}{n - k + 2k} \\ &= k - \frac{k^2 + k(k' - 1)(2t + 1)}{n + k} \geq k - \frac{k^2 + k(\beta k - 1)(2t + 1)}{n + k}, \end{aligned}$$

which finishes the proof.  $\square$

## 4 A 2-Approximation in the CONGEST Model

In this section, we show how to achieve a 2-approximate  $k$ -center clustering in the CONGEST model in  $O(kD)$  rounds, where  $D$  is the diameter of the underlying graph. Algorithm 2 presents an overview of our technique. In the following, we explain each part of this algorithm in detail.

---

### Algorithm 2 A 2-approximation in the CONGEST model

---

- 1: Find the node  $v_{\min}$  with minimum id
  - 2:  $S \leftarrow \{v_{\min}\}$
  - 3: **for**  $k - 1$  times **do**
  - 4:     Perform a BFS from all the nodes in  $S$
  - 5:     Let  $v^*$  be the furthest node from  $S$ , breaking ties by id
  - 6:      $S \leftarrow S \cup \{v^*\}$
  - 7: Each node in  $S$  marks itself as a center
- 

As opposed to the LOCAL algorithm, here we do not need to use the assumption that the nodes ids are in  $\{1, \dots, n\}$ , and assuming they are taken from  $\{1, \dots, \text{poly } n\}$  suffices. The first step of Algorithm 2 is finding the node  $v_{\min}$  with the minimum id, which can be done in  $O(D)$  rounds (and can be skipped if the ids are in  $\{1, \dots, n\}$ ). To do this, we start a BFS from all nodes. In each round, each node may receive messages of BFS trees from multiple sources. If this happens, such a node only continue the BFS from the source with the minimum id it has seen so far, and ignores the BFS's for all other sources. Therefore, the only source that its BFS is not paused after  $D$  rounds is  $v_{\min}$ . Each leaf of a BFS tree reports the termination of the construction up the tree, and back to the tree's parent. Hence,  $v_{\min}$  becomes aware that it is the node with the lowest id in  $O(D)$  rounds. In another  $O(D)$  time,  $v_{\min}$  disseminates along its BFS tree the depth of this tree, and all nodes learn this value, to which we refer as  $D'$ . Observe that  $D' \leq D < 2D'$ , and  $2D'$  will be used when an upper bound on  $D$  is needed (e.g., for the time of each iteration described below).

We next set  $S = \{v_{\min}\}$ , and perform  $k - 1$  iterations. In each iteration, a BFS is performed, where the sources are all nodes in  $S$ . Each node chooses to join the first BFS tree which reaches to

it, breaking ties arbitrarily. After  $O(D)$  rounds, the BFS is done and each node knows its distance to  $S$ . We next need to find the node with minimum id among the furthest nodes to  $S$ , which we refer to it as  $v^*$ . This can be done in  $O(D)$  rounds. In each round, every node informs its neighbors which node has the maximum distance among those it is currently aware of, and if it knows more than one node with maximum distance, it only reports the one with the smallest id. After  $D'$  such rounds, node  $v^*$  knows that it is the furthest node to  $S$  with minimum id. Finally,  $v^*$  can be added to  $S$ . At the end of the algorithm, the nodes in  $S$  mark themselves as centers.

The proof for the approximation ratio of our algorithm comes from the approximation ratio of the known greedy approach by Gonzalez [15], which we applied in our algorithm. In Lemma 2 we formally state this approximation.

**Lemma 2** ([15]). *Let  $G$  be a graph, and  $k \geq 1$  be an integer. Assume  $S_1 = v_1$ , where  $v_1 \in G$  is an arbitrary node. For each  $1 < i \leq k$ , we have  $S_i = S_{i-1} \cup \{v_i\}$ , where  $v_i$  is a node of  $G$  with maximum distance to  $S_{i-1}$ . Then,  $S_k$  is a 2-approximate solution for  $k$ -center of  $G$ .*

We can now summarize this section in Theorem 2.

**Theorem 2.** *There exists a deterministic 2-approximation algorithm for the  $k$ -center problem in the CONGEST model that needs  $O(kD)$  communication rounds.*

## 5 A Lower Bound in the CONGEST Model

In this section, we show a lower bound on the number of communication rounds for any algorithm for the  $k$ -center problem in the CONGEST model with an approximation ratio better than  $4/3$ . We start with a lower bound for the 1-center problem and then extend it for the  $k$ -center problem for a general  $k$ .

### 5.1 A Lower Bound for the 1-Center Problem

To prove the lower bound, we use a graph that was introduced by Abboud, Censor-Hillel, Khoury, and Paz [1] in order to prove a lower bound for computing the radius of a graph. Let  $x$  and  $y$  be two binary strings of length  $\ell$ . The graph  $G_{x,y}$  is built as follows, on  $n = \Theta(\ell)$  nodes (see Fig. 2). On a high level, the graph consists of two main sets of  $\ell$  nodes each,  $A$  and  $B$ , and a path on 4 nodes:  $c_A, \bar{c}_A, c_B, \bar{c}_B$ . The input strings  $x$  and  $y$  for a set-disjointness problem are used to set which edges from  $A$  to  $\bar{c}_A$  exist (representing  $x$ ), and which edges from  $B$  to  $\bar{c}_B$  exist (representing  $y$ ). The rest of the graph is built in order to guarantee that the optimal solution for the 1-center problem,  $\text{OPT}_1$  satisfies  $\text{OPT}_1 = 4$  if and only if  $x$  and  $y$  are disjoint. Hence, by simulating a 1-center algorithm and finding the distance from the chosen center to all other nodes, Alice and Bob can decide if  $x$  and  $y$  are disjoint. Finally, if the algorithm is too fast, Alice and Bob can simulate it with too little communication, contradicting the communication-complexity lower bound for disjointness.

We now present the graph and the proof in detail. The graph  $G_{x,y}$  consists of node sets  $A, B, F_A, T_A, F_B,$  and  $T_B$ , as well as nodes  $c_A, \bar{c}_A, c_B,$  and  $\bar{c}_B$ . The set  $A$  is the set of  $\ell$  nodes  $a^0, a^1, \dots, a^{\ell-1}$  and the set  $B$  is the set of  $\ell$  nodes  $b^0, b^1, \dots, b^{\ell-1}$ . For  $S \in \{A, B\}$ , the set  $F_S$  consists of  $\lceil \log_2 \ell \rceil$  nodes  $f_S^0, f_S^1, \dots, f_S^{\lceil \log_2 \ell - 1 \rceil}$ . Similarly, the set  $T_S$  contains node  $t_S^0, t_S^1, \dots, t_S^{\lceil \log_2 \ell - 1 \rceil}$ . The edges of the graph  $G_{x,y}$  are described in the following.

- **Edges from  $A$  to  $F_A$  and  $T_A$  and edges from  $B$  to  $F_B$  and  $T_B$ .** Let  $0 \leq i < \ell$  and  $0 \leq h < \ell$ , and let  $\text{bin}_h^i$  be the  $h$ -th bit in the binary representation of  $i$ . If  $\text{bin}_h^i = 0$ , we connect  $a^i$  to  $f_A^h$  and we connect  $b^i$  to  $f_B^h$ . Otherwise, if  $\text{bin}_h^i = 1$ , we connect  $a^i$  to  $t_A^h$  and we



that  $d(a^i, b^i) \leq 3$ , and therefore, the second part of the claim holds. The rigorous proof of this claim can be found in [1, Claim 4].

Recall that  $x$  and  $y$  are not disjoint if there is a  $0 \leq i < \ell$  such  $x[i] = y[i] = 1$ , and otherwise, they are disjoint. In Lemma 3, we show that  $\text{OPT}_1(G_{x,y}) = 4$  if and only if  $x$  and  $y$  are disjoint.

**Lemma 3.** *If the strings  $x$  and  $y$  are disjoint then  $\text{OPT}_1(G_{x,y}) = 4$ . Otherwise,  $\text{OPT}_1(G_{x,y}) = 3$ .*

*Proof.* To prove the lemma, we first show that if  $\text{OPT}_1 < 4$ , then an optimal center of  $G_{x,y}$  for the 1-center problem is a node  $a^i \in A$ , such that  $x[i] = y[i] = 1$ . Assume that  $\text{OPT}_1 < 4$ . Then Claim 1(1) implies that the optimal center is in  $A$ . Let  $a^i \in A$  be an optimal center. Then,  $d(a^i, b^i) \leq 3$  should hold since we assume  $\text{OPT}_1 < 4$ . But  $d(a^i, b^i) \leq 3$  holds only if  $a^i$  is connected to  $\bar{c}_A$  and  $b^i$  is connected to  $\bar{c}_B$ , which means  $x[i] = y[i] = 1$ .

We next consider the case that  $x$  and  $y$  are disjoint. In this case, we have  $\text{OPT}_1 \geq 4$  since otherwise, our claim implies that there exists  $0 \leq i < \ell$  such that  $x[i] = y[i] = 1$  and contradicts the disjointness of  $x$  and  $y$ . Hence,  $\text{OPT}_1 \geq 4$  holds. On the other hand, all nodes are within distance 4 of  $a^0$ , and therefore,  $\text{OPT}_1 = 4$ .

Now, we consider the case that  $x$  and  $y$  are not disjoint. This means that there exist  $0 \leq i < \ell$  such that  $x[i] = y[i] = 1$ . We show that  $\text{ecc}(a^i) = 3$  and then  $\text{OPT}_1 = 3$ . Claim 1(2) states that  $d(a^i, u) \leq 3$  holds for any  $u \in V \setminus \{b^i, c_B\}$ . Besides,  $a^i$  is connected to  $\bar{c}_A$  as  $x[i] = 1$  and  $b^i$  is connected to  $\bar{c}_B$  since  $y[i] = 1$ . Thus,  $d(a^i, b^i) = d(a^i, c_B) = 3$ . Putting everything together we have  $\text{OPT}_1 = 3$ .  $\square$

**Theorem 3.** *Any CONGEST algorithm that returns an  $\alpha$ -approximation for the 1-center problem with  $\alpha < 4/3$ , even randomized, must take  $\Omega(n/\log^2 n)$  rounds to complete.*

*Proof.* Assume for contradiction a CONGEST algorithm faster than in the theorem's statement. Let  $x, y \in \{0, 1\}^\ell$  be two inputs of Alice and Bob for the disjointness problem. Alice and Bob simulate the algorithm on the graph  $G_{x,y}$  described above.

To this end, we split the graph node by  $V_A = A \cup F_A \cup T_A \cup \{c_A, \bar{C}_a, w^0, w^1, w^2\}$  and  $V_B = V \setminus V_A$  (the nodes on the left and right sides of Fig. 2, respectively). Alice is in charge of the nodes of  $V_A$  and Bob on  $V_B$ . To simulate a round, they locally simulate the exchange of messages in each of their node sets, and exchange bits to simulate the messages between nodes of  $V_A$  and  $V_B$ . Since there are  $O(\log n)$  edges between the sets, a simulation of a round requires  $O(\log^2 n)$  bits of communication.

If the output of the algorithm is a node not in  $A$ , Alice and Bob return 1 for the disjointness problem. If it is some node  $a^i \in A$ , they exchange the bits  $x[i]$  and  $y[i]$  (for the same index  $i$ ) and return 0 for the disjointness problem if and only if  $x[i] = y[i] = 1$ .

For correctness, first note that Alice and Bob return 0 for the disjointness problem only if they find an index such that  $x[i] = y[i] = 1$ , so this answer must be correct. When they return 1, on the other hand, it might be since the algorithm returned a center  $u \notin A$ , or a center  $a^i$  such that  $x[i] = 0$  or  $y[i] = 0$  (or both).

If Alice and Bob return 1 because of a center  $u \notin A$ , then Claim 1(1) guarantees that  $\text{ecc}(u) \geq 4$ , and by  $\alpha < 4/3$  we get that  $\text{OPT}_1 > 3$ . Have the sets not been disjoint, Lemma 3 guarantees that  $\text{OPT}_1 = 3$ , a contradiction.

If Alice and Bob return 1 because of a center  $a^i \in A$ , then we also know that  $x[i] = 0$  or  $y[i] = 0$ . Note that  $d(a^i, b^i) \geq 4$ : the nodes of  $\text{bin}(a^i)$  and  $\text{bin}(b^i)$  are never neighbors, so any 3-path connecting  $a^i$  and  $b^i$  must go through the edge  $(\bar{c}_A, \bar{c}_B)$ ; but since  $x[i] = 0$  or  $y[i] = 0$ , no such path can exist. Hence,  $\text{ecc}(a^i) \geq 4$ . As before, if the sets were not disjoint, Lemma 3 guarantees that  $\text{OPT}_1 = 3$ , and we would have got a contradiction.

For the complexity, let  $T$  be the number of rounds used by the algorithm. To simulate these rounds, Alice and Bob exchange  $O(T \log^2 n)$  bits. As shown above, they solve the disjointness

problem on  $\ell$  bits, which requires them to communicate  $\Omega(\ell)$  bits, even when using a randomized algorithm. As  $\ell = \Theta(n)$ , we get  $T = \Omega(n/\log^2 n)$ . In fact, Alice and Bob must also exchange the id of the node chosen as the center, but these  $\log n$  bits of communication do not affect the asymptotic complexity.

Finally, note that Alice and Bob cannot solve disjointness solely by the output of the algorithm, and need extra communication after it. Hence, we cannot utilize standard reductions such as [1, Theorem 6] and use a non-standard one.  $\square$

## 5.2 Extending the Lower Bound to the $k$ -Center Problem

To generalize the lower bound for the 1-center problem to the  $k$ -center problem, we consider  $k$  copies of the graph  $G_{x,y}$  introduced in Section 5.1, such that all copies share the node  $w^2$ . We refer to this graph as  $G_{x,y}^k$ .

**Claim 2.** *In any  $\alpha$ -approximate solution for  $k$ -center of  $G_{x,y}^k$ ,  $\alpha < 3/2$ , exactly one center is chosen from each copy.*

*Proof.* Note that Lemma 3 guarantees that in each copy we can choose a center and get a solution with distance at most 4 from each node to a center, and  $\text{OPT}_k \leq 4$ . If some copy contains no center, then  $b^0$  in this copy has distance at least 6 to  $w^2$ , and the solution is at least  $(6/4)$ -approximation, a contradiction.  $\square$

**Theorem 4.** *Any CONGEST algorithm that returns an  $\alpha$ -approximation for the  $k$ -center problem with  $\alpha < 4/3$ , even randomized, must take  $\Omega(n/(k^2 \log^2 n))$  rounds to complete.*

The proof extends the proof of Theorem 3 by using the graph  $G_{x,y}^k$  described above. Note that all copies are for the same input pair  $(x, y)$ .

*Proof.* Assume a faster algorithm than in the theorem statement. Alice and Bob simulate this algorithm for  $\ell$  input bits, where  $\ell = \Theta(n/k)$ , and get an output at the form of  $k$  center nodes. If none of the output nodes is of the form  $a^i \in A$  for some copy, Alice and Bob answer 1 for the disjointness problems. Otherwise, for every node  $a^i \in A$  in the output (perhaps from different copies), they exchange  $x[i], y[i]$ , and answer 0 for the disjointness problem only if  $x[i] = y[i] = 1$  for at least one such  $i$ .

For correctness, recall first that Claim 2 guarantees that each copy contains exactly one center.

If Alice and Bob answer 0, this is because they found an index  $i$  of intersection and the algorithm is correct. If they answer 1, the proof follows the same lines of the proof of Theorem 3: if the sets are not disjoint then Lemma 3 guarantees each copy has a center  $a^i$  with  $\text{ecc}(a^i) \leq 3$ , and we get  $\text{OPT}_k \leq 3$ . On the other hand, the algorithm returns at each copy a node  $u \notin A$  or a node  $a^i$  with  $x[i] = 0$  or  $y[i] = 0$ , which has  $d(a^i, b^i) \geq 4$ ; in both cases, the solution is not an  $\alpha$ -approximation for  $\alpha < 4/3$ .

For the complexity, assume the algorithm takes  $T$  rounds. Alice and Bob communicate, for each copy and each round,  $O(\log(n/k) \log n)$  bits, which are  $O(\log^2 n)$  bits. At the end of the algorithm, they may exchange the indices of at most  $k$  centers of the form  $a^i$  and their inputs  $x[i]$  and  $y[i]$  in these locations, with takes requires another  $O(k \log n)$  bits. In total, they exchange  $O(Tk \log^2 n + k \log n)$  bits, which is  $O(Tk \log^2 n)$ .

On the other hand, they solve disjointness on  $\ell = \Theta(n/k)$  input bits, so they must communicate  $\Omega(n/k)$  bits. Hence,

$$Tk \log^2 n \geq cn/k$$

for some constant  $c$ , and we get

$$T = \Omega\left(\frac{n}{k^2 \log^2 n}\right)$$

as claimed. □

## 6 A 2-Approximation Algorithm in the CLIQUE Model

We now present Algorithm 3, which provides an approximate solution to the  $k$ -center clustering problem in the CLIQUE model. It is applicable both when the distance between the nodes are weighted and when they are unweighted.

The algorithm consists of two phases. The first phase is computing all-pairs shortest distances or approximating them, using known algorithms.

The second phase is to greedily find the centers, which is done in additional  $k$  rounds. In this phase, we first find  $v_{\min}$ , the node with the minimum id. This is trivial when the ids are in  $\{1, \dots, n\}$ , but easy also without this assumption, by having each node send its id to all other nodes. We then set  $S \leftarrow \{v_{\min}\}$ . Next, we have  $k - 1$  iteration, and in each iteration we find  $v^*$ , which is the furthest node from  $S$  (if there is more than one furthest node, we define  $v^*$  as the furthest node with minimum id). To accomplish this, it is enough that each node sends its distance to  $S$  to all nodes. Hence, all nodes can know who  $v^*$  is in one communication round. At the end of each iteration, we set  $S$  to  $S \cup \{v^*\}$ , and report  $S$  as the set of centers at the end of the algorithm. This process is formalized in the following algorithm.

---

**Algorithm 3** Varying approximation ratios in the CLIQUE model

---

- 1: Compute all-pairs shortest paths
  - 2: Let  $v_{\min}$  be the node with minimum id
  - 3:  $S \leftarrow \{v_{\min}\}$
  - 4: **for**  $k - 1$  times **do**
  - 5:     Let  $v^*$  be the furthest node to  $S$
  - 6:      $S \leftarrow S \cup \{v^*\}$
  - 7: Each node in  $S$  marks itself as a center
- 

If we calculate the exact all-pairs shortest paths in the first phase of Algorithm 3, Lemma 2 indicates that the set  $S$  computed by the algorithm is a 2-approximate solution for  $k$ -center.

Approximate distance can be computed much faster than exact ones, and we next show that Algorithm 3 can also work with approximate distances, in which case it computes an approximate  $k$ -center solution. To prove this, we now prove Lemma 4, which extends Lemma 2 to the scenario where only a multiplicative  $\alpha$ -approximation with one-sided error of all-pairs shortest paths is computed in the first phase.

**Lemma 4.** *For any  $\alpha \geq 1$ , if an  $\alpha$ -approximation of all-pairs shortest paths is computed in the first phase of Algorithm 3, then the set  $S$  obtained by the algorithm is a  $(2\alpha)$ -approximate solution for the  $k$ -center problem.*

*Proof.* If  $n \leq k$ , the algorithm trivially returns all the nodes as the centers, so we assume  $n > k$  in the following. Let  $S$  be the set returned by Algorithm 3,  $\hat{v} \in V \setminus S$  the furthest node from  $S$ , and  $\hat{r}$  its distance.

First, we claim that  $\text{OPT}_k \geq \hat{r}/(2\alpha)$ . Since  $\hat{v}$  is not added to  $S$ , we can conclude that the distance between any two nodes in  $S$  is at least  $\hat{r}/\alpha$ . Thus, the distance between any two nodes

Table 1: Approximation times in the CLIQUE model, for constant  $0 < \epsilon < 1$

Approx. ratio	Time	Weighted?	Deter.?	Ref.
2	$O(n^{1/3} + k)$	No	Yes	[5] & Thm. 5
2	$O(n^{0.158} + k)$	Yes	Yes	[5] & Thm. 5
$2 + \epsilon$	$O(n^{0.158} + k)$	No	Yes	[5] & Thm. 5
$2 + \epsilon$	$O(\text{poly log } n)$	Yes	No	[2]
$4 + \epsilon$	$O(\log^4 \log n + k)$	No	Yes	[12] & Thm. 5
$6 + \epsilon$	$O(\log^2 n + k)$	Yes	Yes	[6] & Thm. 5
$O(\log n)$	$O(k)$	Yes	No	[7] & Thm. 5

in  $S \cup \{\hat{v}\}$  is at least  $\hat{r}/\alpha$  as well. In addition, in any optimal solution there exist two nodes in  $S \cup \{\hat{v}\}$  that have the same nearest center in that optimal solution since  $|S \cup \{\hat{v}\}| = k + 1$ . Since the distance between these two nodes is at least  $\hat{r}/\alpha$ , at least one of them has distance at least  $\hat{r}/(2\alpha)$  from their common center, implying  $\text{OPT}_k \geq \hat{r}/(2\alpha)$ . On the other hand, all nodes in  $V$  are within distance  $\hat{r}$  from  $S$  by the definition of  $\hat{r}$ . Hence,  $S$  is a solution with the approximation ratio of at most  $\frac{\hat{r}}{\hat{r}/(2\alpha)} = 2\alpha$ .  $\square$

All is left now is to plug fast CLIQUE distance computation algorithms [5, 6, 7, 12] in the lemma, and we get fast algorithms for exact and approximate  $k$ -center. This yields the results detailed in Theorem 5 and Table 1, showing that  $k$ -center can be approximated in the same times as all pairs shortest paths computation, up to an additive  $O(k)$  time. Note that for the specific case of  $(2 + \epsilon)$ -approximation, a much faster (randomized) algorithm exists [2], running in  $O(\text{poly log } n)$  rounds.

**Theorem 5.** *There exists a deterministic CLIQUE algorithm for the  $k$ -center problem which gives a 2-approximation in  $O(n^{1/3} + k)$  rounds on weighted graphs and  $O(n^{0.158} + k)$  rounds on unweighted graphs.*

*For every constant  $0 < \epsilon < 1$ , there exist deterministic algorithms that give a  $(2 + \epsilon)$ -approximation in  $O(n^{0.158} + k)$  rounds on unweighted graphs, a  $(4 + \epsilon)$ -approximation in  $O(\log^4 \log n + k)$  rounds on unweighted graphs, and a  $(6 + \epsilon)$ -approximation in  $O(\log^2 n + k)$  rounds on weighted graphs.*

*There exists a randomized algorithm which gives a  $O(\log n)$ -approximation in  $O(k)$  rounds on both weighted and unweighted graphs.*

## Acknowledgments

We thanks Morteza Monemizadeh for many valuable conversations on this project, Michal Dory for discussions regarding distance computation in the CLIQUE model, and the anonymous reviewers of SIROCCO'24 for their helpful comments.

## References

- [1] A. Abboud, K. Censor-Hillel, S. Khoury, and A. Paz. Smaller cuts, higher lower bounds. *ACM Trans. Algorithms*, 17(4):30:1–30:40, 2021. doi: 10.1145/3469834. URL <https://doi.org/10.1145/3469834>.
- [2] S. Bandyapadhyay, T. Inamdar, S. Pai, and S. V. Pemmaraju. Near-optimal clustering in the  $k$ -machine model. *Theor. Comput. Sci.*, 899:80–97, 2022.

- [3] A. Berns, J. Hegeman, and S. V. Pemmaraju. Super-fast distributed algorithms for metric facility location. In *ICALP (2)*, volume 7392 of *Lecture Notes in Computer Science*, pages 428–439. Springer, 2012.
- [4] P. Briest, B. Degener, B. Kempkes, P. Kling, and P. Pietrzyk. A distributed approximation algorithm for the metric uncapacitated facility location problem in the congest model. *CoRR*, abs/1105.1248, 2011.
- [5] K. Censor-Hillel, P. Kaski, J. H. Korhonen, C. Lenzen, A. Paz, and J. Suomela. Algebraic methods in the congested clique. *Distributed Comput.*, 32(6):461–478, 2019.
- [6] K. Censor-Hillel, M. Dory, J. H. Korhonen, and D. Leitersdorf. Fast approximate shortest paths in the congested clique. *Distributed Comput.*, 34(6):463–487, 2021.
- [7] S. Chechik and T. Zhang. Constant-round near-optimal spanners in congested clique. In *PODC*, pages 325–334. ACM, 2022.
- [8] A. Chiplunkar, S. S. Kale, and S. N. Ramamoorthy. How to solve fair k-center in massive data models. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 1877–1886. PMLR, 2020.
- [9] E. Cruciani, S. Forster, G. Goranci, Y. Nazari, and A. Skarlatos. Dynamic algorithms for k-center on graphs. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3441–3462. SIAM, 2024.
- [10] A. Das Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, and R. Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.*, 41(5):1235–1265, 2012.
- [11] D. Dolev, C. Lenzen, and S. Peled. "tri, tri again": Finding triangles and small subgraphs in a distributed setting - (extended abstract). In *DISC*, volume 7611 of *Lecture Notes in Computer Science*, pages 195–209. Springer, 2012.
- [12] M. Dory and M. Parter. Exponentially faster shortest paths in the congested clique. *J. ACM*, 69(4):29:1–29:42, 2022.
- [13] A. Drucker, F. Kuhn, and R. Oshman. On the power of the congested clique model. In *PODC*, pages 367–376. ACM, 2014.
- [14] J. Gehweiler, C. Lammersen, and C. Sohler. A distributed  $o(1)$ -approximation algorithm for the uniform facility location problem. *Algorithmica*, 68(3):643–670, 2014.
- [15] T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theor. Comput. Sci.*, 38:293–306, 1985. doi: 10.1016/0304-3975(85)90224-5. URL [https://doi.org/10.1016/0304-3975\(85\)90224-5](https://doi.org/10.1016/0304-3975(85)90224-5).
- [16] J. W. Hegeman and S. V. Pemmaraju. Sub-logarithmic distributed algorithms for metric facility location. *Distributed Comput.*, 28(5):351–374, 2015.
- [17] S. Holzer and R. Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *PODC*, pages 355–364. ACM, 2012.



- [18] T. Inamdar, S. Pai, and S. V. Pemmaraju. Large-scale distributed algorithms for facility location with outliers. In *OPODIS*, volume 125 of *LIPICs*, pages 5:1–5:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [19] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997. ISBN 0-521-56067-5.
- [20] C. Lenzen. Optimal deterministic routing and sorting on the congested clique. In *Proc. 32nd ACM Symposium on Principles of Distributed Computing (PODC 2013)*, pages 42–50, 2013. doi: 10.1145/2484239.2501983.
- [21] C. Lenzen, B. Patt-Shamir, and D. Peleg. Distributed distance computation and routing with small messages. *Distributed Comput.*, 32(2):133–157, 2019.
- [22] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [23] Z. Lotker, E. Pavlov, B. Patt-Shamir, and D. Peleg. MST construction in  $o(\log \log n)$  communication rounds. In *SPAA*, pages 94–100. ACM, 2003.
- [24] R. R. Mettu and C. G. Plaxton. The online median problem. *SIAM J. Comput.*, 32(3):816–832, 2003.
- [25] T. Moscibroda and R. Wattenhofer. Facility location: distributed approximation. In *PODC*, pages 108–117. ACM, 2005.
- [26] S. Pandit and S. V. Pemmaraju. Return of the primal-dual: distributed metric facility location. In *PODC*, pages 180–189. ACM, 2009.
- [27] D. Peleg. *Distributed Computing: A Locality-sensitive Approach*. Society for Industrial and Applied Mathematics, 2000. ISBN 0-89871-464-8.
- [28] D. Peleg and V. Rubinovich. Near-tight lower bound on the time complexity of distributed MST construction. *SIAM Journal on Computing*, 30(5):1427–1442, 2000. doi: 10.1137/S0097539700369740.
- [29] P. Pietrzyk. *Local and online algorithms for facility location*. PhD thesis, University of Paderborn, 2013.
- [30] A. Rao and A. Yehudayoff. *Communication Complexity: and Applications*. Cambridge University Press, 2020. ISBN 9781108497985.