# MODERN ASPECTS OF UNSUPERVISED LEARNING 

A Thesis<br>Presented to<br>The Academic Faculty<br>by<br>Yingyu Liang

In Partial Fulfillment of the Requirements for the Degree<br>Doctor of Philosophy in the School of Computer Science

Georgia Institute of Technology
August 2014

## MODERN ASPECTS OF UNSUPERVISED LEARNING

Approved by:

Maria-Florina Balcan, Advisor
School of Computer Science Georgia Institute of Technology

Avrim Blum
School of Computer Science
Carnegie Mellon University
Lance Fortnow
School of Computer Science
Georgia Institute of Technology

Charles L. Isbell, Jr.
School of Interactive Computing
Georgia Institute of Technology
Dana Randall
School of Computer Science
Georgia Institute of Technology
Le Song
School of Computational Science and
Engineering
Georgia Institute of Technology
Date Approved: 21 May 2014

To my parents,

Qunhong Liang and Guixian Qin.

## ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor, Maria-Florina Balcan, for her excellent guidance, caring, and patience. She has been a tremendous mentor for me, providing me with an excellent model for doing research with great enthusiasm and sharp insights. Her advice on both research as well as on my career has been priceless. I would also like to thank my committee members, Avrim Blum, Lance Fortnow, Charles L. Isbell, Jr., Dana Randall, and Le Song for serving as my committee members and providing stimulating suggestions and comments for this thesis.

I would like to thank Lev Reyzin, Steven Ehrlich, and Christopher Berlind, who were always willing to help and give their best suggestions on my research and my writing. Many thanks to Anand Louis, Ying Xiao, Arindam Khan, Abhishek Banerjee, Shang-Tse Chen, Nan Du, Bo Dai, He Niao and many others in the laboratory who have made my years at Georgia Institute of Technology wonderful and productive.

I would like to express my special thanks to my family. Words cannot express how grateful I am to my parents, parents-in-law, my sisters and brother for all of the sacrifices that they have made on my behalf. They are always supporting me and encouraging me with their best wishes. At the end, I would like to thank my beloved wife Ting Jiao, who is always there cheering me up and stands by me through all times.

## TABLE OF CONTENTS

DEDICATION ..... iii
ACKNOWLEDGEMENTS ..... iv
LIST OF TABLES ..... vii
LIST OF FIGURES ..... viii
SUMMARY ..... x
I INTRODUCTION ..... 1
1.1 Clustering under Perturbation Resilience ..... 5
1.2 Distributed Clustering ..... 6
1.3 Community Detection ..... 8
1.4 Summary and Bibliographic Information ..... 8
II CLUSTERING UNDER PERTURBATION RESILIENCE ..... 10
2.1 Preliminaries ..... 15
$2.2 \quad \alpha$-Perturbation Resilience for Center-Based Objectives ..... 17
$2.3(\alpha, \epsilon)$-Perturbation Resilience for $k$-Median ..... 21
2.3.1 Structure of $(\alpha, \epsilon)$-Perturbation Resilient Instances ..... 22
2.3.2 Approximating the Optimum Clustering ..... 29
2.3.3 Sublinear Time Algorithm ..... 37
$2.4 \alpha$-Perturbation Resilience for Min-Sum ..... 39
2.4.1 Sublinear Time Algorithm ..... 43
$2.5(\alpha, \epsilon)$-Perturbation Resilience for Min-Sum ..... 47
2.5.1 Structure of $(\alpha, \epsilon)$-Perturbation Resilient Instances ..... 48
2.5.2 Approximating the Optimal Clustering ..... 56
III DISTRIBUTED CLUSTERING ..... 63
3.1 Preliminaries ..... 65
3.1.1 Coresets ..... 66
3.1.2 Principal Component Analysis ..... 67
3.2 Distributed Coreset-Based Clustering ..... 68
3.2.1 Distributed Coreset Construction ..... 71
3.2.2 Effect of Network Topology on Communication Cost ..... 78
3.3 Distributed $k$-Means Clustering of High Dimensional Data ..... 81
3.3.1 Distributed PCA ..... 83
3.3.2 Distributed Clustering ..... 85
3.4 Experiments ..... 86
3.4.1 Experiments on Distributed Coreset-Based Clustering ..... 86
3.4.2 Experiments on High Dimensional Data ..... 90
IV COMMUNITY DETECTION ..... 93
4.1 Hierarchical Community Model ..... 94
4.2 Hierarchical Community Detection Algorithm ..... 99
4.3 Local Community Detection Algorithm ..... 106
4.4 Experiments ..... 109
4.4.1 Evaluation on Real-World Networks ..... 110
4.4.2 Evaluation on Synthetic Networks ..... 112
APPENDIX A - CLUSTERING UNDER PERTURBATION RE- SILIENCE ..... 115
APPENDIX B - DISTRIBUTED CLUSTERING ..... 139
REFERENCES ..... 154
VITA ..... 162

## LIST OF TABLES

1 The parameters of the synthetic data sets for community detection. . 112

## LIST OF FIGURES

1 The closure distance. ..... 19
2 Comparing $d$ and $d_{P}\left(A, A^{\prime}\right)$ in closure linkage. ..... 20
3 Different types of points when bounding the number of bad points for $(\alpha, \epsilon)$-perturbation resilient $k$-median instances. ..... 23
4 The optimality of the bound on the number of bad points for $(\alpha, \epsilon)$ - perturbation resilient $k$-median instances ..... 27
5 Perturbation construction. ..... 50
6 Different types of points when bounding the number of bad points for $(\alpha, \epsilon)$-perturbation resilient min-sum instances. ..... 51
7 Properties of the potential good points. ..... 55
$8 \quad$ Notations in Claim 4 and 7. ..... 59
9 Notations in Claim 8 ..... 59
10 Illustration of different coreset construction approaches. ..... 69
11 The key points of the distributed PCA algorithm. ..... 83
12 Normalized $k$-means cost v.s. communication cost over graphs ..... 89
13 Normalized $k$-means cost v.s. communication cost over the spanning trees ..... 90
14 Normalized $k$-means cost v.s. communication cost for different projec- tion dimensions ..... 92
$15 \alpha$-compact blob ..... 95
$16(\alpha, \beta)$-stable community ..... 96
$17(\alpha, \beta, \nu)$-stable community ..... 98
18 Results of community detection on real world networks ..... 111
19 The average error of community detection on synthetic networks ..... 112
20 The running time of community detection on synthetic networks ..... 113
21 Notations in Lemma 24. ..... 128
22 Notations in Claim 4. ..... 132
23 Notations in Claim 6. ..... 134
24 Notations in Lemma 15 and Claim 9. ..... 137
25 Experimental results of distributed clustering on random graphs ..... 143
26 Experimental results of distributed clustering on grid and preferential graphs ..... 144
27 Experimental results of distributed clustering on the spanning trees of the random graphs ..... 145
28 Experimental results of distributed clustering on the spanning trees of the grid and preferential graphs ..... 146

## SUMMARY

Unsupervised learning has become more and more important due to the recent explosion of data. Clustering, a key topic in unsupervised learning, is a well-studied task arising in many applications ranging from computer vision to computational biology to the social sciences. This thesis is a collection of work exploring two modern aspects of clustering: stability and scalability.

In the first part, we study clustering under a stability property called perturbation resilience. As an alternative approach to worst case analysis, this novel theoretical framework aims at understanding the complexity of clustering instances that satisfy natural stability assumptions. In particular, we show how to correctly cluster instances whose optimal solutions are resilient to small multiplicative perturbations on the distances between data points, significantly improving existing guarantees. We further propose a generalized property that allows small changes in the optimal solutions after perturbations, and provide the first known positive results in this more challenging setting.

In the second part, we study the problem of clustering large scale data distributed across nodes which communicate over the edges of a connected graph. We provide algorithms with small communication cost and provable guarantees on the clustering quality. We also propose algorithms for distributed principal component analysis, which can be used to reduce the communication cost of clustering high dimensional data while merely comprising the clustering quality.

In the third part, we study community detection, the modern extension of clustering to network data. We propose a theoretical model of communities that are
stable in the presence of noisy nodes in the network, and design an algorithm that provably detects all such communities. We also provide a local algorithm for large scale networks, whose running time depends on the sizes of the output communities but not that of the entire network.

## CHAPTER I

## INTRODUCTION

This thesis develops new frameworks and designs algorithms for new aspects of clustering, a key topic in unsupervised learning. These modern aspects of clustering, including a focus on scalability and stability, are areas of significant practical importance and rising concerns. Traditional clustering analysis tends not to capture these new aspects, hence new frameworks and algorithms are desirable.

Generally, the goal of clustering is to identify meaningful subsets (called clusters) within a set of data points. The clusters are selected in such a way so that the points in the same clusters are more similar to each other than to the points in different clusters. Clustering is unsupervised because it does not make use of annotated data in order to estimate the aforementioned clusters. Instead, the clusters are identified only by using the characteristics of the data points, such as the feature vector representation of the points, or the distances between them.

There are different paradigms for clustering in the literature, among which are the following common approaches.

- Objective-based clustering, which imposes a quantitative objective and assumes that the target clustering is equal or close to the partition that optimizes the objective. For example, in the classic $k$-means clustering [81, 5], the points are partitioned into $k$ clusters and each cluster is assigned a center, and the objective function is the sum of the squared distances between the points and their centers. The goal is then to find the partitioning and cluster centers that minimize the given objective. Another example is $k$-median clustering, whose objective is the sum of the distances between the points and their centers.
- Hierarchical clustering, which builds a hierarchy (typically a tree) of clusters instead of a partition. This includes both agglomerative and divisive algorithms. The agglomerative algorithms typically begin with each point being a cluster, and then repeatedly merge the two closest clusters, where the closeness between clusters is specified by some criterion such as single-linkage, complete-linkage, or average-linkage [98, 66, 44]. The divisive algorithms typically begin with all points in one cluster, and then repeatedly split one current cluster into several clusters based on some subroutine, such as $k$-means [100] or spectral clustering [31].
- Distribution-based clustering, which models the generative distribution of the data by statistics and probabilities. For example, the mixture of Gaussians model [37] assumes the points are generated from a mixture of Gaussian distributions, and aims at recovering the parameters of the mixture.
- Graph clustering, which organizes the data on the basis of the edge structure between the points and identifies the clusters by studying the edge structure. The edges can be deterministically given by the application, such as the links between the members of a social network. The clusters can then be identified by optimizing some criterion, such as conductance [61], or cut ratio [85]. In this case, the approach is closely related to objective-based clustering. The edges can also be generated from some distribution. For example, in the planted $k$-partition model [34], the points are divided into $k$ clusters, and points in the same clusters are connected with probability $p$ while points in different clusters are connected with probability $q$, where $p, q$ are two parameters. The clusters can then be identified by estimating the distribution parameters. In this case, the approach is closely related to distribution-based clustering. In some literature on network analysis, graph clustering is also called community
detection [51, 90].
- Set system clustering, which specifies some properties the collection of clusters should satisfy and aims at finding such a collection. For example, [22] aimed at finding the collection of clusters such that for any two points $p$ and $q$ in the same cluster and any point $w$ outside the cluster, $w$ is farther away from at least one of $p$ and $q$ than $p$ is from $q$. Some other work $[67,1]$ specifies a set of axioms for partitioning the data, and studies the existence and uniqueness of such partition procedures.

This thesis considers objective-based clustering where we are given the number $k$ of clusters and the distances between the data points. First, it is one of the most frequently used and studied approaches. Second, as described in the paragraphs above, other approaches are closely related to the objective-based clustering. For example, divisive hierarchical clustering algorithms may have a subroutine that is objectivebased, and estimating the parameters in distribution-based clustering may reduce to optimizing some objective function over the sampled points. Third, although it has been well studied, there are new challenges that are not (fully) addressed by existing analysis. This thesis also considers graph clustering (named community detection in some network analysis literature) due to the recent explosion of network data and the increasing interest in understanding them.

We focus on two modern aspects of the two clustering approaches considered: stability and scalability.

First, since optimizing most natural clustering objectives is NP-hard, there has been an increasing interest in moving beyond worst-case analysis of clustering based on the stability properties of the target $[92,14,9,26,69,8]$. Such stability properties formalize what is implicitly assumed about the practically interesting instances. Bilu and Linial [26] proposed such a property called perturbation resilience, which assumes
that the optimal solution does not change under small multiplicative perturbations to the pairwise distances. The first part of this thesis provides a study of this property for center-based clustering (including $k$-median and $k$-means) and min-sum clustering, and the generalization of this property that allows small changes in the optimal solutions after perturbations.

Another trend of modern clustering tasks is to deal with large scale data, especially those collected in the distributed setting. A natural question in this setting is how to compute high quality solutions with low communication. Due to the constraint on communication, most distributed applications only ask for a good approximate solution. Ideally, there is a tradeoff between the communication and the solution quality, so that the application can choose its sweet spot between the two: with more communication, the solution quality gets better; with a suitable amount of communication, the error can be within any given small distance from that of a given centralized algorithm on the global data. The second part of the thesis provides a study of such a tradeoff, and further discusses how to reduce the communication of $k$-means clustering of high dimensional data by distributed principal component analysis.

Finally, the third part of the thesis turns to community detection on network data. We propose a model of communities based on a stability property with respect to noisy nodes in the network, and design an algorithm which provably detects all such stable communities. Furthermore, we design a local algorithm to handle large scale networks, whose running time of finding a community for a given node is independent of the size of the entire network.

### 1.1 Clustering under Perturbation Resilience

For most natural clustering objectives, finding the optimal solution to the objective function is NP-hard. As a consequence, there has been substantial work on approximation algorithms [59, 29, 23, 40, 6] with both upper and lower bounds on the approximability of these objective functions on worst case instances. Bilu and Linial [26] suggested an interesting alternative approach aimed at understanding the complexity of clustering instances which arise in practice. They argued that interesting instances should be resilient to small perturbations in the distances, and specifically defined an instance to be $\alpha$-perturbation resilient for an objective if perturbing pairwise distances by multiplicative factors in the range $[1, \alpha]$ does not change the optimum clustering under the objective. Two important questions raised are: (1) how much resilience is required so that one can develop algorithms for important clustering objectives? (2) the resilience definition requires the optimal solution to remain exactly the same after perturbation: can one succeed under weaker conditions?

In Chapter 2, we address both of these questions. First, for center-based objectives, we design a polynomial time algorithm for finding the optimal solution for instances resilient to perturbations of value $\alpha=1+\sqrt{2}$, thus beating the previously best known factor of 3 in [9]. Second, for $k$-median (which is a specific center-based objective), we consider a weaker, relaxed, and more realistic notion of perturbationresilience where we allow the optimal clustering of the perturbed instance to differ from the optimal of the original in a small $\epsilon$ fraction of the points. Compared to the original perturbation resilience assumption, this is arguably a more natural though also more difficult condition to deal with. We give positive results for this case as well, showing for $\alpha>4$ we can in polynomial time compute a $(1+O(\epsilon / \rho))$-approximation to the optimum, where $\rho$ is the fraction of the points in the smallest cluster.

We further give positive results for min-sum clustering, which is a generally much harder objective than center-based objectives. For $\alpha$-perturbation resilient min-sum
instances, we provide the first polynomial time algorithm for optimally clustering when $\alpha$ is in the order of the ratio between the sizes of the largest and smallest clusters. For $(\alpha, \epsilon)$-perturbation resilient min-sum instances with $\alpha$ in the order of the ratio between the sizes of the largest and smallest clusters and $\epsilon=\tilde{O}(\rho)$, we provide a polynomial time algorithm that outputs a clustering that is both a $(1+\tilde{O}(\epsilon / \rho))$ approximation and $\tilde{O}(\epsilon)$-close to the optimal clustering.

We additionally provide sublinear time algorithms both for the $k$-median and minsum objectives, showing algorithms that can return an implicit clustering from only access to a small random sample.

### 1.2 Distributed Clustering

Most classic clustering algorithms are designed for the centralized setting, but in recent years large scale data has become distributed over different locations. As a consequence, it has become crucial to develop clustering algorithms which are effective in the distributed setting. Several algorithms for distributed clustering have been proposed and empirically tested. Some of these algorithms [48, 101, 38] are direct adaptations of centralized algorithms which rely on statistics that are easy to compute in a distributed manner. Other algorithms [60,64] generate summaries of local data and transmit them to a central coordinator which then performs the clustering algorithm. They do not provide theoretical guarantees on the clustering quality, or the reduction in communication cost. Additionally, most of these algorithms assume that the distributed nodes can communicate with all other sites or that there is a central coordinator that communicates with all other sites.

In Chapter 3, we study the problem of distributed clustering where the data is distributed across nodes which communicate over the edges of an arbitrary connected graph. We provide algorithms for $k$-means and $k$-median clustering, which have low communication costs and provable guarantees on the clustering quality. Our technique
for reducing communication in general graphs is based on the construction of a small set of points called coresets [57], which act as a proxy for the entire data set. An $\epsilon$-coreset is a weighted set of points whose cost on any set of centers is approximately the cost of the original data on those same centers (up to a multiplicative factor $1+\epsilon$ ), thus an $\alpha$-approximate solution for the coreset is also an $\alpha(1+O(\epsilon)$ )-approximate solution for the original data. In our distributed algorithms for $k$-means and $k$ median, each node constructs a local portion of a global coreset. Communicating the total cost of local approximate solutions to each node is enough for the construction, leading to low communication cost overall. The nodes then share the local portions of the coreset, which can be done efficiently in general graphs using a message passing approach. The total communication cost then is proportional to the number of points in the coreset, which is $\tilde{O}(k d+s k)$ for constant $\epsilon$, where $d$ is the dimension of the data and $s$ is the number of the nodes in the communication graph.

The number of points in the coreset is independent of the number of points in the original data, which is useful for large scale applications. However, it is linear in the dimension of the data, leading to a communication cost quadratic in the dimension, which is not scalable to high dimensional data. We propose a distributed principal component analysis algorithm, which projects the data to a subspace of dimension $O\left(k / \epsilon^{2}\right)$ with a communication cost of $O\left(s k d / \epsilon^{2}\right)$ words. We show that its output represents the original data in the sense that any $\alpha$-approximation solution of $k$-means clustering on the projected data is an $\alpha(1+\epsilon)$-approximation solution on the original data. We can then apply the aforementioned coreset based distributed clustering algorithm, whose communication cost is now independent of the number of points in the original data and their dimension.

### 1.3 Community Detection

In analyzing a social network, it is meaningful to identify interesting subsets called communities based on the affinities between the members (nodes) of the social network. In Chapter 4, we propose a theoretical model that formalizes the collection of target communities. In our model, each node of a community falls into a subcommunity and the sub-communities within this community have active interactions with each other, while entities outside this community have fewer interactions with nodes inside. To deal with practical noise, we only require the communities to satisfy the above property after removing a few nodes from the network. This means that such communities are stable in the presence of a few exceptional or irregular outliers (such as a member connected to almost all the other members, or a member whose affinity data with others have been corrupted). Given this formalization, we then propose an efficient algorithm that detects all the communities in this model, and prove that all the communities form a tree hierarchy. Furthermore, a local algorithm is designed to handle large scale networks. Such an algorithm takes a node in the network as input, and outputs a community containing this node. Its running time depends on the size of the output community but not that of the entire network, and thus it is particularly suitable for large scale networks.

### 1.4 Summary and Bibliographic Information

The thesis is organized as follows.

- In Chapter 2 we study the objective-based clustering under a stability notion called perturbation resilience. We improve the existing bound for center-based objectives including $k$-median and $k$-means, and provide the first known bound for the min-sum objective. We further propose a generalization of perturbation resilience, and provide the first known bounds under this generalization for the $k$-median and min-sum objectives. Finally, we provide sublinear time algorithms
for the $k$-median and min-sum objectives. Part of this chapter is based on the work that appears in [19].
- In Chapter 3 we provide algorithms for $k$-means and $k$-median clustering in the distributed setting, where the data is distributed across nodes which communicate over the edges of a connected graph. We bound the communication cost and clustering quality in our algorithms. For $k$-means clustering of high dimensional data, we further provide an algorithm that first reduces the dimension by distributed PCA and then performs clustering. This chapter is mostly based on the work that appears in [78] and [79].
- In Chapter 4 we propose a model for communities over networks, and provide an algorithm that provably detects all such communities. We further propose a local algorithm for large scale networks. Part of this chapter is based on the work that appears in [20].

Other work that we have done but not included in the thesis includes the following. In [17], we study learning disjunctions in the semi-supervised and active setting, and provide efficient algorithms with nearly optimal label complexity. In [24], we study the problem of learning sparse combinations of elements distributed across a network, and propose a distributed Frank-Wolfe algorithm that solves this class of problems in a scalable and communication-efficient way. In [43], we provide an efficient algorithm for learning the spread of information over networks, based on the insight that the information influence functions in many models are special combinatorial functions called coverage functions and can be learned by convex combinations of random basis functions. In [42], we provide a novel formulation for influence maximization of multiple information items as a submodular maximization under the intersection of a matroid and multiple knapsack constraints, the special structure of which leads to an efficient algorithm with an approximation factor better than known guarantees.

## CHAPTER II

## CLUSTERING UNDER PERTURBATION RESILIENCE

As discussed in the introduction, a common approach for solving clustering problems is objective-based clustering. We are given the number $k$ of clusters and the distances between the data points, and we want to optimize various objective functions such as $k$-median or min-sum. In the $k$-median clustering approach, the goal is to partition the data into $k$ clusters $C_{i}$, giving each a center $c_{i}$, in order to minimize the sum of the distances of all data points to the centers of their cluster. In the min-sum clustering approach, the goal is to partition the data into $k$ clusters $C_{i}$ that minimize the sum of all intra-cluster pairwise distances. Yet unfortunately, for these natural clustering objectives [55, 59], finding the optimal solution to the objective function is NP-hard. As a consequence, there has been substantial work on approximation algorithms $[59,29,23,40,6]$ with both upper and lower bounds on the approximability of these objective functions on worst case instances.

Recently, Bilu and Linial [26] suggested an exciting, alternative approach aimed at understanding the complexity of clustering instances which arise in practice. Motivated by the fact that distances between data points in clustering instances are often based on a heuristic measure, they argue that interesting instances should be resilient to small perturbations in these distances. In particular, if small perturbations can cause the optimum clustering for a given objective to change drastically, then that probably is not a meaningful objective to be optimizing. Bilu and Linial [26] specifically define an instance to be $\alpha$-perturbation resilient ${ }^{1}$ for an objective $\Phi$ if perturbing pairwise distances by multiplicative factors in the range $[1, \alpha]$ does not change the

[^0]optimum clustering under $\Phi .{ }^{2}$ They consider in detail the case of Max-Cut clustering and give an efficient algorithm to recover the optimum when the instance is resilient to perturbations on the order of $\alpha=O(\sqrt{n})$.

Two important questions raised by the work of Bilu and Linial [26] are: (1) the degree of resilience needed for their algorithm to succeed is quite high: can one develop algorithms for important clustering objectives that require much less resilience? (2) the resilience definition requires the optimum solution to remain exactly the same after perturbation: can one succeed under weaker conditions? In the context of center-based clustering objectives such as $k$-median and $k$-center, [9] partially address the first of these questions and show that an algorithm based on the single-linkage heuristic can be used to find the optimal clustering for $\alpha$-perturbation-resilient instances for $\alpha=3$. They also conjecture it to be NP-hard to beat 3 and prove beating 3 is NP-hard for a related notion.

In this work, we address both questions raised by [26] and additionally improve over [9]. First, for the center-based objectives we design a polynomial time algorithm for finding the optimum solution for instances resilient to perturbations of value $\alpha=$ $1+\sqrt{2}$, thus beating the previously best known factor of 3 in [9]. Second, for $k$-median (which is a specific center-based objective), we consider a weaker, relaxed, and more realistic notion of perturbation-resilience where we allow the optimal clustering of the perturbed instance to differ from the optimal of the original in a small $\epsilon$ fraction of the points. Compared to the original perturbation resilience assumption, this is arguably a more natural though also more difficult condition to deal with. We give positive results for this case as well, showing for somewhat larger values of $\alpha$ that we can still achieve a near-optimal clustering on the given instance (see Section 1.1 below for precise results). We additionally give positive results for min-sum clustering which is

[^1]a generally much harder objective than center-based objectives. For example, the best known guarantee for min-sum clustering on worst-case instances is an $O\left(\delta^{-1} \log ^{1+\delta} n\right)$ approximation algorithm that runs in time $n^{O(1 / \delta)}$ due to Bartal et al. [23]; by contrast, the best guarantee known for $k$-median is factor $1+\sqrt{3}+\epsilon[76]$.

Our results are achieved by carefully deriving structural properties of perturbationresilience. At a high level, all the algorithms we introduce work by first running appropriate linkage procedures to produce a hierarchical clustering, and then running dynamic programming to retrieve the best $k$-clustering present in the tree. To ensure that (under perturbation resilient instances) the hierarchy output in the first step has a pruning of low cost, we derive new linkage procedures (closure linkage and robust average linkage) which are of independent interest. While the overall analysis is quite involved, the clustering algorithms we devise are simple and robust. This simplicity and robustness allow us to show how our algorithms can be made sublinear time by returning an implicit clustering from only a small random sample of the input.

From a learning theory perspective, the resilience parameter, $\alpha$, can also be seen as an analog to a margin for clustering. In supervised learning, the margin of a data point is the distance, after scaling, between the data point and the decision boundary of its classifier, and many algorithms have stronger guarantees when the smallest margin over the entire data set is sufficiently large [96, 102]. The $\alpha$ parameter, similarly controls the magnitude of the perturbation the data can withstand before being clustered differently, which is, in essence, the data's distance to the decision boundary for the given clustering objective. Hence, perturbation resilience is also a natural and interesting assumption to study from a learning theory perspective.

Our Results: In this work, we greatly advance the line of work of [26] by solving a number of important problems of clustering perturbation-resilient instances under metric center-based and min-sum objectives.

In Section 2.2 we improve on the bounds of [9] for $\alpha$-perturbation resilient instances
for center-based objectives, giving an algorithm that efficiently ${ }^{3}$ finds the optimum clustering for $\alpha=1+\sqrt{2}$. Most of the frequently used center-based objectives, such as $k$-median, are NP-hard to even approximate, yet we can recover the exact solution for perturbation resilient instances. Our algorithm is based on a new linkage procedure using a new notion of distance (closure distance) between sets that may be of independent interest.

In Section 2.3 we consider the more challenging and more general notion of $(\alpha, \epsilon)$ perturbation resilience for $k$-median, where we allow the optimal solution after perturbation to be $\epsilon$-close to the original. We provide an efficient algorithm which for $\alpha>4$ produces $(1+O(\epsilon / \rho))$-approximation to the optimum, where $\rho$ is the fraction of the points in the smallest cluster. The key structural property we derive and exploit is that, except for $\epsilon n$ bad points, most points are $\alpha$ times closer to their own center than to any other center. To eliminate the noise introduced by the bad points, we carefully partition the points into a list of sufficiently large blobs, each of which contains only good points from one optimal cluster. This then allows us to construct a tree on the blobs with a low-cost pruning that is a good approximation to the optimum.

In Section 2.4 we provide the first efficient algorithm for optimally clustering $\alpha$ perturbation resilient min-sum instances. Our algorithm is based on an appropriate modification of average linkage that exploits the structure of min-sum perturbation resilient instances.

In Section 2.5, we show that for $(\alpha, \epsilon)$-perturbation resilient min-sum instances with $\alpha$ in the order of the ratio between the sizes of the largest and smallest clusters and $\epsilon=\tilde{O}(\rho)$, there exists a polynomial time algorithm that outputs a clustering that is both a $(1+\tilde{O}(\epsilon / \rho))$-approximation and $\tilde{O}(\epsilon)$-close to the optimal clustering.

We also provide sublinear-time algorithms both for the $k$-median and min-sum

[^2]objectives (Sections 2.3.3 and 2.4.1), showing algorithms that can return an implicit clustering from only access to a small random sample.

Related Work: A subsequent work of [26] by Bilu, Daniely, Linial and Saks [27] studied the Max-Cut problem under Bilu and Linial stability, and showed how to solve in polynomial time $(1+\epsilon)$-stable instances of metric and dense Max-Cut, and $\Omega(\sqrt{n})$-stable instances of general Max-Cut. The later bound is further improved by Makarychev, Makarychev and Vijayaraghavan [84]. They proposed a polynomial time exact algorithm for $\Omega(\sqrt{\log n} \log \log n)$-stable Max-Cut instances based on semidefinite programming. They also proved that for Max $k$-Cut with $k \geq 3$, there is no polynomial-time algorithm that solves $\infty$-stable instances of Max $k$-Cut unless $\mathbf{N P}=\mathbf{R P}$. Here an instance is $\infty$-stable if it is $\alpha$-stable for every $\alpha$.

In the context of objective-based clustering, several recent papers have shown how to exploit other notions of stability for overcoming the existing hardness results on worst case instances. The ORSS stability notion of Ostrovsky, Rabani, Schulman and Swamy $[92,9]$ assumes that the cost of the optimal $k$-means solution is small compared to the cost of the optimal $(k-1)$-means solution. The BBG approximation stability condition of Balcan, Blum and Gupta [14] assumes that every nearly optimal solution is close to the target clustering. Awasthi, Sheffet and Blum [8] proposed a stability condition called weak-deletion stability, and showed that it is implied by both the ORSS stability and the BBG stability. Kumar and Kannan [69] proposed a proximity condition which assumes that in the target clustering, most data points satisfy that they are closer to their center than to any other center by an additive factor in the order of the maximal standard variance of their clusters in any direction. Their results are improved by Awasthi and Sheffet [10], which proposed a weaker version of the proximity condition called center separation, and designed algorithms achieving stronger guarantees under this weaker condition.

Several recent papers have shown how to exploit the structure of perturbation
resilient instances in order to obtain better approximation guarantees (than those possible on worst case instances) for other difficult optimization problems. These include the game theoretic problem of finding Nash equilibria [ 15,80 ] and the classic traveling salesman problem [86].

### 2.1 Preliminaries

In a clustering instance, we are given a set $P$ of $n$ points in a finite metric space, and we denote $d: P \times P \rightarrow \mathbf{R}_{\geq 0}$ as the distance function. $\Phi$ denotes the objective function over a partition of $P$ into $k<n$ clusters which we want to optimize over the metric, ie. $\Phi$ assigns a score to every clustering. The optimal clustering with respect to $\Phi$ is denoted as $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{k}\right\}$, and its cost is denoted as $\mathcal{O P} \mathcal{T}$. The core concept we study in this paper is the perturbation resilience notion introduced by [26]. Formally:

Definition 1. A clustering instance $(P, d)$ is $\alpha$-perturbation resilient to a given objective $\Phi$ if for any function $d^{\prime}: P \times P \rightarrow \mathbf{R}_{\geq 0}$ s.t. $\forall p, q \in P, d(p, q) \leq d^{\prime}(p, q) \leq \alpha d(p, q)$, there is a unique optimal clustering $\mathcal{C}^{\prime}$ for $\Phi$ under $d^{\prime}$ and this clustering is equal to the optimal clustering $\mathcal{C}$ for $\Phi$ under $d$.

In this paper, we focus on the center-based and min-sum objectives. For the center-based objectives, we consider separable center-based objectives defined by [9].

Definition 2. A clustering objective is center-based if the optimal solution can be defined by $k$ points $c_{1}, \cdots, c_{k}$ in the metric space called centers such that every data point is assigned to its nearest center. Such a clustering objective is separable if it furthermore satisfies the following two conditions:

- The objective function value of a given clustering is either a (weighted) sum or the maximum of the individual cluster scores.
- Given a proposed single cluster, its score can be computed in polynomial time.

One particular center-based objective is the $k$-median objective. We partition $P$ into $k$ disjoint subsets $\mathcal{P}=\left\{P_{1}, P_{2}, \ldots, P_{k}\right\}$ and assign a set of centers $\mathbf{p}=$ $\left\{p_{1}, p_{2}, \ldots, p_{k}\right\} \subseteq P$ for the subsets. The goal is to minimize $\Phi(\mathcal{P}, \mathbf{p})=\sum_{i=1}^{k} \sum_{p \in P_{i}} d\left(p, p_{i}\right)$. The centers in the optimal clustering are denoted as $\mathbf{c}=\left\{c_{1}, \ldots, c_{k}\right\}$. Clearly, in an optimal solution, each point is assigned to its nearest center. In such cases, the objective is denoted as $\Phi(\mathbf{c})$.

For the min-sum objective, we partition $P$ into $k$ disjoint subsets $\mathcal{P}=\left\{P_{1}, P_{2}, \ldots, P_{k}\right\}$, and the goal is to minimize $\Phi(\mathcal{P})=\sum_{i=1}^{k} \sum_{p, q \in P_{i}} d(p, q)$. Note that we sometimes denote $\Phi$ as $\Phi_{P}$ in the case where the distinction is necessary, such as in Section 2.3.3.

In Section 2.3 we consider a generalization of Definition 1 where we allow a small difference between the original optimum and the new optimum after perturbation. Formally:

Definition 3. Let $\mathcal{C}$ be the optimal $k$-clustering and $\mathcal{C}^{\prime}$ be another $k$-clustering of $a$ set of $n$ points. We say $\mathcal{C}^{\prime}$ is $\epsilon$-close to $\mathcal{C}$ if $\min _{\sigma \in \mathcal{S}_{k}} \sum_{i=1}^{k}\left|C_{i} \backslash C_{\sigma(i)}^{\prime}\right| \leq \epsilon n$, where $\sigma$ is a matching between indices of clusters of $\mathcal{C}^{\prime}$ and those of $\mathcal{C}$.

Definition 4. A clustering instance $(P, d)$ is $(\alpha, \epsilon)$-perturbation resilient to a given objective $\Phi$ if for any function $d^{\prime}: P \times P \rightarrow \mathbf{R}_{\geq 0}$ s.t. $\forall p, q \in P, d(p, q) \leq d^{\prime}(p, q) \leq$ $\alpha d(p, q)$, the optimal clustering $\mathcal{C}^{\prime}$ for $\Phi$ under $d^{\prime}$ is $\epsilon$-close to the optimal clustering $\mathcal{C}$ for $\Phi$ under $d$.

For $A, B \subseteq P$ we define $d(A, B):=\sum_{p \in A} \sum_{q \in B} d(p, q)$ and $d(p, B):=d(\{p\}, B)$. Also, we define $d_{\mathrm{a}}(A, B):=d(A, B) /(|A||B|)$ and $d_{\mathrm{a}}(p, B):=d_{\mathrm{a}}(\{p\}, B)$. For simplicity, we will sometimes assume that $\min _{i}\left|C_{i}\right|$ is known. (Otherwise, we can simply search over the $n$ possible different values.)

Finally, if $A \cap B=\emptyset$, we sometimes write $A \cup B$ as $A+B$ to emphasize that they are disjoint.

## 2.2 $\alpha$-Perturbation Resilience for Center-Based Objectives

In this section we show that, for $\alpha \geq 1+\sqrt{2}$, if the clustering instance is $\alpha$ perturbation resilient for center-based objectives, then we can in polynomial time find the optimal clustering. This improves on the $\alpha \geq 3$ bound of [9] and stands in sharp contrast to the NP-Hardness results on worst-case instances. Our algorithm succeeds for an even weaker property, the $\alpha$-center proximity, introduced in [9].

Definition 5. A clustering instance ( $P, d$ ) satisfies the $\alpha$-center proximity property if for any optimal cluster $C_{i} \in \mathcal{C}$ with center $c_{i}, C_{j} \in \mathcal{C}(j \neq i)$ with center $c_{j}$, any point $p \in C_{i}$ satisfies $\alpha d\left(p, c_{i}\right)<d\left(p, c_{j}\right)$.

Lemma 1. Any clustering instance that is $\alpha$-perturbation resilient to center-based objectives also satisfies the $\alpha$-center proximity.

The proof follows easily by constructing a specific perturbation that blows up all the pairwise distances within cluster $C_{i}$ by a factor of $\alpha$. By $\alpha$-perturbation resilience, the optimal clustering remains the same after this perturbation. This then implies the desired result. The full proof appears in [9]. In the remainder of this section, we prove our results for $\alpha$-center proximity, but because it is a weaker condition, our upper bounds also hold for $\alpha$-perturbation resilience.

We begin with some key properties of $\alpha$-center proximity instances.

Lemma 2. For any points $p \in C_{i}$ and $q \in C_{j}(j \neq i)$ in the optimal clustering of an $\alpha$-center proximity instance, when $\alpha \geq 1+\sqrt{2}$, we have: (1) $d\left(c_{i}, q\right)>d\left(c_{i}, p\right)$, (2) $d\left(p, c_{i}\right)<d(p, q)$.

Proof. (1) Lemma 1 gives us that $d\left(q, c_{i}\right)>\alpha d\left(q, c_{j}\right)$. By the triangle inequality, we have $d\left(c_{i}, c_{j}\right) \leq d\left(q, c_{j}\right)+d\left(q, c_{i}\right)<\left(1+\frac{1}{\alpha}\right) d\left(q, c_{i}\right)$. On the other hand, $d\left(p, c_{j}\right)>$ $\alpha d\left(p, c_{i}\right)$ and therefore $d\left(c_{i}, c_{j}\right) \geq d\left(p, c_{j}\right)-d\left(p, c_{i}\right)>(\alpha-1) d\left(p, c_{i}\right)$. Combining these inequalities, we get (1).
(2) It suffices to prove $d(p, q)>(\alpha-1) \max \left\{d\left(p, c_{i}\right), d\left(q, c_{j}\right)\right\}$. The proof first appears in [9], and we include it for completeness. Without loss of generality, we can assume that $d\left(p, c_{i}\right) \geq d\left(q, c_{j}\right)$. By triangle inequality we have $d(p, q) \geq d\left(p, c_{j}\right)-$ $d\left(q, c_{j}\right)$. From Lemma 1 we have $d\left(p, c_{j}\right)>\alpha d\left(p, c_{i}\right)$. Hence $d(p, q)>\alpha d\left(p, c_{i}\right)-$ $d\left(q, c_{j}\right) \geq(\alpha-1) d\left(p, c_{i}\right) \geq(\alpha-1) d\left(q, c_{j}\right)$.

Lemma 2 implies that for any optimal cluster $C_{i}$, the ball of radius $\max _{p \in C_{i}} d\left(c_{i}, p\right)$ around the center $c_{i}$ contains only points from $C_{i}$, and moreover, points inside the ball are each closer to the center than to any point outside the ball. Inspired by this structural property, we define the notion of closure distance between two sets as the radius of the minimum ball that covers the sets and has some margin from points outside the ball. We show that any (strict) subset of an optimal cluster has smaller closure distance to another subset in the same cluster than to any subset of other clusters or to unions of other clusters. Using this, we will be able to define an appropriate linkage procedure that, when applied to the data, produces a tree on subsets that will all be laminar with respect to the clusters in the optimal solution. This will then allow us to extract the optimal solution using dynamic programming applied to the tree.

We now define the notion of closure distance and then present our algorithm for $\alpha$-center proximity instances (Algorithm 1$)$. Let $\mathbb{B}(p, r)=\{q: d(q, p) \leq r\}$.

Definition 6. The closure distance $d_{P}\left(A, A^{\prime}\right)$ between two disjoint non-empty subsets $A$ and $A^{\prime}$ of point set $P$ is the minimum $d \geq 0$ such that there is a point $c \in A \cup A^{\prime}$ satisfying the following requirements:
(1) coverage: the ball $\mathbb{B}(c, d)$ covers $A$ and $A^{\prime}$, i.e. $A \cup A^{\prime} \subseteq \mathbb{B}(c, d)$;
(2) margin: points inside $\mathbb{B}(c, d)$ are closer to the center $c$ than to points outside, i.e. $\forall p \in \mathbb{B}(c, d), q \notin \mathbb{B}(c, d)$, we have $d(c, p)<d(p, q)$.

Note that for any $A$ and $A^{\prime}$, we have $d_{P}\left(A, A^{\prime}\right)=d_{P}\left(A^{\prime}, A\right) \leq \max _{p, q \in P} d(p, q)$.


Figure 1: The closure distance.

Furthermore, $d_{P}\left(A, A^{\prime}\right)$ can be computed in polynomial time.

```
Algorithm 1 Center-based objectives, \(\alpha\) perturbation resilience
Input: Data set \(P\), distance function \(d(\cdot, \cdot)\) on \(P\).
    1: Begin with \(n\) singleton clusters.
    2: Repeat till only one cluster remains:
    merge clusters \(C, C^{\prime}\) which minimize \(d_{P}\left(C, C^{\prime}\right)\).
    3: Let \(\mathcal{T}\) be the tree with single points as leaves and internal nodes corresponding
    to the merges performed.
    4: Run dynamic programming on \(\mathcal{T}\) to get the minimum cost pruning \(\tilde{\mathcal{C}}\).
Output: Clustering \(\tilde{\mathcal{C}}\).
```

Theorem 1. For $(1+\sqrt{2})$-center proximity instances, Algorithm 1 outputs the optimal clustering in polynomial time.

The proof follows immediately from the following key property of Algorithm 1. The details of dynamic programming are presented in Appendix A.1, and an efficient implementation of the algorithm is presented in Appendix A.2.

Theorem 2. For $(1+\sqrt{2})$-center proximity instances, Algorithm 1 constructs a binary tree $\mathcal{T}$ such that the optimal clustering is a pruning of this tree.

Proof. We prove correctness by induction. In particular, assume that our current clustering is laminar with respect to the optimal clustering - that is, for each cluster $A$ in our current clustering and each $C$ in the optimal clustering, we have either $A \subseteq C$, or $C \subseteq A$ or $A \cap C=\varnothing$. This is clearly true at the start. To prove that the merge steps keep the laminarity, we need to show the following: if $A$ is a
strict subset of an optimal cluster $C_{i}, A^{\prime}$ is a subset of another optimal cluster or the union of one or more other clusters, then there exists $B$ from $C_{i} \backslash A$, such that $d_{P}(A, B)<d_{P}\left(A, A^{\prime}\right)=d_{P}\left(A^{\prime}, A\right)$.

We first prove that there is a cluster $B \subseteq C_{i} \backslash A$ in the current cluster list such that $d_{P}(A, B) \leq \tilde{d}:=\max _{p \in C_{i}} d\left(c_{i}, p\right)$. There are two cases. First, if $c_{i} \notin A$, then define $B$ to be the cluster in the current cluster list that contains $c_{i}$. By induction, $B \subseteq C_{i}$ and thus $B \subseteq C_{i} \backslash A$. Then we have $d_{P}(B, A) \leq \tilde{d}$ since there is $c_{i} \in B$, and (1) for any $p \in A \cup B, d\left(c_{i}, p\right) \leq \tilde{d}$, (2) for any $p \in P$ satisfying $d\left(c_{i}, p\right) \leq \tilde{d}$, and any $q \in P$ satisfying $d\left(c_{i}, q\right)>\tilde{d}$, by Lemma 2 we know $p \in C_{i}$ and $q \notin C_{i}$, and thus $d\left(c_{i}, p\right)<d(p, q)$. In the second case when $c_{i} \in A$, we pick any $B \subseteq C_{i} \backslash A$ and a similar argument gives $d_{P}(A, B) \leq \tilde{d}$.

case 2: $c \in A^{\prime}$


Figure 2: Comparing $d$ and $d_{P}\left(A, A^{\prime}\right)$ in closure linkage.

As a second step, we need to show that $d<\hat{d}:=d_{P}\left(A, A^{\prime}\right)$. There are two cases: the center for $d_{P}\left(A, A^{\prime}\right)$ is in $A$ or in $A^{\prime}$. See Figure 2 for an illustration. In the first case, there is a point $c \in A$ such that $c$ and $\hat{d}$ satisfy the requirements of the closure distance. Pick a point $q \in A^{\prime}$, and define $C_{j}$ to be the cluster in the optimal
clustering that contains $q$. As $d(c, q) \leq \hat{d}$, and by Lemma 2 we have $d\left(c_{j}, q\right)<d(c, q)$, then $d\left(c_{j}, c\right) \leq \hat{d}$ (otherwise it violates the second requirement of closure distance). Suppose $p=\arg \max _{p^{\prime} \in C_{i}} d\left(c_{i}, p^{\prime}\right)$. Then we have $\tilde{d}=d\left(p, c_{i}\right)<d\left(p, c_{j}\right) / \alpha \leq(\tilde{d}+$ $\left.d\left(c_{i}, c\right)+d\left(c, c_{j}\right)\right) / \alpha$ where the first inequality comes from Lemma 1 and the second from the triangle inequality. Since $d\left(c_{i}, c\right)<d\left(c, c_{j}\right) / \alpha$, we can combine the above inequalities and compare $\tilde{d}$ and $d\left(c, c_{j}\right)$, and when $\alpha \geq 1+\sqrt{2}$ we have $\tilde{d}<d\left(c, c_{j}\right) \leq \hat{d}$.

Now consider the second case, when there is a point $c \in A^{\prime}$ such that $c$ and $\hat{d}$ satisfy the requirements in the definition of the closure distance. Select an arbitrary point $q \in A$. We have $\hat{d} \geq d(c, q)$ from the first requirement, and $d(c, q)>d\left(c_{i}, q\right)$ by Lemma 2. Then from the second requirement of closure distance $d\left(c_{i}, c\right) \leq \hat{d}$. And by Lemma $2, \tilde{d}=d\left(c_{i}, p\right)<d\left(c_{i}, c\right)$, we have $\tilde{d}<d\left(c_{i}, c\right) \leq \hat{d}$.

Note: Our factor of $\alpha=1+\sqrt{2}$ beats the NP-hardness lower bound of $\alpha=3$ of [9] for center-proximity instances. The reason is that the lower bound of [9] requires the addition of Steiner points that can act as centers but are not part of the data to be clustered (though the upper bound of [9] does not allow such Steiner points). One can also show a lower bound for center-proximity instances without Steiner points. In particular one can show that for any $\epsilon>0$, the problem of solving $(2-\epsilon)$-center proximity $k$-median instances is NP-hard [95].

## $2.3(\alpha, \epsilon)$-Perturbation Resilience for $k$-Median

In this section we consider a natural relaxation of the $\alpha$-perturbation resilience, the $(\alpha, \epsilon)$-perturbation resilience property, that requires the optimum after perturbation of up to a multiplicative factor $\alpha$ to be $\epsilon$-close to the original (one should think of $\epsilon$ as sub-constant). We show that if the instance is ( $\alpha, \epsilon$ )-perturbation resilient, with $\alpha>4$ and $\epsilon=O\left(\epsilon^{\prime} \rho\right)$ where $\rho$ is the fraction of the points in the smallest cluster, then we can in polynomial time output a clustering that provides a $\left(1+\epsilon^{\prime}\right)$-approximation to the optimum. Thus this improves over the best worst-case approximation guarantees
known [76] when $\epsilon^{\prime} \leq \sqrt{3}$ and also beats the lower bound of $(1+1 / e)$ on the best approximation achievable on worst case instances for the metric $k$-median objective $[55,59]$ when $\epsilon^{\prime} \leq 1 / e$.

The key idea is to understand and leverage the structure implied by $(\alpha, \epsilon)$-perturbation resilience. We show that perturbation resilience implies that there exists only a small fraction of points that are bad in the sense that their distance to their own center is not $\alpha$ times smaller than their distance to any other centers in the optimal solution. We then use this bounded number of bad points in our clustering algorithm.

### 2.3.1 Structure of $(\alpha, \epsilon)$-Perturbation Resilient Instances

To understand $(\alpha, \epsilon)$-perturbation resilience, we need to consider the difference between the optimal clustering $\mathcal{C}$ under $d$ and the optimal clustering $\mathcal{C}^{\prime}$ under $d^{\prime}$, defined as $\min _{\sigma \in \mathcal{S}_{k}} \sum_{i=1}^{k}\left|C_{i} \backslash C_{\sigma(i)}^{\prime}\right|$. Without loss of generality, we assume in this subsection that $\mathcal{C}^{\prime}$ is indexed so that the $\operatorname{argmin} \sigma$ is the identity, and the distance between $\mathcal{C}$ and $\mathcal{C}^{\prime}$ is $\sum_{i=1}^{k}\left|C_{i} \backslash C_{i}^{\prime}\right|$. We denote by $c_{i}^{\prime}$ the center of $C_{i}^{\prime}$.

In the following we call a point good if it is $\alpha$ times closer to its own center than to any other center in the optimal clustering; otherwise we call it bad. Let $B_{i}$ be the set of bad points in $C_{i}$. That is, $B_{i}=\left\{p \in C_{i}: \exists j \neq i, \alpha d\left(c_{i}, p\right)>d\left(c_{j}, p\right)\right\}$. Let $G_{i}=C_{i} \backslash B_{i}$ be the good points in cluster $C_{i}$. Let $B=\cup_{i} B_{i}$ and $G=\cup_{i} G_{i}$. We show that under perturbation resilience we do not have too many bad points. Formally:

Theorem 3. Suppose the clustering instance is $(\alpha, \epsilon)$-perturbation resilient for the $k$-median objective and $\min _{i}\left|C_{i}\right|>\left(3+\frac{2 \alpha}{\alpha-1}\right) \epsilon n+9 \alpha$. Then $|B| \leq \epsilon n$.

The main idea is to construct a specific perturbation that forces certain selected bad points to move from their original optimal clusters. Then the $(\alpha, \epsilon)$-perturbation resilience leads to a bound on the number of selected bad points, which can also be proved to be a bound on all the bad points. The selected bad points $\hat{B}_{i}$ in cluster $i$ are defined by arbitrarily selecting $\min \left(\epsilon n+1,\left|B_{i}\right|\right)$ points from $B_{i}$. Let $c(p)$ denote the


Figure 3: Different types of points when bounding the number of bad points for $(\alpha, \epsilon)$-perturbation resilient $k$-median instances.
second nearest center for $p \in \hat{B}_{i}$ and the nearest center for $p \notin \hat{B}_{i}$. The perturbation we consider blows up all distances by a factor of $\alpha$ except for those distances between $p$ and $c(p)$. Suppose after perturbation, the optimal cluster $C_{i}^{\prime}$ is obtained by adding point set $A_{i}$ and removing point set $M_{i}$ from $C_{i}$, i.e. $A_{i}=C_{i}^{\prime} \backslash C_{i}, M_{i}=C_{i} \backslash C_{i}^{\prime}$. $C_{i} \cap C_{i}^{\prime}$ can be divided into two parts: $W_{i}=\left(C_{i} \cap C_{i}^{\prime}\right) \backslash \hat{B}_{i}$ and $V_{i}=\left(C_{i} \cap C_{i}^{\prime}\right) \cap \hat{B}_{i}$. Then we have $C_{i}=W_{i}+V_{i}+M_{i}, C_{i}^{\prime}=W_{i}+V_{i}+A_{i}$. See Figure 3 for an illustration.

The key challenge in proving a bound on the selected bad points is to show that $c_{i}^{\prime}=c_{i}$ for all $i$. This means the optimal centers do not change after the perturbation. Then in the optimum under $d^{\prime}$ each point $p$ is assigned to the center $c(p)$, and therefore the selected bad points will move from their original optimal clusters. By $(\alpha, \epsilon)$ perturbation resilience, we get an upper bound on the number of selected bad points.

At a high level, we prove that $c_{i}=c_{i}^{\prime}$ for all $i$ as follows. We first show that for each cluster, its new center is close to its old center, roughly speaking since the new and old cluster have a lot in common (Claim 1). We then show if $c_{i} \neq c_{i}^{\prime}$ for some $i$, then the weighted sum of the distances $\sum_{1 \leq i \leq k}\left|C_{i}\right| d\left(c_{i}, c_{i}^{\prime}\right)$ should be large (Claim 2). However, this contradicts Claim 1, so the centers do not move after the perturbation.

In the following, we will prove the two claims that imply $c_{i}^{\prime}=c_{i}(\forall i)$ and then use them to prove the theorem. The proofs make frequent use of the translation from $d^{\prime}$ to $d$, which is summarized in Fact 1 .

Fact 1. Suppose the clustering instance is $(\alpha, \epsilon)$-perturbation resilient and $\min _{i}\left|C_{i}\right|>$
$\left(\frac{2}{\alpha-1}+3\right) \epsilon n+1$. If $c_{i}^{\prime} \neq c_{i}$, then we have
$d^{\prime}\left(c_{i}^{\prime}, W_{i}\right) \geq \alpha d\left(c_{i}^{\prime}, W_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right), \quad d^{\prime}\left(c_{i}, W_{i}\right)=d\left(c_{i}, W_{i}\right)$,
$d^{\prime}\left(c_{i}^{\prime}, V_{i}\right)=\alpha d\left(c_{i}^{\prime}, V_{i}\right), \quad d^{\prime}\left(c_{i}, V_{i}\right)=\alpha d\left(c_{i}, V_{i}\right)$,
$d^{\prime}\left(c_{i}^{\prime}, A_{i}\right) \geq \alpha d\left(c_{i}^{\prime}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right), \quad d^{\prime}\left(c_{i}, A_{i}\right) \leq \alpha d\left(c_{i}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)+\alpha(1+\alpha) d\left(c_{i}^{\prime}, c_{i}\right)$.
Proof Sketch. If $d^{\prime}\left(c_{i}^{\prime}, C\right)$ does not involve distances between $p$ and $c(p)$ for any $p \in P$, then $d^{\prime}\left(c_{i}^{\prime}, C\right)=\alpha d\left(c_{i}^{\prime}, C\right)$. To apply this idea, we first need to show that $c_{i}^{\prime} \neq c_{j}(\forall j \neq i)$. Intuitively, if $c_{i}^{\prime}=c_{j}$, then under $d^{\prime}$, points in $W_{j}$ should be closer to $c_{j}^{\prime}$ than to $c_{i}^{\prime}=c_{j}$. So under $d$, these points are $\alpha$ time closer to $c_{j}^{\prime}$ than to $c_{j}$, which means the distance between $c_{j}$ and $c_{j}^{\prime}$ is not so large compared to the average distance between $c_{j}$ and $W_{j}$ by the triangle inequality. On the other hand, it also means $c_{j}$ has $(1-1 / \alpha) d\left(c_{j}, W_{j}\right)$ more cost than $c_{j}^{\prime}$ on $W_{j}$. Then $c_{j}^{\prime}$ should have at least $(1-1 / \alpha) d\left(c_{j}, W_{j}\right)$ more cost on $C_{j} \backslash W_{j}$. By the triangle inequality, the distance between $c_{j}$ and $c_{j}^{\prime}$ is much larger than the average distance between $c_{j}$ and $W_{j}$, which is contradictory. Therefore, $c_{i}^{\prime} \neq c_{j}$.

Then we only need to check if $c\left(c_{i}^{\prime}\right) \in C$ when translating $d^{\prime}\left(c_{i}^{\prime}, C\right)$ to $d\left(c_{i}^{\prime}, C\right)$ by a case-by-case study. The same idea can be applied to $d^{\prime}\left(c_{i}, C\right)$. The formal proof is presented in Appendix A.3.

Claim 1. For each i, $d\left(c_{i}, W_{i}\right) \geq \frac{\alpha+2}{\alpha+1} \frac{\left|C_{i}\right|}{3} d\left(c_{i}, c_{i}^{\prime}\right)$.

Proof. The key idea is that under $d^{\prime}, c_{i}^{\prime}$ is the optimal center for $C_{i}^{\prime}$, so it has no more cost than $c_{i}$ on $C_{i}^{\prime}$. Since $V_{i}$ and $A_{i}$ are small compared to $W_{i}, c_{i}^{\prime}$ cannot save much cost on $V_{i}$ and $A_{i}$, so it cannot have much more cost on $W_{i}$ than $c_{i}$. Then $c_{i}^{\prime}$ is close to $W_{i}$ (compared to the distance between $c_{i}$ and $W_{i}$ ). By triangle inequality, $c_{i}^{\prime}$ is close to $c_{i}$.

Formally, if $c_{i}^{\prime}=c_{i}, d\left(c_{i}^{\prime}, c_{i}\right)=0$, which immediately implies the bound. Otherwise, we need to use the fact that $c_{i}^{\prime}$ has smaller cost than $c_{i}$ on $C_{i}^{\prime}$ under $d^{\prime}: d^{\prime}\left(c_{i}^{\prime}, C_{i}^{\prime}\right) \leq$ $d^{\prime}\left(c_{i}, C_{i}^{\prime}\right)$. We divide $C_{i}^{\prime}$ into three parts $W_{i}, V_{i}$ and $A_{i}$, and move terms on $W_{i}$ to one
side ( the cost more than $c_{i}$ on $W_{i}$ ), the rest terms to another side (the cost saved on $V_{i}$ and $A_{i}$ ):

$$
d^{\prime}\left(c_{i}^{\prime}, W_{i}\right)-d^{\prime}\left(c_{i}, W_{i}\right) \leq d^{\prime}\left(c_{i}, A_{i}\right)-d^{\prime}\left(c_{i}^{\prime}, A_{i}\right)+d^{\prime}\left(c_{i}, V_{i}\right)-d^{\prime}\left(c_{i}^{\prime}, V_{i}\right)
$$

Translating $d^{\prime}$ to $d$ by Fact 1, we have

$$
\begin{aligned}
& \alpha d\left(c_{i}^{\prime}, W_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)-d\left(c_{i}, W_{i}\right) \\
\leq & \alpha d\left(c_{i}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)+\alpha(1+\alpha) d\left(c_{i}^{\prime}, c_{i}\right)-\alpha d\left(c_{i}^{\prime}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right) \\
& +\alpha d\left(c_{i}, V_{i}\right)-\alpha d\left(c_{i}^{\prime}, V_{i}\right) .
\end{aligned}
$$

By the triangle inequality,

$$
\alpha d\left(c_{i}, c_{i}^{\prime}\right)\left(\left|W_{i}\right|-1\right)-(\alpha+1) d\left(c_{i}, W_{i}\right) \leq \alpha d\left(c_{i}, c_{i}^{\prime}\right)\left[\left|V_{i}\right|+\left|A_{i}\right|+(1+\alpha)\right]
$$

which implies the desired result.

Claim 2. Let $I_{i}=1$ if $c_{i} \neq c_{i}^{\prime}$ and $I_{i}=0$ otherwise. Then $\sum_{1 \leq i \leq k} I_{i} d\left(c_{i}, W_{i}\right) \leq$ $\sum_{1 \leq i \leq k} \frac{\left|C_{i}\right|}{3} d\left(c_{i}, c_{i}^{\prime}\right)$.

Proof. The key idea is that the clustering that under $d^{\prime}$ assigns points in $W_{i}$ and $A_{i}$ to $c_{i}$ and points $p$ in $V_{i}$ to $c(p)$, saves as much cost as $d^{\prime}\left(c_{i}^{\prime}, W_{i}\right)-d^{\prime}\left(c_{i}, W_{i}\right) \approx$ $(\alpha-1) d\left(c_{i}, W_{i}\right)$ on $W_{i}$ compared to the optimal clustering $\left\{C_{i}^{\prime}\right\}$ under $d^{\prime}$, if $c_{i}^{\prime} \neq c_{i}$. Then $\left\{C_{i}^{\prime}\right\}$ must save this cost on other parts of points. So $\left\{c_{i}^{\prime}\right\}$ should be near these points and $\left\{c_{i}\right\}$ should be far away. By the triangle inequality, the weighted sum of the distances between $\left\{c_{i}^{\prime}\right\}$ and $\left\{c_{i}\right\}$ should be large.

Formally, we have the following inequality from the fact that $\left\{c_{i}^{\prime}\right\}$ are the optimal centers under $d^{\prime}$, thus have no more cost than the clustering that under $d^{\prime}$ assigns points in $W_{i} \cup A_{i}$ to $c_{i}$ and points $p$ in $V_{i}$ to $c(p)$ :

$$
\sum_{i=1}^{k} d^{\prime}\left(c_{i}^{\prime}, C_{i}^{\prime}\right) \leq \sum_{i=1}^{k}\left[d^{\prime}\left(c_{i}, C_{i}^{\prime} \backslash V_{i}\right)+\sum_{p \in V_{i}} d^{\prime}(c(p), p)\right]
$$

We divide $C_{i}^{\prime}$ into $A_{i}, V_{i}$ and $W_{i}$, and divide $C_{i}^{\prime} \backslash V_{i}$ into $A_{i}$ and $W_{i}$. Then we move terms on $W_{i}$ to one side (the cost more than $c_{i}$ on $W_{i}$ ), and move the rest terms to another side (the cost saved on $V_{i}$ and $A_{i}$ ):
$\sum_{i=1}^{k}\left[d^{\prime}\left(c_{i}^{\prime}, W_{i}\right)-d^{\prime}\left(c_{i}, W_{i}\right)\right] \leq \sum_{i=1}^{k}\left[d^{\prime}\left(c_{i}, A_{i}\right)-d^{\prime}\left(c_{i}^{\prime}, A_{i}\right)+\sum_{p \in V_{i}} d^{\prime}(c(p), p)-d^{\prime}\left(c_{i}^{\prime}, V_{i}\right)\right]$.
To estimate $d^{\prime}\left(c_{i}^{\prime}, W_{i}\right)-d^{\prime}\left(c_{i}, W_{i}\right)$, note that $d^{\prime}\left(c_{i}^{\prime}, W_{i}\right)-d^{\prime}\left(c_{i}, W_{i}\right) \approx \alpha d\left(c_{i}^{\prime}, W_{i}\right)-$ $d\left(c_{i}, W_{i}\right)$ by Fact 1 , so it suffices to show that $d\left(c_{i}, W_{i}\right)$ is not much larger than $d\left(c_{i}^{\prime}, W_{i}\right)$. This follows from the fact that $\left\{c_{i}\right\}$ are the optimal centers under $d$. Formally, $d\left(c_{i}, C_{i}\right) \leq d\left(c_{i}^{\prime}, C_{i}\right)$, which leads to

$$
\begin{equation*}
d\left(c_{i}, W_{i}\right)-d\left(c_{i}^{\prime}, W_{i}\right) \leq d\left(c_{i}^{\prime}, M_{i}\right)-d\left(c_{i}, M_{i}\right)+d\left(c_{i}^{\prime}, V_{i}\right)-d\left(c_{i}, V_{i}\right) \tag{2}
\end{equation*}
$$

Now, we are ready to estimate $d^{\prime}\left(c_{i}^{\prime}, W_{i}\right)-d^{\prime}\left(c_{i}, W_{i}\right)$. Multiply Inequality (2) by $\alpha$ and sum over all $i$, then add it to Inequality (1):

$$
\begin{aligned}
& \sum_{i=1}^{k}\left[d^{\prime}\left(c_{i}^{\prime}, W_{i}\right)-d^{\prime}\left(c_{i}, W_{i}\right)+\alpha d\left(c_{i}, W_{i}\right)-\alpha d\left(c_{i}^{\prime}, W_{i}\right)\right] \\
\leq & \sum_{i=1}^{k}\left[d^{\prime}\left(c_{i}, A_{i}\right)-d^{\prime}\left(c_{i}^{\prime}, A_{i}\right)+\alpha d\left(c_{i}^{\prime}, M_{i}\right)-\alpha d\left(c_{i}, M_{i}\right)\right. \\
& \left.+\sum_{p \in V_{i}} d^{\prime}(c(p), p)-d^{\prime}\left(c_{i}^{\prime}, V_{i}\right)+\alpha d\left(c_{i}^{\prime}, V_{i}\right)-\alpha d\left(c_{i}, V_{i}\right)\right] .
\end{aligned}
$$

Rewrite it as $\sum_{i} S_{i} \leq \sum_{i} T_{i}$ where $S_{i}$ and $T_{i}$ are the terms related to index $i$ on the two sides respectively. Then the claim follows from the following bounds on $S_{i}$ and $T_{i}$ and the fact that $\left|C_{i}\right| / 3 \geq\left|A_{i}\right|+\left|M_{i}\right|+\alpha+2$ :

$$
S_{i} \geq(\alpha-1) I_{i} d\left(c_{i}, W_{i}\right)-\alpha d\left(c_{i}, c_{i}^{\prime}\right), \quad T_{i} \leq \alpha\left(\left|A_{i}\right|+\left|M_{i}\right|+\alpha+1\right) d\left(c_{i}, c_{i}^{\prime}\right) .
$$

It remains to prove the two bounds by a case-by-case study. For $S_{i}$, we have three cases: if $c_{i}^{\prime}=c_{i}$, then $S_{i}=0$; if $c_{i}^{\prime} \neq c_{i}$ and $c\left(c_{i}^{\prime}\right) \neq c_{i}$, then it is $0+(\alpha-1) d\left(c_{i}, W_{i}\right)$; if $c\left(c_{i}^{\prime}\right)=c_{i}$, then it is $-\alpha d\left(c_{i}, c_{i}^{\prime}\right)+(\alpha-1) d\left(c_{i}, W_{i}\right)$. In conclusion, in any case $S_{i} \geq(\alpha-1) I_{i} d\left(c_{i}, W_{i}\right)-\alpha d\left(c_{i}, c_{i}^{\prime}\right)$.

For $T_{i}$, in the easy case when $c_{i}^{\prime}=c_{i}$, most terms cancel out, and we get $T_{i}=$ $\sum_{p \in V_{i}} d(p, c(p))-\alpha d\left(c_{i}, V_{i}\right)$. Then $T_{i} \leq(\alpha-\alpha) d\left(c_{i}, V_{i}\right)=0$ since $V_{i}$ are the selected bad points. If $c_{i}^{\prime} \neq c_{i}$, Fact 1 leads to

$$
\begin{aligned}
T_{i} & \leq \alpha d\left(c_{i}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)+\alpha(1+\alpha) d\left(c_{i}^{\prime}, c_{i}\right)-\alpha d\left(c_{i}^{\prime}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right) \\
& +\alpha d\left(c_{i}^{\prime}, M_{i}\right)-\alpha d\left(c_{i}, M_{i}\right) \\
& +\sum_{p \in V_{i}} d(p, c(p))-\alpha d\left(c_{i}^{\prime}, V_{i}\right)+\alpha d\left(c_{i}^{\prime}, V_{i}\right)-\alpha d\left(c_{i}, V_{i}\right)
\end{aligned}
$$

The first line is bounded by $\alpha\left(\alpha+1+\left|A_{i}\right|\right) d\left(c_{i}, c_{i}^{\prime}\right)$. The second line is bounded by $\alpha d\left(c_{i}, c_{i}^{\prime}\right)\left|M_{i}\right|$. The third line is bounded by 0 , since $V_{i}$ are the selected bad points and thus $\sum_{p \in V_{i}} d(p, c(p)) \leq \alpha d\left(c_{i}, V_{i}\right)$. In conclusion, in any case $T_{i} \leq \alpha\left(\left|A_{i}\right|+\left|M_{i}\right|+\right.$ $\alpha+1) d\left(c_{i}, c_{i}^{\prime}\right)$.

Proof of Theorem 3. Combining Claim 1 and 2 leads to

$$
\sum_{1 \leq i \leq k}\left|C_{i}\right| d\left(c_{i}, c_{i}^{\prime}\right)\left[1-(\alpha+2) I_{i} /(\alpha+1)\right] \geq 0 .
$$

If $I_{i}=0$, then $d\left(c_{i}, c_{i}^{\prime}\right)=0$; if $I_{i}=1$, the coefficient of $d\left(c_{i}, c_{i}^{\prime}\right)$ is negative. So all terms on the left hand side are equal to 0 , i.e. $d\left(c_{i}, c_{i}^{\prime}\right)=0(1 \leq i \leq k)$. Then points in $\hat{B}_{i}$ will move to other clusters after perturbation, which means that $\hat{B}_{i} \subseteq M_{i}$. So $\left|\cup_{i} \hat{B}_{i}\right| \leq\left|\cup_{i} M_{i}\right| \leq \epsilon n$. In particular, $\left|\hat{B}_{i}\right| \leq \epsilon n$ for any $i$. Then $\left|B_{i}\right| \leq \epsilon n$, otherwise $\left|\hat{B}_{i}\right|$ would be $\epsilon n+1$. Therefore, $\hat{B}_{i}=B_{i}$, and $|B|=\left|\cup_{i} \hat{B}_{i}\right| \leq \epsilon n$.

$$
\left|G_{1}\right|=\frac{1-2 \epsilon}{2} n \quad|B|=\epsilon n \quad\left|G_{2}\right|=\frac{1-2 \epsilon}{2} n
$$



Figure 4: The optimality of the bound on the number of bad points for $(\alpha, \epsilon)$ perturbation resilient $k$-median instances.

Note 1: The bound in Theorem 3 is optimal in the sense that for any $\alpha>1, \epsilon<\frac{1}{5}$, we can easily construct an ( $\alpha, \epsilon$ )-perturbation resilient 2 -median instance which has $\epsilon n$ bad points.

The instance is shown in Figure 4. It has 3 groups of points: $G_{1}, G_{2}$, and $B$. Both $G_{1}$ and $G_{2}$ have $(1-\epsilon) n / 2$ points, and $B$ has $\epsilon n$ points. Let $M$ be a sufficiently large constant, say, $M>n^{2} / \epsilon$. The distances within the same group are 1 , while those between the points in $G_{1}$ and $G_{2}$ are $M$, those between the points in $B$ and $G_{1}$ are $\frac{M}{\alpha+1}+1$, and those between the points in $B$ and $G_{2}$ are $\frac{\alpha M}{\alpha+1}-1$. The instance satisfies the triangle inequality, which can be verified by a case-by-case study. The optimal clustering before perturbation has one center in $G_{1}$ and the other in $G_{2}$. Then $B$ are trivially bad points, and thus we have $\epsilon n$ bad points in this instance.

Now we show that the instance is $(\alpha, \epsilon)$-perturbation resilient. To prove that the optimal clustering after perturbation $\mathcal{C}^{\prime}$ is $\epsilon$-close to the original optimal clustering, it suffices to show that $\mathcal{C}^{\prime}$ has one center from $G_{1} \cup B$ and the other center from $G_{2}$. Assume for contradiction that this is not true. If both centers come from $G_{2}$, the cost of points in $G_{1}$ is $\frac{(1-\epsilon) n}{2} M$. On the other hand, the optimal cost before perturbation is $(1-\epsilon) n-2+\epsilon n\left(\frac{M}{\alpha+1}+1\right)$, so the optimal cost after perturbation is no more than $\alpha\left[(1-\epsilon) n-2+\epsilon n\left(\frac{M}{\alpha+1}+1\right)\right]$. But this is smaller than $\frac{(1-\epsilon) n}{2} M$, which is a contradiction. Similarly, we get a contradiction if both centers come from $G_{1} \cup B$.

Note 2: Theorem 3 requires that the sizes of the optimal clusters are sufficiently large. This makes sure that a majority of points in each optimal cluster stay after moving at most $\epsilon n$ points, which means $W_{i}$ is significantly larger than $V_{i}, A_{i}$ and $M_{i}$. This fact is crucial for proving the theorem. In the following subsections, we assume $n>\epsilon / \alpha$, so that when $\alpha-1=O(1)$, the requirement in Theorem 3 is simplified as $\min _{i}\left|C_{i}\right|=\Omega(\epsilon n)$.

### 2.3.2 Approximating the Optimum Clustering

Now, we consider approximating the cost of the optimum clustering. We can see that after removing the bad points, the optimal clusters are far apart from each other. In order to get rid of the influence of the bad points, we generate a list of blobs, which form a partition of the data points, and each of which contains only good points from one optimal cluster. Then we construct a tree on the list of blobs with a pruning that assigns all good points correctly. We will show that this pruning has low cost, so the lowest cost pruning of the tree is a good approximation. The details are described in Algorithm 2.

A key step is to generate the list of almost "pure" blobs, which is described in Algorithm 3. Suppose for any $i$ and any good point $p \in G_{i}$, its $\gamma\left|G_{i}\right|$ nearest neighbors contain no good points outside $C_{i}$. Also suppose the algorithm knows the value of $\gamma$. Informally, the algorithm maintains a threshold $t$. At each threshold, for each point $p$ that has not been added to the list, the algorithm checks its $\gamma t$ nearest neighbors $N_{\gamma t}(p)$. It constructs a graph $F_{t}$ by connecting any two points that have most neighbors in common. It then builds another graph $H_{t}$ by connecting any two points that have sufficiently many neighbors in $F_{t}$, and adds sufficiently large components in $H_{t}$ to the list. Finally, for each remaining point $p$, it checks if most of $p$ 's neighbors are in the list and if there are blobs containing a significant amount of $p$ 's neighbors. If so, it inserts $p$ into such a blob with the smallest median distance. Then the threshold is increased and the above steps are repeated.

The intuition behind Algorithm 3 is as follows. As mentioned above, the algorithm works when for any $i$ and any good point $p \in G_{i}$, the $\gamma\left|G_{i}\right|$ nearest neighbors of $p$ contain no good points outside $C_{i}(\gamma=1$ for the $k$-median instances considered in this section, as shown in Lemma 3; $\gamma=\frac{4}{5}$ for the min-sum instances considered in Section 2.5, as shown in Claim 5). Without loss of generality, assume $\left|C_{1}\right| \leq$ $\left|C_{2}\right| \leq \cdots \leq\left|C_{k}\right|$. When $t \leq\left|C_{1}\right|$, good points in different clusters do not have
most neighbors in common and thus are not connected in $F_{t}$. However, they may be connected by a path of bad points. So we further build the graph $H_{t}$ to disconnect such paths, which ensures that the blobs added into the list contain only good points from one optimal cluster. The final insert step (Step 6) makes sure that when $t=\left|C_{1}\right|$, all remaining good points in $C_{1}$ will be added to the list and will not affect the construction of blobs from other optimal clusters. We can show by induction that, at the end of the iteration $t=\left|C_{i}\right|$, all good points in $C_{j}(j \leq i)$ are added to the list. When $t$ is large enough, any remaining bad points are inserted into the list, so the points are partitioned into a list of almost pure blobs. The formal guarantee for Algorithm 3 is stated in Lemma 4.

Another key step is to construct a tree on these blobs. Since good points are closer to good points in the same optimal cluster than to those in other clusters (Lemma 3), there exist algorithms that can build a tree with a pruning that assigns all good points correctly. In particular, we can use the robust linkage procedure in [16], which repeatedly merges the two blobs $C, C^{\prime}$ with the maximum score $\left(C, C^{\prime}\right)$ defined as follows. For each $p \in C$, sort the other blobs in decreasing order of the median distance between $p$ and points in the blob, and let $\operatorname{rank}\left(p, C^{\prime}\right)$ denote the rank of $C^{\prime}$. Then define $\operatorname{rank}\left(C, C^{\prime}\right)=\operatorname{median}_{x \in C}\left[\operatorname{rank}\left(x, C^{\prime}\right)\right]$ and $\operatorname{score}\left(C, C^{\prime}\right)=$ $\min \left[\operatorname{rank}\left(C, C^{\prime}\right), \operatorname{rank}\left(C^{\prime}, C\right)\right]$. Intuitively, for any blobs $A, A^{\prime}$ from the same optimal cluster and $D$ from a different cluster, good points in $A$ always rank $A^{\prime}$ later than $D$ in the sorted list, so $\operatorname{rank}\left(A, A^{\prime}\right)>\operatorname{rank}(A, D) . \operatorname{Similarly}, \operatorname{rank}\left(A^{\prime}, A\right)>\operatorname{rank}\left(A^{\prime}, D\right)$, and thus $\operatorname{score}\left(A^{\prime}, A\right)>\operatorname{score}(A, D)$. This means the algorithm will always merge blobs from the same cluster before merging them with blobs outside, so there is a pruning that assigns all good points correctly.

In the following, we prove that Algorithm 2 outputs a good approximation. We begin with a key property of ( $\alpha, \epsilon$ )-perturbation resilience that ensures the success of Algorithm 3.

Algorithm $2 k$-median, $(\alpha, \epsilon)$ perturbation resilience
Input: Data set $P$, distance function $d(\cdot, \cdot)$ on $P, \min _{i}\left|C_{i}\right|, \epsilon>0$
1: Run Algorithm 3 to generate a list $\mathcal{L}$ of blobs with parameters $m_{B}=\epsilon n, \gamma=1$.
2: Run the robust linkage procedure in [16] to get a cluster tree $\mathcal{T}$.
3: Run dynamic programming on $\mathcal{T}$ to get the minimum cost pruning $\tilde{\mathcal{C}}$ and its centers $\tilde{\mathbf{c}}$.
Output: Clustering $\tilde{\mathcal{C}}$ and its centers $\tilde{\mathbf{c}}$.

## Algorithm 3 Generating interesting blobs

Input: Data set $P$, distance function $d(\cdot, \cdot)$ on $P$, the size of the smallest optimal cluster $\min _{i}\left|C_{i}\right|$, the upper bound on the number of bad points $m_{B}$, a parameter $\gamma \in[4 / 5,1]$
Let $N_{r}(p)$ denote the $r$ nearest neighbors of $p$ in $P$.
Let $\mathcal{L}=\emptyset, A_{P}=P$. Let the initial threshold $t=\min _{i}\left|C_{i}\right|$.
Construct a graph $F_{t}$ by connecting $p, q \in A_{P}$ if
$\left|N_{\gamma t}(p) \cap N_{\gamma t}(q)\right|>(2 \gamma-1) t-2 m_{B}$.
4: Construct a graph $H_{t}$ by connecting points $p, q \in A_{P}$ if $p, q$ share more than $m_{B}$ neighbors in $F_{t}$.
5: Add to $\mathcal{L}$ all the components $C$ of $H_{t}$ with $|C| \geq \frac{1}{2} \min _{i}\left|C_{i}\right|$ and remove them from $A_{P}$.
6: For each point $p \in A_{P}$, check if most of $N_{\gamma t}(p)$ are in $\mathcal{L}$ and if there exists $C \in \mathcal{L}$ containing a significant number of points in $N_{\gamma t}(p)$. More precisely, check if
(1) $\left|N_{\gamma t}(p) \backslash \mathcal{L}\right| \leq \frac{1}{2} \min _{i}\left|C_{i}\right|+2 m_{B}$;
(2) $\mathcal{L}_{p} \neq \emptyset$ where $\mathcal{L}_{p}=\left\{C \in \mathcal{L}:\left|C \cap N_{\gamma t}(p)\right| \geq\left(\gamma-\frac{3}{5}\right)|C|\right\}$.

If so, assign $p$ to the blob in $\mathcal{L}_{p}$ of smallest median distance, remove $p$ from $A_{P}$.
7: While $\left|A_{P}\right|>0$, increase $t$ by 1 and go to Step 3 .
Output: The list $\mathcal{L}$.

Lemma 3. When $\alpha>4$, for any good points $p_{1}, p_{2} \in G_{i}, q \in G_{j}(j \neq i)$, we have $d\left(p_{1}, p_{2}\right)<d\left(p_{1}, q\right)$. Consequently, for any good point $p \in G_{i}$, all its $\left|G_{i}\right|$ nearest neighbors belong to $C_{i} \cup B$.

Proof. By the same proof in Lemma 2(2), we have

$$
d\left(p_{1}, q\right)>(\alpha-1) d\left(p_{1}, c_{i}\right) \quad \text { and } \quad d\left(p_{2}, q\right)>(\alpha-1) d\left(p_{2}, c_{i}\right) .
$$

These then lead to

$$
d\left(p_{1}, p_{2}\right) \leq d\left(p_{1}, c_{i}\right)+d\left(p_{2}, c_{i}\right)<\frac{d\left(p_{1}, q\right)+d\left(p_{2}, q\right)}{\alpha-1} \leq \frac{2 d\left(p_{1}, q\right)+d\left(p_{1}, p_{2}\right)}{\alpha-1}
$$

and thus $d\left(p_{1}, p_{2}\right)<\frac{2}{\alpha-2} d\left(p_{1}, q\right)$. When $\alpha>4$, we have $d\left(p_{1}, p_{2}\right)<d\left(p_{1}, q\right)$.

Lemma 4. Suppose the number of bad points is bounded by $m_{B}$, and for any $i$ and any good point $p \in G_{i}$, all its $\gamma\left|G_{i}\right|$ nearest neighbors in $P$ are from $C_{i} \cup B$. If $\min _{i}\left|C_{i}\right|>30 m_{B}$, then Algorithm 3 generates a list $\mathcal{L}$ of blobs each of size at least $\frac{1}{2} \min _{i}\left|C_{i}\right|$ such that:

- The blobs in $\mathcal{L}$ form a partition of $P$.
- Each blob in $\mathcal{L}$ contains good points from only one optimal cluster.

Proof. We prove the following two claims by induction on $i \leq k$ :

- For any $t \leq\left|G_{i}\right|$, any blob in the list $\mathcal{L}$ only contains good points from only one optimal cluster; all blobs have size at least $\frac{1}{2} \min _{i}\left|C_{i}\right|$.
- At the beginning of the iteration $t=\left|G_{i}\right|+1$, any good point $p \in G_{j}(j \leq i)$ has been assigned to a blob in the list that contains good points only from $C_{j}$.

The first two claims imply that each blob in the list contains good points from only one optimal cluster. Moreover, at the beginning of the iteration $t=\left|G_{k}\right|+1$, all good points have been assigned to one of the blobs in $\mathcal{L}$, so there are only bad points left, the number of which is smaller than $\frac{1}{2} \min _{i}\left|C_{i}\right|$. These remaining points will eventually be assigned to the blobs before $\gamma t>n$, so the blobs form a partition of $P$.

The claims are clearly both true initially. We show now that as long as $t \leq\left|G_{1}\right|$, the graphs $F_{t}$ and $H_{t}$ have the following properties.

- No good point $p_{i}$ in cluster $C_{i}$ is connected in $F_{t}$ to a good point $p_{j}$ in a different cluster $C_{j}$. By assumption, $p_{i}$ has no neighbors outside $C_{i} \cup B$ and $p_{j}$ has no neighbors outside $C_{j} \cup B$, so they share at most $m_{B}<(2 \gamma-1) t-2 m_{B}$ neighbors.
- No point $q$ is connected in $F_{t}$ to both a good point $p_{i}$ in $C_{i}$ and a good point $p_{j}$ in a different cluster $C_{j}$. If $q$ is connected to $p_{i}$, then $\left|N_{\gamma t}\left(p_{i}\right) \cap N_{\gamma t}(q)\right|>$
$(2 \gamma-1) t-2 m_{B}$. Since $p_{i}$ has no neighbors outside $C_{i} \cup B, N_{\gamma t}(q)$ contains more than $(2 \gamma-1) t-3 m_{B} \geq \gamma t / 2$ points from $G_{i}$. Similarly, if $q$ is connected to $p_{j}$, then $N_{\gamma t}(q)$ contains more than $\gamma t / 2$ points from $G_{j}$, which is contradictory.
- All the components in $H_{t}$ of size at least $\frac{1}{2} \min _{i}\left|C_{i}\right|$ will only contain good points from one optimal cluster. As there are at most $m_{B}$ bad points, any two points connected in $H_{t}$ must be connected in $F_{t}$ to one good point. Then by the above two properties, points on a path in $H_{t}$ must be connected in $F_{t}$ to good points in the same cluster, so there is no path connecting good points from different clusters.

We can use the three properties to argue the first claim: as long as $t \leq\left|G_{1}\right|$, each blob in $\mathcal{L}$ contains good points from at most one optimal cluster. This is true at the beginning and by the third property, for any $t \leq\left|G_{1}\right|$, anytime we insert a whole new blob in the list in Step 5, that blob must contain point from at most one optimal cluster. We now argue that this property is never violated as we assign points to blobs already in the list in Step 6. Suppose a good point $p \in C_{i}$ is inserted into $C \in \mathcal{L}$. Then $C \in \mathcal{L}_{p}$, which means $\left|N_{\gamma t}(p) \cap C\right| \geq|C| / 2>m_{B}$. So $N_{\gamma t}(p) \cap C$ contains at least one good point, which must be from $C_{i}$ since $N_{\gamma t}(p)$ contains no good points outside $C_{i}$. Then by induction $C$ must contain only good points from $C_{i}$, and thus adding $p$ to $C$ does not violate the first claim.

We now show the second claim: after the iteration $t=\left|G_{1}\right|$, all the good points in $C_{1}$ have already been assigned to a blob in the list that only contains good points from $C_{1}$. There are two cases. First, if at the beginning of the iteration $t=\left|G_{1}\right|$, there are still at least $\frac{1}{2} \min _{i}\left|C_{i}\right|$ points from the good point set $G_{1}$ that do not belong to blobs in the list. Any such good point has all $\gamma\left|G_{1}\right|$ neighbors in $C_{1} \cup B$. Then any two such good points share at least $2 \gamma\left|G_{1}\right|-\left|C_{1} \cup B\right| \geq(2 \gamma-1)\left|G_{1}\right|-|B| \geq(2 \gamma-1) t-2 m_{B}$ neighbors. So they will connect to each other in $F_{t}$ and then in $H_{t}$, and thus we
will add one blob to $\mathcal{L}$ containing all these points. Second, it could be that at the beginning of the iteration $t=\left|G_{1}\right|$, all but less than $\frac{1}{2} \min _{i}\left|C_{i}\right|$ good points in $G_{1}$ have been assigned to a blob in the list. Denote these points as $E$. Any point $p \in E$ has no neighbors outside $C_{1} \cup B$. Then $\left|N_{\gamma t}(p) \backslash \mathcal{L}\right| \leq|E|+|B| \leq \frac{1}{2} \min _{i}\left|C_{i}\right|+2 m_{B}$. Also, there exists a blob $C$ containing good points from $C_{1}$ such that $C \in \mathcal{L}_{p}$. Otherwise, $N_{\gamma t}(p)$ contains at most $\left(\gamma-\frac{3}{5}\right)\left(\left|C_{1} \cup B\right|\right)<\gamma\left|C_{1}\right|-\frac{1}{2}\left|C_{1}\right|-2 m_{B}$ points in $C_{1} \cap \mathcal{L}$, while it contains at most $|E|$ good points in $C_{1} \backslash \mathcal{L}$ and contains no points outside $C_{1} \cup B$. In total, $N_{\gamma t}(p)$ has less than $\gamma t$ points, which is contradictory. So $\mathcal{L}_{p} \neq \emptyset$ and $p$ will be added to the list in Step 6.

We then iterate the argument on the remaining set $A_{P}$. The key point is that for $t \geq\left|G_{i}\right|, i>1$, we have that all the good points in $C_{1}, C_{2}, \ldots, C_{i}$ have already been assigned to blobs in $\mathcal{L}$.

Lemma 3 and 4 show that Algorithm 3 with parameters $m_{B}=\epsilon n$ and $\gamma=1$ produces a list of sufficiently large, almost pure blobs. Then the robust linkage procedure in [16] can build a tree on these blobs with a pruning that assigns all good points correctly. Now it suffices to show that this pruning is a good approximation, for which we need to bound the cost increased by the bad points assigned incorrectly. The following property of these bad points turns out to be useful. Intuitively, Algorithm 3 is designed such that whenever a bad point is added to a blob containing good points from a different cluster, it must be closer to a significantly amount of points in that cluster than to a significantly amount of points in its own cluster. Then the cost increased by incorrectly assigning each such bad point is small, resulting in a good approximation.

Lemma 5. Suppose for any good point $p \in G_{i}$, all its $\left|G_{i}\right|$ nearest neighbors in $P$ are from $C_{i} \cup B$, and $\min _{i}\left|C_{i}\right|>30 m_{B}$. When running Algorithm 3 with $\gamma=1$, if a bad point $q \in B_{i}$ is assigned to a blob $C$ containing good points from a different optimal
clustering $C_{j}$, then there exist $m=\frac{1}{5} \min _{i}\left|C_{i}\right|$ points $Z_{i}$ from $C_{i}$, and $m$ points $Z_{j}$ from $C_{j}$, such that $d\left(q, Z_{i}\right) \geq d\left(q, Z_{j}\right)$.

Proof. There are two cases: $q$ is added into $C$ in (1) Step 5 or (2) Step 6.
Case 1 There must be a path in $H_{t}$ connecting $q$ to a good point in $C_{j}$ at threshold $t$. For any edge $(x, y)$ in $H_{t}$, since $x, y$ share at least $\epsilon n$ neighbors in $F_{t}$ and there are at most $\epsilon n$ bad points, they share at least one good point as neighbor in $F_{t}$. As shown in the proof of Lemma 4, no point can connect to good points from different clusters, so in $F_{t}$ all points on the path must connect to good points in $C_{j}$. In particular, $q$ is connected in $F_{t}$ to a good point $p \in G_{j}$. Then $\left|N_{t}(p) \cap N_{t}(q)\right|>t-2 m_{B}$. Since $p$ is still in $A_{P}, t \leq\left|G_{j}\right|$, and thus $N_{t}(p)$ contains no points outside $C_{j} \cup B$. This means that at least $t-3 m_{B} \geq m$ points in $N_{t}(q)$ are good points in $C_{j}$, then we can select $m$ points $Z_{j}$ from $N_{t}(q) \cap G_{j}$. We also have that at most $2 m_{B}$ points in $N_{t}(q)$ are points in $C_{i}$, so we can select $m$ points $Z_{i}$ from $C_{i} \backslash N_{t}(q)$.

Case 2 There are three subcases when $q$ is inserted into $C$ at threshold $t$.
(1) There is no good points from $C_{i}$ in the list. Since $\left|N_{t}(q) \backslash \mathcal{L}\right| \leq \frac{1}{2} \min _{i}\left|C_{i}\right|+$ $2 m_{B}, N_{t}(q)$ contains at most this number of good points in $C_{i}$. This means at least $\frac{1}{2} \min _{i}\left|C_{i}\right|-2 m_{B}>m$ good points in $C_{i}$ are outside $N_{t}(q)$, from which we can select $Z_{i}$. On the other hand, we can select $Z_{j}$ as follows. When inserting $q$ into $C$, we have $\left|N_{t}(q) \cap C\right| \geq \frac{2}{5}|C| \geq m+m_{B}$. Since $C$ contains only good points from $C_{j}$ and some bad points, $N_{t}(q) \cap C$ contains at least $m$ good points in $C_{j}$, from which we can select $Z_{j}$. Since $Z_{j}$ are from $N_{t}(q)$ and $Z_{i}$ are outside $N_{t}(q)$, we have $d\left(q, Z_{i}\right) \geq d\left(q, Z_{j}\right)$.
(2) There exists $C^{\prime} \in \mathcal{L}$ containing good points from $C_{i}$, but $C^{\prime} \notin \mathcal{L}_{p}$. This means $\left|B(q, t) \cap C^{\prime}\right| \leq \frac{2}{5}\left|C^{\prime}\right|$, so there are at least $\frac{3}{5}\left|C^{\prime}\right| \geq m+m_{B}$ points in $C^{\prime}$ are outside $N_{t}(q)$. At least $m$ of these points are good points from $C_{i}$, since $C^{\prime}$ contains only good points from $C_{i}$ and at most $m_{B}$ bad points. On the other hand, we can select $Z_{j}$ as in the first subcase.
(3) There exists $C^{\prime} \in \mathcal{L}_{p}$ containing good points from $C_{i}$. Since $q$ is assigned to $C$
rather than $C^{\prime}$ according to median distances, we know that at least half of the points $Z_{j}^{\prime}$ from $C$ are closer to $q$ than at least half of the points $Z_{i}^{\prime}$ from $C^{\prime}$. Since there are at most $m_{B}$ bad points, we can select $m$ good points $Z_{j}$ from $Z_{j}^{\prime}$ and select $m$ good points $Z_{i}$ from $Z_{i}^{\prime}$. Note that $Z_{j}$ are all from $G_{j}$ and $Z_{i}$ are all from $G_{i}$, so $d\left(q, Z_{i}\right) \geq d\left(q, Z_{j}\right)$.

Theorem 4. If the clustering instance is ( $\alpha, \epsilon$ )-perturbation resilient for $\alpha>4$ and $\epsilon \leq \rho / 30$ where $\rho=\frac{\min _{i}\left|C_{i}\right|}{n}$, then Algorithm 2 produces a clustering which is $(1+$ $\frac{5 \epsilon}{\rho}$ )-approximation to the optimal clustering with respect to the $k$-median objective in polynomial time.

Proof. By Lemma 3 and 4, Algorithm 3 partition the points into a list of blobs, each of which has size at least $\frac{1}{2} \min _{i}\left|C_{i}\right|$ and contains only good points from one optimal cluster. Let $B_{i}^{\prime}$ denote the bad points that are assigned to blobs containing good points in $C_{i}$. By Lemma 3, Theorem 9 in [16] can be applied to $\mathcal{L}$, by which we know that $\left\{\left(C_{i} \cap G\right) \cup B_{i}^{\prime}\right\}$ is a pruning of the tree. Suppose the cost of the optimum is $\mathcal{O P} \mathcal{T}$. We now show that this pruning, using the original centers $\left\{c_{i}\right\}$, is a $\left(1+\frac{5 \epsilon}{\rho}\right)$-approximation to $\mathcal{O P} \mathcal{T}$.

Suppose a bad point $q \in C_{i}$ is assigned to a blob $C$ containing good points from a different optimal cluster $C_{j}$. By Lemma 5 , there exist $m=\frac{1}{5} \min _{i}\left|C_{i}\right|$ points $Z_{i}$ from $C_{i}$, and $m$ points $Z_{j}$ from $C_{j}$, such that $d\left(q, Z_{i}\right) \geq d\left(q, Z_{j}\right)$. Then the increase in cost due to $q$ is bounded as follows:

$$
\begin{aligned}
d\left(q, c_{j}\right)-d\left(q, c_{i}\right) & \leq \frac{d\left(q, Z_{j}\right)+d\left(c_{j}, Z_{j}\right)}{m}-\frac{d\left(q, Z_{i}\right)-d\left(c_{i}, Z_{i}\right)}{m} \\
& \leq \frac{1}{m}\left[d\left(c_{j}, Z_{j}\right)+d\left(c_{i}, Z_{i}\right)\right] \leq \frac{\mathcal{O} \mathcal{P} \mathcal{T}}{m}
\end{aligned}
$$

As there are at most $\epsilon n$ bad points and $m=\frac{\min _{i}\left|C_{i}\right|}{5}$, the increase of cost is at most $\frac{\epsilon n}{m} \mathcal{O P} \mathcal{T}=\frac{5 \epsilon}{\rho} \mathcal{O} \mathcal{P} \mathcal{T}$.

Running Time In Algorithm 3, for each $p \in P$, we first sort all the other points in ascending order of distances in time $O\left(n^{2} \log n\right)$. At each threshold $t$, think of
a directed $t$-regular graph $E_{t}$, where, for each point $q$ in the $t$ nearest neighbors of a point $p$, there is a directed edge from $p$ to $q$ in $E_{t}$. Let $A_{E}$ denote the adjacent matrix for $E_{t}$, and let $N=A_{E} A_{E}^{T}$. Then $N_{p q}$ is the number of common neighbors between $p$ and $q$, which can be used in constructing $F_{t}$. Computing $N$ takes time $O\left(n^{\omega}\right)$, the state of the art for matrix multiplication. The same method can be used to compute the number of common neighbors in $F_{t}$ and construct $H_{t}$. Since there are $O(n)$ thresholds, the total time for constructing $F_{t}$ and $H_{t}$ is $O\left(n^{\omega+1}\right)$. For the other steps, adding a blob takes time $O\left(n^{2}\right)$ and inserting a point takes time $O\left(n^{2}\right)$. These steps can be performed at most $O(n)$ times, so they take $O\left(n^{3}\right)$ time. In total, Algorithm 3 takes time $O\left(n^{\omega+1}\right)$. Since the robust linkage algorithm takes time at most $O\left(n^{\omega+1}\right)([16])$, and the dynamic programming takes time $O\left(n^{3}\right)$ (see Appendix A.1), the running time of Algorithm 2 is $O\left(n^{\omega+1}\right)$.

We observe that for $\alpha>4$ we can tolerate $\epsilon$ up to approximately $\rho$ (recalling that $\rho$ is the fraction of points in the smallest cluster), and beat the lower bound of $(1+1 / e)$ on the best approximation achievable on worst case instances for the metric $k$-median objective $[55,59]$ when $\epsilon<\frac{\rho}{5 e}$.

### 2.3.3 Sublinear Time Algorithm

Consider a clustering instance $(X, d)$ that is $(\alpha, \epsilon)$-perturbation resilient to $k$-median. For simplicity, suppose the distances are normalized to $[0,1]$. Let $N=|X|$. Let $\rho=\min _{i}\left|C_{i}\right| / N$ denote the fraction of the points in the smallest cluster, $\zeta=\Phi_{X}(\mathbf{c}) / N$ denote the average cost of the points in the optimum clustering.

Theorem 5. Suppose $(X, d)$ is ( $\alpha, \epsilon$ )-perturbation resilient for $\alpha>4, \epsilon<\rho / 100$. Then with probability $\geq 1-\delta$, we can get an implicit clustering that is $2\left(1+\frac{16 \epsilon}{\rho}\right)$ approximation in time poly $\left(\log \frac{N}{\delta}, k, \frac{1}{\epsilon}, \frac{1}{\zeta}\right)$.

Intuition. The main idea is to run Algorithm 2 on a sufficiently large sample $P$ to obtain the minimum cost pruning and the centers $\tilde{\mathbf{c}}$. Then the implicit clustering of
the whole space $X$ assigns each point in $X$ to its nearest center in $\tilde{\mathbf{c}}$. To show that this is a good approximation, it suffices to show that $\Phi_{P}(\tilde{\mathbf{c}})$ is close to $\Phi_{P}(\mathbf{c})$ where c are the optimal centers for $X$. Note that Algorithm 2 builds a tree with a pruning $\mathcal{P}^{\prime}$ that assigns all good points correctly. The key is to use the cost of this pruning as a bridge for $\Phi_{P}(\tilde{\mathbf{c}})$ and $\Phi_{P}(\mathbf{c})$ : on one hand, $\Phi_{P}(\tilde{\mathbf{c}})$ is no more than the cost of $\mathcal{P}^{\prime}$ since $\tilde{\mathbf{c}}$ is the centers in the minimum cost pruning; on the other hand, the cost of $\mathcal{P}^{\prime}$ is roughly bounded by twice $\Phi_{P}(\mathbf{c})$ by triangle inequality.

Proof. We sample a set $P$ of size $n=\Theta\left(\frac{k}{\epsilon^{2} \zeta^{2}} \ln \frac{N}{\delta}\right)$ and run Algorithm 2 on $P$ to obtain the minimum cost pruning $\tilde{\mathcal{C}}$ and its centers $\tilde{\mathbf{c}}$. The implicit clustering of the whole space $X$ then assigns each point in $X$ to its nearest neighbor in $\tilde{\mathbf{c}}$. Recall that if we partition $A$ into $\mathcal{P}$, the cost using centers $\mathbf{p}$ is denoted as $\Phi_{A}(\mathcal{P}, \mathbf{p})$. If we partition $A$ by assigning points to nearest centers in $\mathbf{p}$, the cost is denoted as $\Phi_{A}(\mathbf{p})$. We will show that the cost of implicit clustering $\Phi_{X}(\tilde{\mathbf{c}})$ approximates the optimum $\Phi_{X}(\mathbf{c})$.

First, we will prove that when $n$ is sufficiently large, with high probability, $\Phi_{X}(\tilde{\mathbf{c}}) / N \approx$ $\Phi_{P}(\tilde{\mathbf{c}}) / n$ and $\Phi_{X}(\mathbf{c}) / N \approx \Phi_{P}(\mathbf{c}) / n$. Formally, for every set of centers $\mathbf{p}$, if $n=$ $\Theta\left(\frac{k}{v^{2} \zeta^{2}} \log \frac{N}{\delta}\right)$ where $0<v<1$, then

$$
\operatorname{Pr}\left[\left|\frac{\Phi_{P}(\mathbf{p})}{n}-\frac{\Phi_{X}(\mathbf{p})}{N}\right|>v \frac{\Phi_{X}(\mathbf{p})}{N}\right] \leq 2 \exp \left\{-2 v^{2} \zeta^{2} n\right\} \leq \frac{\delta}{4 N^{k}}
$$

By the union bound, we have with probability at least $1-\delta / 4,(1-v) \Phi_{X}(\tilde{\mathbf{c}}) / N \leq$ $\Phi_{P}(\tilde{\mathbf{c}}) / n$ and $\Phi_{P}(\mathbf{c}) / n \leq(1+v) \Phi_{X}(\mathbf{c}) / N$. We can choose $v=\epsilon / 20$, then it is sufficient to show $\Phi_{P}(\tilde{\mathbf{c}}) \leq 2\left(1+\frac{12 \epsilon}{\rho}\right) \Phi_{P}(\mathbf{c})$.

Next, since $\tilde{\mathcal{C}}$ may be different from $\mathcal{C} \cap P$, we need to find a bridge for comparing $\Phi_{P}(\tilde{\mathbf{c}})$ and $\Phi_{P}(\mathbf{c})$. Now, we turn to analyze Algorithm 2 on $P$ to find such a bridge. First, we know that $X$ has at most $\epsilon N$ bad points. Since $n$ is sufficiently large, with probability at least $1-\delta / 4, P$ has at most $2 \epsilon n$ bad points. Similarly, with probability at least $1-\delta / 4$, for any $i,\left|C_{i} \cap P\right|>60 \epsilon n$. These ensure that Algorithm 2 can successfully produce a tree with a pruning $\mathcal{P}^{\prime}$ that assigns all good points in $P$
correctly, as shown in Theorem 4. Suppose in $P, \mathbf{c}^{\prime}$ are the optimal centers for $\mathcal{P}^{\prime}$. Then we can use $\Phi_{P}\left(\mathcal{P}^{\prime}, \mathbf{c}^{\prime}\right)$ as a bridge for comparing $\Phi_{P}(\tilde{\mathbf{c}})$ and $\Phi_{P}(\mathbf{c})$.

On one hand, $\Phi_{P}(\tilde{\mathbf{c}}) \leq \Phi_{P}(\tilde{\mathcal{C}}, \tilde{\mathbf{c}}) \leq \Phi_{P}\left(\mathcal{P}^{\prime}, \mathbf{c}^{\prime}\right)$. The first inequality comes from the fact that in $\Phi_{P}(\tilde{\mathbf{c}})$ each point is assigned to its nearest center and the second comes from that $\tilde{\mathcal{C}}$ is the minimum cost pruning.

On the other hand, $\Phi_{P}\left(\mathcal{P}^{\prime}, \mathbf{c}^{\prime}\right) \leq 2 \Phi_{P}\left(\mathcal{P}^{\prime}, \mathbf{c}\right) \leq 2\left(1+\frac{12 \epsilon}{\rho}\right) \Phi_{P}(\mathbf{c})$. The second inequality comes from an argument similar to that in Theorem 4 and the fact that $\Phi_{P}\left(\mathcal{P}^{\prime}, \mathbf{c}\right)$ is different from $\Phi_{P}(\mathbf{c})$ only on the bad points. The first inequality comes from the triangle inequality. More precisely, for any $N_{i}^{\prime} \in \mathcal{P}^{\prime}$,

$$
\begin{aligned}
2\left|N_{i}^{\prime}\right| \sum_{p \in N_{i}^{\prime}} d\left(p, c_{i}\right) & =\sum_{q \in N_{i}^{\prime}} \sum_{p \in N_{i}^{\prime}}\left[d\left(p, c_{i}\right)+d\left(q, c_{i}\right)\right] \geq \sum_{p \in N_{i}^{\prime}} \sum_{q \in N_{i}^{\prime}} d(p, q) \\
& \geq \sum_{p \in N_{i}^{\prime}} \sum_{q \in N_{i}^{\prime}} d\left(q, c_{i}^{\prime}\right)=\left|N_{i}^{\prime}\right| \sum_{q \in N_{i}^{\prime}} d\left(q, c_{i}^{\prime}\right)
\end{aligned}
$$

where the third step follows from the fact that $c_{i}^{\prime}$ is the optimal center for $N_{i}^{\prime}$.

Note: If we have an oracle that given a set of points $C_{i}^{\prime}$ finds the best center in $X$ for that set, then we can save a factor of 2 in the bound.

## $2.4 \alpha$-Perturbation Resilience for Min-Sum

In this section we provide an efficient algorithm for clustering $\alpha$-perturbation resilient instances for the metric min-sum $k$-clustering problem (Algorithm 4).

```
Algorithm 4 min-sum, \(\alpha\) perturbation resilience
Input: Data set \(P\), distance function \(d(\cdot, \cdot)\) on \(P, \min _{i}\left|C_{i}\right|\).
    1: Connect each point with its \(\frac{1}{2} \min _{i}\left|C_{i}\right|\) nearest neighbors.
    2: Initialize the clustering \(\mathcal{C}^{\prime}\) with each connected component being a cluster.
    3: Repeat till only one cluster remains in \(\mathcal{C}^{\prime}\) :
    merge clusters \(C, C^{\prime}\) in \(\mathcal{C}^{\prime}\) which minimize \(d_{\mathrm{a}}\left(C, C^{\prime}\right)\).
    4: Let \(\mathcal{T}\) be the tree with components as leaves and internal nodes corresponding to
    the merges performed.
    5: Run dynamic programming on \(\mathcal{T}\) to get the minimum min-sum cost pruning \(\tilde{\mathcal{C}}\).
Output: \(\tilde{\mathcal{C}}\).
```

Theorem 6. For $\left(3 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}\right)$-perturbation resilient instances, Algorithm 4 outputs the optimal min-sum $k$-clustering in polynomial time.

Intuition. To prove the theorem, first we show that the $\alpha$-perturbation resilience property implies the following (Lemma 6): for any two different optimal clusters $C_{i}$ and $C_{j}$ and any $A \subseteq C_{i}$, we have $\alpha d\left(A, C_{i} \backslash A\right)<d\left(A, C_{j}\right)$. This follows by considering the perturbation where $d^{\prime}(p, q)=\alpha d(p, q)$ if $p \in A, q \in C_{i} \backslash A$ and $d^{\prime}(p, q)=d(p, q)$ otherwise, and using the fact that the optimum does not change after the perturbation. This can be used to show that when $\alpha>3 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$, we have the following (Lemma 7): (1) for any optimal clusters $C_{i}$ and $C_{j}$ and any $A_{i} \subseteq C_{i}$, $A_{j} \subseteq C_{j}$ such that $\min \left(\left|C_{i} \backslash A_{i}\right|,\left|C_{j} \backslash A_{j}\right|\right)>\min _{i}\left|C_{i}\right| / 2$ we have $d_{\mathrm{a}}\left(A_{i}, A_{j}\right)>$ $\min \left\{d_{\mathrm{a}}\left(A_{i}, C_{i} \backslash A_{i}\right), d_{\mathrm{a}}\left(A_{j}, C_{j} \backslash A_{j}\right)\right\} ;(2)$ for any point $p$ in the optimal cluster $C_{i}$, twice its average distance to points in $C_{i} \backslash\{p\}$ is smaller than the distance to any point in other optimal cluster $C_{j}$. Claim (2) implies that for any point $p \in C_{i}$ its $\left|C_{i}\right| / 2$ nearest neighbors are in the same optimal cluster, so the leaves of the tree $\mathcal{T}$ are laminar to the optimum clustering. Claim (1) can be used to show that the merge steps preserve the laminarity with the optimal clustering, so the minimum cost pruning of $\mathcal{T}$ will be the optimal clustering, as desired.

We now prove the key lemmas mentioned above, and then present the detailed proof for the theorem.

Fact 2. For any nonempty set $A$ and $D \subseteq C$, we have $|D| d(A, C) \leq|C| d(A, D)+$ $|A| d(D, C \backslash D)$.

Proof. By the triangle inequality, $d_{\mathrm{a}}(A, C \backslash D) \leq d_{\mathrm{a}}(A, D)+d_{\mathrm{a}}(D, C \backslash D)$. The fact then follows from $d(A, C)=d(A, D)+d(A, C \backslash D)$.

Lemma 6. Suppose the clustering instance is $\alpha$-perturbation resilient to the min-sum objective. For any two different optimal clusters $C_{i}$ and $C_{j}$ and any $A \subseteq C_{i}$, we have $\alpha d\left(A, C_{i} \backslash A\right)<d\left(A, C_{j}\right)$.

Proof. This follows by considering a specific perturbation and using the fact that the optimum does not change after the perturbation. We define a perturbation as follows: $d^{\prime}(p, q)=\alpha d(p, q)$ if $p \in A, q \in C_{i} \backslash A$ or $q \in A, p \in C_{i} \backslash A$, and $d^{\prime}(p, q)=$ $d(p, q)$ otherwise. $d^{\prime}$ is a valid $\alpha$-perturbation of $d$, so the optimal clustering after perturbation should remain the same. Specially, its cost should be smaller than that of the clustering obtained by replacing $C_{i}, C_{j}$ with $C_{i} \backslash A, A \cup C_{j}$. After canceling the terms common in the two costs, we have $2 d^{\prime}\left(A, C_{i} \backslash A\right)<2 d^{\prime}\left(A, C_{j}\right)$, which implies $\alpha d\left(A, C_{i} \backslash A\right)<d\left(A, C_{j}\right)$.

Lemma 7. Suppose the clustering instance is $\alpha$-perturbation resilient to min-sum for $\alpha>3 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$.
(1) For any two different optimal clusters $C_{i}$ and $C_{j}$ and any $A_{i} \subseteq C_{i}, A_{j} \subseteq C_{j}$, if $\left|C_{i} \backslash A_{i}\right|$ and $\left|C_{j} \backslash A_{j}\right|$ are larger than $\min _{i}\left|C_{i}\right| / 2$, then

$$
d_{\mathrm{a}}\left(A_{i}, A_{j}\right)>\min \left[d_{\mathrm{a}}\left(A_{i}, C_{i} \backslash A_{i}\right), d_{\mathrm{a}}\left(A_{j}, C_{j} \backslash A_{j}\right)\right]
$$

(2) For any point $p$, all its $\min _{i}\left|C_{i}\right| / 2$ nearest neighbors are in the same optimal cluster.

Proof. (1) Let $\bar{A}_{i}=C_{i} \backslash A_{i}, \bar{A}_{j}=C_{j} \backslash A_{j}$. We have from Lemma 6 and Fact 2:

$$
\begin{array}{r}
\alpha d\left(A_{i}, \bar{A}_{i}\right)<d\left(A_{i}, C_{j}\right) \leq \frac{1}{\left|A_{j}\right|}\left[\left|C_{j}\right| d\left(A_{i}, A_{j}\right)+\left|A_{i}\right| d\left(A_{j}, \bar{A}_{j}\right)\right] \\
\alpha d\left(A_{j}, \bar{A}_{j}\right)<d\left(A_{j}, C_{i}\right) \leq \frac{1}{\left|A_{i}\right|}\left[\left|C_{i}\right| d\left(A_{j}, A_{i}\right)+\left|A_{j}\right| d\left(A_{i}, \bar{A}_{i}\right)\right] \tag{4}
\end{array}
$$

Divide Inequality (3) by $\left|A_{i}\right|$, divide Inequality (4) by $\left|A_{j}\right|$, and add them up:

$$
(\alpha-1)\left|\bar{A}_{j}\right| d_{\mathrm{a}}\left(A_{j}, \bar{A}_{j}\right)+(\alpha-1)\left|\bar{A}_{i}\right| d_{\mathrm{a}}\left(A_{i}, \bar{A}_{i}\right)<\left(\left|C_{i}\right|+\left|C_{j}\right|\right) d_{\mathrm{a}}\left(\bar{A}_{i}, \bar{A}_{j}\right)
$$

Since $\alpha,\left|\bar{A}_{j}\right|$ and $\left|\bar{A}_{i}\right|$ are large enough, $(\alpha-1)\left|\bar{A}_{i}\right|>\left|C_{i}\right|$ and $(\alpha-1)\left|\bar{A}_{j}\right|>\left|C_{j}\right|$. Then the claim follows.
(2) Suppose $p$ comes from the optimal cluster $C_{i}$. Let $q=\arg \min _{p^{\prime} \notin C_{i}} d\left(p, p^{\prime}\right)$, and suppose $q \in C_{j}$.

If $d_{\mathrm{a}}\left(p, C_{i}\right) \geq d_{\mathrm{a}}\left(q, C_{j}\right)$, then by Inequality (3),

$$
\begin{aligned}
\alpha d\left(p, C_{i}\right) & \leq\left|C_{j}\right| d(p, q)+d\left(q, C_{j}\right)=\left|C_{j}\right| d(p, q)+\left|C_{j}\right| d_{\mathrm{a}}\left(q, C_{j}\right) \\
& \leq\left|C_{j}\right| d(p, q)+\left|C_{j}\right| d_{\mathrm{a}}\left(p, C_{i}\right)
\end{aligned}
$$

which leads to $d_{\mathrm{a}}\left(p, C_{i}\right)<d(p, q) / 2$ since $\alpha$ is sufficiently large.
If $d_{\mathrm{a}}\left(p, C_{i}\right)<d_{\mathrm{a}}\left(q, C_{j}\right)$, then we have $d_{\mathrm{a}}\left(q, C_{j}\right)<d(p, q) / 2$ by a similar argument.
In conclusion, we always have $d_{\mathrm{a}}\left(p, C_{i}\right)<d(p, q) / 2$, which means that more than $\left|C_{i}\right| / 2$ points in $C_{i}$ are within distance less than $d(p, q)=\min _{p^{\prime} \notin C_{i}} d\left(p, p^{\prime}\right)$.

We are now ready to use Lemma 6 and Lemma 7 to prove the correctness of our theorem.

Proof of Theorem 6. It is sufficient to show that in Algorithm 4, $\mathcal{C}^{\prime}$ is always laminar to the optimal clustering $\mathcal{C}$, that is, for any $A \in \mathcal{C}^{\prime}$ and $C \in \mathcal{C}$, we have either $A \subseteq C$, or $C \subseteq A$, or $A \cap C=\varnothing$. Then the minimum cost pruning of $\mathcal{T}$ will be the optimal clustering, which can be obtained by dynamic programming.

Intuitively, Lemma $7(2)$ implies that $\mathcal{C}^{\prime}$ is laminar initially, and Lemma 7(1) can be used to show that the merge steps preserve the laminarity, so $\mathcal{C}^{\prime}$ is always laminar to the optimal clustering.

Formally, we prove the laminarity by induction. By Lemma $7(2), \mathcal{C}^{\prime}$ is laminar initially. It is sufficient to prove that if the current clustering is laminar, then the merge step keeps the laminiarity. Assume that our current clustering $\mathcal{C}^{\prime}$ is laminar to the optimal clustering. Consider a merge of two clusters $A$ and $A^{\prime}$. There are two cases when laminarity could fail to be satisfied after the merge: (1) they are strict subsets from different optimal clusters, i.e. $A \subsetneq C_{i}, A^{\prime} \subsetneq C_{j} \neq C_{i} ;(2) A$ is a strict subset of an optimal cluster $C_{i}$ and $A^{\prime}$ is the union of one or several other optimal cluster(s). By Lemma $7(1)$, the first case cannot happen. In the second case, for any $E$ that is
a subset of $C_{i} \backslash A$ in the current clustering, we have $d_{\mathrm{a}}(A, E) \geq d_{\mathrm{a}}\left(A, A^{\prime}\right)$. We know that $d_{\mathrm{a}}\left(A, C_{i} \backslash A\right)$ is a weighted average of the average distances between $A$ and the clusters that are subsets of $C_{i} \backslash A$ in the current clustering, so $d_{\mathrm{a}}\left(A, C_{i} \backslash A\right) \geq d_{\mathrm{a}}\left(A, A^{\prime}\right)$. Also, $d_{\mathrm{a}}\left(A, A^{\prime}\right)$ is a weighted average of the average distances between $A$ and the optimal clusters in $A^{\prime}$, so there must exist an optimal cluster $C_{j} \subseteq A^{\prime}$ such that $d_{\mathrm{a}}\left(A, C_{j}\right) \leq d_{\mathrm{a}}\left(A, A^{\prime}\right) \leq d_{\mathrm{a}}\left(A, C_{i} \backslash A\right)$. This means

$$
d\left(A, C_{j}\right) \leq \frac{\left|C_{j}\right|}{\left|C_{i} \backslash A\right|} d\left(A, C_{i} \backslash A\right) \leq \alpha d\left(A, C_{i} \backslash A\right)
$$

where the last inequality comes from $\alpha \geq 3 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$ and $\left|C_{i} \backslash A\right| \geq \min _{i}\left|C_{i}\right| / 2$. This contradicts Lemma 6. So the merge of the two clusters $A$ and $A^{\prime}$ will preserve the laminarity.

Running Time Finding the nearest neighbors for each point takes $O(n \log n)$ time, so the step of constructing components takes $O\left(n^{2} \log n\right)$ time. To compute average distances between clusters, we can record the size of each cluster, and $d\left(C_{i}^{\prime}, C_{j}^{\prime}\right)$ for any $C_{i}^{\prime}, C_{j}^{\prime}$ in the current clustering, and update $d\left(C_{i}^{\prime} \cup C_{j}^{\prime}, C_{l}^{\prime}\right)=d\left(C_{i}^{\prime}, C_{l}^{\prime}\right)+d\left(C_{j}^{\prime}, C_{l}^{\prime}\right)$ for any other cluster $C_{l}^{\prime}$ when merging $C_{i}^{\prime}$ and $C_{j}^{\prime}$. So the merge steps take $O\left(n^{3}\right)$ time. As dynamic programming takes $O\left(n^{3}\right)$ time, we can find the optimum clustering in $O\left(n^{3}\right)$ time.

### 2.4.1 Sublinear Time Algorithm

Here we provide a sublinear algorithm for a clustering instance $(X, d)$ that is $\alpha$ perturbation resilient to the min-sum objective. For simplicity, suppose the distances are normalized to $[0,1]$. Let $N=|X|$. Let $\rho=\frac{\min _{i}\left|C_{i}\right|}{N}$ denote the fraction of the points in the smallest optimal cluster, and $\eta=\min _{p \in X, 1 \leq i \leq k} d_{\mathrm{a}}\left(p, C_{i}\right)$ denote the minimum average distance between points and optimal clusters.

Our main result in this subsection is the following.

Theorem 7. Suppose the clustering instance $(X, d)$ is $\alpha$-perturbation resilient to

Algorithm 5 min-sum, $\alpha$ perturbation resilience, sublinear
Input: Data set $X$, distance function $d(\cdot, \cdot)$ on $X, \min _{i}\left|C_{i}\right|$.
1: Draw a sample $P$ of size $n=\Theta\left(\frac{1}{\rho^{2} \eta^{2}} \ln \frac{N k}{\delta}\right)$ i.i.d. from $X$.
2: Run Algorithm 4 on $P$ to obtain $\mathcal{C}$.
Output: The implicit clustering of $X$ obtained by assigning each point $p \in X$ to $\tilde{C}_{i} \in \tilde{\mathcal{C}}$ such that $d\left(p, \tilde{C}_{i}\right)$ is minimized.
the min-sum objective where $\alpha \geq 6 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$. Then with probability at least $1-\delta$, Algorithm 5 outputs an implicit optimum clustering in time poly $\left(\log \frac{N k}{\delta}, \frac{1}{\rho \eta}\right)$.

Proof. To prove the theorem, we first show the following (Lemma 8): with high probability, $\mathcal{C}^{\prime}$ in Algorithm 4 is always laminar to $\mathcal{C} \cap P$. The key idea is that when the sample is sufficiently large, we have that for any $p \in C_{i}$ and $C_{j}(j \neq i)$,

$$
3 \frac{\max _{i}\left|C_{i} \cap P\right|}{\min _{i}\left|C_{i} \cap P\right|} d\left(p, C_{i} \cap P\right)<d\left(p, C_{j} \cap P\right)
$$

since $\frac{d\left(p, C_{i} \cap P\right)}{n} \approx \frac{d\left(p, C_{i}\right)}{N}, \frac{d\left(p, C_{j} \cap P\right)}{n} \approx \frac{d\left(p, C_{j}\right)}{N}$ and $\frac{\max _{i}\left|C_{i} \cap P\right|}{\min _{i}\left|C_{i} \cap P\right|} \approx \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$. Then $\mathcal{C} \cap P$ satisfies the properties for the linkage in Algorithm 4 to succeed, and thus $\mathcal{C}^{\prime}$ in Algorithm 4 is always laminar to $\mathcal{C} \cap P$, i.e. $\mathcal{C} \cap P$ is a pruning of the tree.

Then we show that $\mathcal{C} \cap P$ is actually the minimum cost pruning $\tilde{\mathcal{C}}$ (Lemma 9). The key idea is that clusters in $\mathcal{C} \cap P$ are far apart, the cost increased by joining different clusters in it is larger than that saved by splitting clusters, so any other pruning has larger cost than $\mathcal{C} \cap P$. It immediately follows from the two lemmas that the implicit clustering obtained is the optimum clustering $\mathcal{C}$.

We now present the proofs of the lemmas for the correctness of the theorem.

Lemma 8. Suppose the clustering instance $(X, d)$ is $\alpha$-perturbation resilient to the min-sum objective where $\alpha \geq 6 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$. If the size of the sample $n=\Theta\left(\frac{1}{\rho^{2} \eta^{2}} \ln \frac{N k}{\delta}\right)$, then with probability at least $1-\delta, \mathcal{C}^{\prime}$ in Algorithm 4 is always laminar to $\mathcal{C} \cap P$.

Proof. The intuition is that on $X$, for any $i \neq j$, any $p \in C_{i}$, we have $\alpha d\left(p, C_{i}\right)<$ $d\left(p, C_{j}\right)$. When $n$ is sufficiently large, we can show $d\left(p, C_{i} \cap P\right) \approx \frac{n}{N} d\left(p, C_{i}\right)$ for
any $i$ and $\frac{\max _{i}\left|C_{i} \cap P\right|}{\min _{i}\left|C_{i} \cap P\right|} \approx \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$, and thus we have a similar claim on $P$. Then $\mathcal{C}^{\prime}$ in Algorithm 4 is always laminar to $\mathcal{C} \cap P$.

First, we show that with probability at least $1-\delta / 4$, for any $1 \leq i \leq k$ and $v=1 / 20$,

$$
\begin{equation*}
(1-v) \frac{n}{N}\left|C_{i}\right| \leq\left|C_{i} \cap P\right| \leq(1+v) \frac{n}{N}\left|C_{i}\right| . \tag{5}
\end{equation*}
$$

This follows from the union bound and

$$
\operatorname{Pr}\left[\left|\frac{\left|C_{i} \cap P\right|}{n}-\frac{\left|C_{i}\right|}{N}\right| \geq v \frac{\left|C_{i}\right|}{N}\right] \leq 2 \exp \left\{-2 v^{2} \frac{\left|C_{i}\right|^{2}}{N^{2}} n\right\} \leq 2 \exp \left\{-2 v^{2} \rho^{2} n\right\} \leq \frac{\delta}{4 k}
$$

A similar argument shows that with probability at least $1-\delta / 2$, for any $1 \leq i \leq k$ and $p \in X$,

$$
\begin{equation*}
(1-v) \frac{n}{N} d\left(p, C_{i}\right) \leq d\left(p, C_{i} \cap P\right) \leq(1+v) \frac{n}{N} d\left(p, C_{i}\right) \tag{6}
\end{equation*}
$$

Now, by (5), we have $\max _{i}\left|C_{i} \cap P\right| \leq(1+v) \frac{n}{N} \max _{i}\left|C_{i}\right|, \min _{i}\left|C_{i} \cap P\right| \geq(1-$ v) $\frac{n}{N} \min _{i}\left|C_{i}\right|$. Combined these with (6), we have that with probability at least $1-\delta$, for any $i \neq j$ and any $p \in C_{i}, 3 \frac{\max _{i}\left|C_{i} \cap P\right|}{\min _{i}\left|C_{i} \cap P\right|} d\left(p, C_{i} \cap P\right)<d\left(p, C_{j} \cap P\right)$, which guarantees the success of Algorithm 4.

Lemma 9. Suppose the clustering instance ( $X, d$ ) is $\alpha$-perturbation resilient to the min-sum objective where $\alpha \geq 6 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$. If the size of the sample $n=\Theta\left(\frac{1}{\rho^{2} \eta^{2}} \ln \frac{N k}{\delta}\right)$, then with probability at least $1-\delta$, the minimum min-sum cost pruning of the tree in Algorithm 4 is $\mathcal{C} \cap P$.

Proof. Since the tree is laminar to $\mathcal{C} \cap P$, we know that $\mathcal{C} \cap P$ is a pruning of the tree, and any other pruning can be obtained by splitting some clusters in $\mathcal{C} \cap P$ and joining some others into unions. Intuitively, the clusters in $\mathcal{C} \cap P$ are far apart, so the cost increased by joining different clusters is larger than the cost saved by splitting clusters. This claim then implies $\mathcal{C} \cap P$ is the minimum cost pruning. We first prove a
similar claim for $\mathcal{C}$ by the $\alpha$-perturbation resilience, i.e. for any three different clusters $C_{i}, C_{j}, C_{l} \in \mathcal{C}$, any $A_{X} \subseteq C_{i}, \alpha d\left(A_{X}, C_{i} \backslash A_{X}\right)<d\left(C_{j}, C_{l}\right)$. Then we prove the claim for $\mathcal{C} \cap P$ : for any $A \subseteq C_{i} \cap P, d\left(A, C_{i} \cap P \backslash A\right)<d\left(C_{j} \cap P, C_{l} \cap P\right) / 2$. Finally we use it to prove $\mathcal{C} \cap P$ is the minimum cost pruning.

First, for any $A_{X} \subseteq C_{i}$, we define a perturbation as follows: blow up the distances between the points in $A_{X}$ and those in $C_{i} \backslash A_{X}$ by a factor of $\alpha$, and keep all the other pairwise distances unchanged. By the $\alpha$-perturbation resilience, we know that $\mathcal{C}$ is still the optimum clustering after perturbation. Therefore, it has lower cost than the clustering obtained by replacing $C_{i}$ with $A_{X}$ and $C_{i} \backslash A_{X}$, and replacing $C_{j}$ and $C_{l}$ with $C_{j} \cup C_{l}$. After canceling the common terms in the costs of the two clusterings, we have $2 d^{\prime}\left(A_{X}, C_{i} \backslash A_{X}\right)<2 d^{\prime}\left(C_{j}, C_{l}\right)$, which leads to

$$
\alpha d\left(A_{X}, C_{i} \backslash A_{X}\right)<d\left(C_{j}, C_{l}\right)
$$

Second, we prove the following claim: for any $A \subseteq C_{i} \cap P$,

$$
2 d\left(A, C_{i} \cap P \backslash A\right) \leq 2 d\left(C_{i} \cap P, C_{i} \cap P\right)<d\left(C_{j} \cap P, C_{l} \cap P\right)
$$

On one hand, by summing over all the subsets of $C_{i}$, we have $\sum_{A_{X} \subseteq C_{i}} d\left(A_{X}, C_{i} \backslash A_{X}\right)=$ $2^{\left|C_{i}\right|} d\left(C_{i}, C_{i}\right) / 2$. Then $\frac{\alpha}{2} d\left(C_{i}, C_{i}\right)<d\left(C_{j}, C_{l}\right)$. On the other hand, similar to the proof of Lemma 8, we can show that with high probability, for any $p \in C_{i}, d\left(p, C_{i} \cap P\right) \leq$ $(1+v) \frac{n}{N} d\left(p, C_{i}\right)$ for $v=1 / 20$. So we have

$$
\begin{aligned}
d\left(C_{i} \cap P, C_{i} \cap P\right) & =\sum_{p \in C_{i} \cap P} d\left(p, C_{i} \cap P\right) \leq(1+v) \frac{n}{N} \sum_{p \in C_{i} \cap P} d\left(p, C_{i}\right) \\
& =(1+v) \frac{n}{N} \sum_{q \in C_{i}} d\left(C_{i} \cap P, q\right) \leq(1+v)^{2} \frac{n^{2}}{N^{2}} d\left(C_{i}, C_{i}\right) .
\end{aligned}
$$

A similar argument shows that for any $C_{j}$ and $C_{l}, d\left(C_{j} \cap P, C_{l} \cap P\right) \geq(1-v)^{2} \frac{n^{2}}{N^{2}} d\left(C_{j}, C_{l}\right)$. The claim then follows by combining the three inequalities and noting $v=1 / 20$.

Now, we use the claim to prove the optimality of $\mathcal{C} \cap P$. Suppose a pruning $\mathcal{P}^{*}$ is obtained by splitting $h$ clusters in $\mathcal{C} \cap P$ and at the same time joining some
other clusters into $g$ unions. Specifically, for $1 \leq i \leq h$, split $C_{i} \cap P$ into $m_{i} \geq 2$ clusters $P_{i, 1}, \ldots, P_{i, m_{i}}$; after that, merge $C_{h+1} \cap P, \ldots, C_{h+l_{g}} \cap P$ into $g$ unions, i.e. for $1 \leq j \leq g, l_{0}=0$, merge $l_{j}-l_{j-1} \geq 2$ clusters $C_{h+l_{j-1}+1} \cap P, \ldots, C_{h+l_{j}} \cap P$ into a union $U_{j}$; the other clusters in $\mathcal{C} \cap P$ remain the same in $\mathcal{P}^{*}$. Since the number of clusters is still $k$, we have $\sum_{i} m_{i}-h=l_{g}-g$. The cost saved by splitting $h$ clusters is

$$
\begin{equation*}
\sum_{1 \leq i \leq h} \sum_{1 \leq p \neq q \leq m_{i}} d\left(P_{i, p}, P_{i, q}\right)=\sum_{1 \leq i \leq h} \sum_{1 \leq p \leq m_{i}} d\left(P_{i, p}, C_{i} \cap P \backslash P_{i, p}\right) . \tag{7}
\end{equation*}
$$

The cost increased by joining clusters is

$$
\begin{equation*}
\sum_{1 \leq j \leq g} \sum_{h+l_{j-1}<p \neq q \leq h+l_{j}} d\left(C_{p} \cap P, C_{q} \cap P\right) . \tag{8}
\end{equation*}
$$

To prove $\mathcal{C} \cap P$ is the minimum cost pruning, we need to show that the saved cost (7) is less than the increased cost (8). Since each term in (8) is twice larger than any term in (7), it suffices to show that the number of the terms in (8) is at least half the number of the terms in (7). Formally, we need to show $2 \sum_{1 \leq j \leq g}\binom{l_{j}-l_{j-1}}{2} \geq \sum_{1 \leq i \leq h} m_{i}$. We have $2 \sum_{j}\binom{l_{j}-l_{j-1}}{2}=\sum_{j}\left(l_{j}-l_{j-1}\right)\left(l_{j}-l_{j-1}-1\right) \geq 2 \sum_{j}\left(l_{j}-l_{j-1}-1\right)=2\left(l_{g}-g\right)$, where the inequality comes from $l_{j}-l_{j-1} \geq 2$. Since $l_{g}-g=\sum_{i} m_{i}-h$, it is sufficient to show $l_{g}-g \geq h$. This comes from $l_{g}-g=\sum_{i} m_{i}-h=\sum_{i}\left(m_{i}-1\right) \geq \sum_{i} 1=h$ since $m_{i} \geq 2$.

## $2.5(\alpha, \epsilon)$-Perturbation Resilience for Min-Sum

For $(\alpha, \epsilon)$-perturbation resilient min-sum instances, we will show that when $\alpha=$ $\Omega\left(\frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}\right), \epsilon=\tilde{O}\left(\frac{\min _{i}\left|C_{i}\right|}{n}\right)$, there exists a polynomial time algorithm that outputs a clustering that is both a good approximation and also $\tilde{O}(\epsilon)$-close to the optimal clustering. Formally,

Theorem 8. Suppose the instance is ( $\alpha, \epsilon$ )-perturbation resilient to the min-sum objective for $\alpha>\frac{8 \max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$, $\epsilon<\frac{\min _{i}\left|C_{i}\right|}{600 n \log n}$. There exists an algorithm that outputs $a$
clustering which is a $\left(1+\frac{40 \epsilon n \log n}{\min _{i}\left|C_{i}\right|}\right)$-approximation to the optimal clustering in polynomial time. Furthermore, the output clustering is also $(6 \epsilon \log n)$-close to the optimal clustering.

Since $\epsilon=O\left(\frac{\min _{i}\left|C_{i}\right|}{n \log n}\right)$, the approximation factor is always $O(1)$ and gets better if $\epsilon$ gets smaller. To prove the theorem, we first derive new useful structural properties implied by $(\alpha, \epsilon)$-perturbation resilience for min-sum, and then use them to design our algorithm achieving the guarantees in the theorem.

### 2.5.1 Structure of $(\alpha, \epsilon)$-Perturbation Resilient Instances

In this subsection, we introduce the notion of bad points for the min-sum objective, and then show that there are just a few bad points while the other points have useful properties. More precisely,

Definition 7. Define bad points to be those that are not $\beta$ times closer to its own cluster than to other clusters, where $\beta=\min \left\{\frac{4}{5} \alpha, 8 n\right\}$. That is

$$
\begin{equation*}
B_{i}=\left\{p \in C_{i}: \exists j, d\left(p, C_{j}\right) \leq \beta d\left(p, C_{i}\right)\right\}, \beta=\min \left\{\frac{4}{5} \alpha, 8 n\right\}, B=\cup_{i} B_{i} . \tag{9}
\end{equation*}
$$

The other points are called good points.
$\beta$ is chosen to be $\min \left\{\frac{4}{5} \alpha, 8 n\right\}$ for reasons that will become clear in the proof bounding the number of bad points. Informally, we will show that when $\alpha$ is sufficiently large and $\epsilon$ is sufficiently small, the number of bad points are bounded by $\tilde{O}(\epsilon n)$ (Theorem 9 in Section 2.5.1.1). The good points in different optimal clusters, by definition, are far from each other. It is then possible to design approximation algorithms if the influence of the few bad points can be eliminated.

However, we do not know the actual bad points. A key observation is that we can introduce a proxy called potential bad points, which can be easily computed. Formally,

Definition 8. (a) Define $m_{B}:=6 \epsilon \log n$.
(b) For a set $A$ with $|A|>2 m_{B}$, define the potential bad points $F(A)$ to be the $2 m_{B}$ points in $A$ that are farthest from $A$. That is, $F(A) \subseteq A,|F(A)|=2 m_{B}$, and for any $p \in F(A), q \in P \backslash F(A), d(p, A) \geq d(q, A)$. The other points $P(A)=A \backslash F(A)$ are called potential good points.
(c) For a cluster $A$, define its robust min-sum cost as $d_{\mathrm{rs}}(A):=d(P(A), P(A))$, where $P(A)$ are the potential good points in $A$. For a clustering $\mathcal{C}$, define its robust min-sum cost as $\sum_{C \in \mathcal{C}} d_{\mathrm{rs}}(C)$.

We show that the robust min-sum cost computed after removing the potential bad points approximates the min-sum cost computed after removing the actual bad points (see Section 2.5.1.2). In other words, we can design approximation algorithms using the potential bad points as if we knew the actual bad points.

### 2.5.1.1 Bounding the Number of Bad Points

Here we bound the number of bad points by $\tilde{O}(\epsilon n)$ when $\alpha$ is sufficiently large and $\epsilon$ is sufficiently small.

Theorem 9. Suppose the clustering instance is $(\alpha, \epsilon)$-perturbation resilient for the min-sum objective where $\alpha>4$ and $\epsilon<\frac{\min _{i}\left|C_{i}\right|}{200 n}$. Then we have $|B| \leq m_{B}=6 \epsilon n \log n$. Proof. Assume for contradiction $|B|>2 \eta \epsilon n$ where $\eta=\left\lceil\log \frac{\max _{i} \max _{p \in B_{i}} d\left(p, C_{i}\right)}{\min _{i} \min _{p \in B_{i}} d\left(p, C_{i}\right)}\right\rceil$. We will first construct a perturbation which leads to a contradiction, and then show that $\eta \leq 3 \log n$, completing the proof.

We begin constructing the perturbation by introducing some notations. Consider the $\eta$ intervals as follows: $\left[2^{t-1} v, 2^{t} v\right]$ where $v=\min _{i} \min _{p \in B_{i}} d\left(p, C_{i}\right), 1 \leq t \leq \eta$. At least one of the intervals, say $[r, 2 r]$, will contain the costs of more than $2 \epsilon n$ bad points. Let $\hat{B}$ denote an arbitrary subset of $2 \epsilon n$ bad points in this interval. Let $\hat{B}_{i}=\hat{B} \cap C_{i}$ denote the selected bad points in the optimal cluster $C_{i}$. Let $K_{i}=C_{i} \backslash \hat{B}_{i}$ denote the other points in $C_{i}$, and set $K=\cup_{i} K_{i}$. Denote as $D_{j}$ all those selected bad points


Figure 5: Perturbation construction.
whose second nearest cluster is $C_{j}$, that is, $D_{j}=\left\{p: p \in \hat{B}_{i}, j=\arg \min _{\ell \neq i} d\left(p, C_{\ell}\right)\right\}$. Finally, let $\tilde{C}_{i}=K_{i} \cup D_{i}$. See Figure 5 for an illustration.

Now we are ready to construct the perturbation, which tries to make the selected bad points move to their second nearest clusters and keep the other points in their original clusters. More precisely, the perturbation is constructed as follows: blow up all distances by a factor of $\alpha$ except those within $\tilde{C}_{i}, 1 \leq i \leq k$. Intuitively, this perturbation favors the clustering $\tilde{C}_{i}$.

To derive a contradiction, consider the optimal clustering after perturbation, denoted as $\left\{C_{i}^{\prime}\right\}$. Since there are more than $\epsilon n$ bad points in $\hat{B}$, by $(\alpha, \epsilon)$-perturbation resilience, not all of them move to new clusters in $\left\{C_{i}^{\prime}\right\}$, and thus $\left\{C_{i}^{\prime}\right\}$ is different from $\left\{\tilde{C}_{i}\right\}$. In fact, we will show that the clustering $\left\{\tilde{C}_{i}\right\}$ has a lower cost than $\left\{C_{i}^{\prime}\right\}$, which is a contradiction. To do so, we consider changing $\left\{C_{i}^{\prime}\right\}$ to $\left\{\tilde{C}_{i}\right\}$ by moving points. It is sufficient to show that by moving these points, the cost saved is larger than the cost added.

To bound the costs, we first divide the points into different types. See Figure 6 for an illustration. First, we need to move out $C_{i}^{\prime} \backslash \tilde{C}_{i}$ from each $C_{i}^{\prime}$. These points can be divided into three types:
(1) $U_{i}=C_{i}^{\prime} \cap \hat{B}_{i}$ are the selected bad points in $C_{i}$ that need to be moved out.
(2) $V_{i}=\left(C_{i}^{\prime} \backslash \tilde{C}_{i}\right) \cap\left(\cup_{j \neq i} \hat{B}_{j}\right)=\cup_{j \neq i}\left(\hat{B}_{j} \cap C_{i}^{\prime}\right)$ are the selected bad points that are


Figure 6: Different types of points when bounding the number of bad points for $(\alpha, \epsilon)$-perturbation resilient min-sum instances.
from other optimal clusters. But their second nearest cluster is not $C_{i}$, so they are not in $\tilde{C}_{i}$.
(3) $W_{i}=\left(C_{i}^{\prime} \backslash \tilde{C}_{i}\right) \cap\left(\cup_{j \neq i} K_{j}\right)=\cup_{j \neq i}\left(K_{j} \cap C_{i}^{\prime}\right)$ are points that are from $K_{j}$ for some $j \neq i$ and are in $C_{i}^{\prime}$. But they are not from $K_{i}$ and thus are not in $\tilde{C}_{i}$.

Second, we need to move in $\tilde{C}_{i} \backslash C_{i}^{\prime}$ for each $\tilde{C}_{i}$. Similarly, these points can also be divided into three types:
(1) $\tilde{W}_{i}=K_{i} \backslash C_{i}^{\prime}=\cup_{j \neq i}\left(K_{i} \cap C_{j}^{\prime}\right)$ are those points in $K_{i}$ and in $C_{j}^{\prime}$ for some $j \neq i$. This means that they are points in $\cup_{j} W_{j}$. More specifically, we have $\cup_{i} \tilde{W}_{i}=\cup_{j} W_{j}$.
(2) $\tilde{V}_{i}=\left(D_{i} \backslash C_{i}^{\prime}\right) \cap\left[\cup_{\ell \neq j}\left(\hat{B}_{\ell} \cap C_{j}^{\prime}\right)\right]$ are part of the selected bad points whose second nearest cluster is $C_{i}$. They are originally in $\hat{B}_{\ell}$ for some $\ell$ but are in $C_{j}^{\prime}$ for some $j \neq \ell$. In other words, they are points from $V_{j}$ for some $j$, and we have $\cup_{i} \tilde{V}_{i}=\cup_{j} V_{j}$.
(3) $\tilde{U}_{i}=\left(D_{i} \backslash C_{i}^{\prime}\right) \cap\left[\cup_{j}\left(\hat{B}_{j} \cap C_{j}^{\prime}\right)\right]$ are also part of the selected bad points whose second nearest cluster is $C_{i}$. They are originally in $\hat{B}_{j}$ for some $j$ and are also in $C_{j}^{\prime}$. In other words, they are points from $U_{j}$ for some $j$, and we have $\cup_{i} \tilde{U}_{i}=\cup_{j} U_{j}$.

In conclusion, we have $C_{i}^{\prime}=C_{i}^{\prime} \cap \tilde{C}_{i}+U_{i}+V_{i}+W_{i}, \tilde{C}_{i}=C_{i}^{\prime} \cap \tilde{C}_{i}+\tilde{U}_{i}+\tilde{V}_{i}+\tilde{W}_{i}$. We also have $\cup_{i} \tilde{U}_{i}=\cup_{j} U_{j}, \cup_{i} \tilde{V}_{i}=\cup_{j} V_{j}$, and $\cup_{i} \tilde{W}_{i}=\cup_{j} W_{j}$. The costs saved and added are then summarized as follows. Suppose we first move out $W$, then $V$, and finally $U$. The cost saved by moving out $W$ is at least $2 \sum_{i} d^{\prime}\left(W_{i}, C_{i}^{\prime} \cap C_{i}\right)$, that saved by moving out $V$ is at least $2 \sum_{i} d^{\prime}\left(V_{i}, C_{i}^{\prime} \cap C_{i}\right)$, and that saved by moving out $U$ is at least $2 \sum_{i} d^{\prime}\left(U_{i}, C_{i}^{\prime} \cap K_{i}\right)$. Next, we move in $\tilde{W}$, then $\tilde{V}$, and finally $\tilde{U}$. The cost added by moving in $\tilde{W}$ is at most $2 \sum_{i} d^{\prime}\left(\tilde{W}_{i}, \tilde{W}_{i}+C_{i}^{\prime} \cap \tilde{C}_{i}\right)$, that added by moving in $\tilde{V}$ is at most $2 \sum_{i} d^{\prime}\left(\tilde{V}_{i}, \tilde{C}_{i}\right)$, and that added by moving in $\tilde{U}$ is at most $2 \sum_{i} d^{\prime}\left(\tilde{U}_{i}, \tilde{C}_{i}\right)$.

We are now ready to show that the cost saved is greater than the cost added. The high level idea is that a significant amount of cost is saved by moving $U_{i}$ to the correct clusters, while the costs added by moving $V_{i}$ and $W_{i}$ are generally small since the number of points moved is bounded by $3 \epsilon n$ and the cost of the selected bad points moved is bounded by $2 r$. Formally, we have the following claim, whose proof is presented in Appendix A.4.1.

Claim 3. The costs saved and added by moving $\left\{U_{i}\right\}_{i=1}^{k},\left\{V_{i}\right\}_{i=1}^{k}$ and $\left\{W_{i}\right\}_{i=1}^{k}$ satisfy:
(a) $\quad 2 \sum_{i} d^{\prime}\left(U_{i}, C_{i}^{\prime} \cap K_{i}\right)-2 \sum_{i} d^{\prime}\left(\tilde{U}_{i}, \tilde{C}_{i}\right)$

$$
\begin{equation*}
\geq \frac{3}{10} \alpha \sum_{i} d\left(U_{i}, C_{i}\right)-\frac{2 \alpha}{100} \sum_{i} d\left(W_{i}, C_{i}\right)-\frac{8 \alpha+16}{100} r \epsilon n \tag{b}
\end{equation*}
$$

$$
\begin{align*}
\geq & \frac{99}{50}(\alpha-2) \sum_{i} d\left(V_{i}, C_{i}\right)-\frac{2 \alpha}{100} \sum_{i} d\left(W_{i}, C_{i}\right)-\frac{8 \alpha+16 \beta}{100} r \epsilon n, \\
& 2 \sum_{i} d^{\prime}\left(W_{i}, C_{i}^{\prime} \cap C_{i}\right)-2 \sum_{i} d^{\prime}\left(\tilde{W}_{i}, \tilde{W}_{i}+C_{i}^{\prime} \cap \tilde{C}_{i}\right)  \tag{c}\\
\geq & \frac{98}{50}(\alpha-2) \sum_{i} d\left(W_{i}, C_{i}\right)-\frac{8 \alpha+8 \beta}{100} r \epsilon n .
\end{align*}
$$

After adding up all the inequalities in the claim, the right hand side is a lower bound on the difference between the cost saved and the cost added, which we now show must be positive when $\alpha>4$ and $\beta \leq \frac{4}{5} \alpha$. The terms about $d\left(W_{i}, C_{i}\right)$ and
$d\left(V_{i}, C_{i}\right)$ are positive, so it suffices to show that $\sum_{i} d\left(U_{i}, C_{i}\right)$ is larger than $r \epsilon n$. First, $d\left(p, C_{i}\right) \geq r$ for any $p \in U_{i}$. Second, $\left|\cup_{i} U_{i}\right| \geq \epsilon n$ since there are $2 \epsilon n$ selected bad points but no more than $\epsilon n$ of them move from their original clusters in $\left\{C_{i}\right\}$ to a different cluster in $\left\{C_{i}^{\prime}\right\}$. Then we have $\sum_{i} d\left(U_{i}, C_{i}\right) \geq \sum_{i} r\left|U_{i}\right|=r \sum_{i}\left|U_{i}\right| \geq r \in n$. Hence, the difference between the cost saved and the cost added is positive. This means the cost of $\left\{\tilde{C}_{i}\right\}$ is smaller than the cost of $\left\{C_{i}^{\prime}\right\}$, which contradicts the assumption that $\left\{C_{i}^{\prime}\right\}$ is the optimal clustering under $d^{\prime}$. Therefore, there can be at most $2 \eta \epsilon n$ bad points.

Finally, what is left is to show $\eta \leq 3 \log n$. Suppose $p_{1}$ is the point that achieves $\max _{i} \max _{p \in B_{i}} d\left(p, C_{i}\right)$ and $p_{2}$ is the point that achieves $\min _{i} \min _{p \in B_{i}} d\left(p, C_{i}\right)$. Without loss of generality, suppose $p_{1} \in C_{1}$ and $p_{2} \in C_{2}$. By definition of bad points, there exists $C_{i} \neq C_{2}$ such that $d\left(p_{2}, C_{i}\right) \leq \beta d\left(p_{2}, C_{2}\right)$. If $C_{i} \neq C_{1}$, we have $d\left(p_{1}, C_{1}\right) \leq$ $d\left(C_{2}, C_{i}\right)$, since otherwise we can get lower cost by splitting $C_{1}$ into $p_{1}$ and $C_{1} \backslash\left\{p_{1}\right\}$ while merging $C_{2}$ and $C_{i}$. If $C_{i}=C_{1}$, we also have $d\left(p_{1}, C_{1}\right) \leq d\left(C_{2}, C_{i}\right)$, since otherwise we can get lower cost by splitting $C_{1}$ into $p_{1}$ and $C_{1} \backslash\left\{p_{1}\right\}$ and then merging $C_{2}$ and $C_{1} \backslash\left\{p_{1}\right\}$. In both cases, we have

$$
\begin{aligned}
d\left(p_{1}, C_{1}\right) \leq d\left(C_{2}, C_{i}\right) & \leq\left|C_{i}\right| d\left(p_{2}, C_{2}\right)+\left|C_{2}\right| d\left(p_{2}, C_{i}\right) \\
& \leq\left|C_{i}\right| d\left(p_{2}, C_{2}\right)+\beta\left|C_{2}\right| d\left(p_{2}, C_{2}\right) \\
& \leq 8 n^{2} d\left(p_{2}, C_{2}\right)
\end{aligned}
$$

where the last inequality follows from $\beta \leq 8 n$. Then we have $\eta \leq 3 \log n$.

### 2.5.1.2 Properties of Actual and Potential Good Points

Since there are just a few bad points and the good points in different clusters are far apart, the cost between sufficiently large subsets of their good points accounts for most of the cost between the two clusters. This means that the min-sum cost can be approximately computed on the good points, and the min-sum clustering can be approximately solved if we knew the actual good points.

Lemma 10. Suppose $W_{i} \subseteq G_{i}, W_{j} \subseteq G_{j}$. When $\left|C_{i}\right| \geq 50\left|C_{i} \backslash W_{i}\right|$ and $\left|C_{j}\right| \geq$ $50\left|C_{j} \backslash W_{j}\right|$, we have $d\left(C_{i}, C_{j}\right) \leq \frac{3}{2} d\left(W_{i}, W_{j}\right)$.

Proof Sketch. Since $\left|W_{i}\right| \approx\left|C_{i}\right|$ and $\left|W_{j}\right| \approx\left|C_{j}\right|$, it suffices to show that $d_{\mathrm{a}}\left(C_{i}, C_{j}\right) \leq$ $d_{\mathrm{a}}\left(W_{i}, W_{j}\right)$ approximately. By the triangle inequality, $d_{\mathrm{a}}\left(C_{i}, C_{j}\right) \leq d_{\mathrm{a}}\left(C_{i}, W_{i}\right)+$ $d_{\mathrm{a}}\left(W_{i}, W_{j}\right)+d_{\mathrm{a}}\left(W_{j}, C_{j}\right)$, so we only need to bound $d_{\mathrm{a}}\left(W_{i}, C_{i}\right)$ and $d_{\mathrm{a}}\left(W_{j}, C_{j}\right)$ by $d_{\mathrm{a}}\left(W_{i}, W_{j}\right)$.

By definition of good points, $d_{\mathrm{a}}\left(W_{i}, C_{i}\right)$ is much less than $d_{\mathrm{a}}\left(W_{i}, C_{j}\right)$, which is approximately $d_{\mathrm{a}}\left(W_{i}, W_{j}\right)$ since $W_{j}$ takes up a majority of points in $C_{j}$. A similar argument bounds $d_{\mathrm{a}}\left(W_{j}, C_{j}\right)$, which then leads to the lemma. The complete proof is provided in Appendix A.4.

Furthermore, the good points in different optimal clusters are far apart in the following sense: the good points from two different clusters have cost much larger than those in a third cluster have. Formally,

Lemma 11. For any three different optimal clusters $C_{i}, C_{j}$, and $C_{l}$, and any $A \subset G_{i}$, $\frac{18}{5} d\left(A, G_{i} \backslash A\right)<d\left(G_{j}, G_{l}\right)$. Consequently, $\frac{9}{5} d\left(G_{i}, G_{i}\right)<d\left(G_{j}, G_{l}\right)$.

If we can construct a tree on the good points with a pruning that assigns all good points correctly, then this property is useful in finding the pruning.

Now we turn to analyze the potential good points. A key property of the potential good points is the following: for any point $p$ and any sufficiently large set $A$, the cost between $p$ and the potential good points in $A$ is roughly bounded by the cost between $p$ and any sufficiently large subset $H$ of $A$. See Lemma 12 for details. A specific case is when $H$ is the actual good points in $A$. In this case, the property says that the cost between $p$ and the potential good points is roughly bounded by the cost between $p$ and the actual good points. This means that in suitable situations, we can regard potential good points as actual good points.


Figure 7: Properties of the potential good points.

Lemma 12. Suppose $H \subseteq A$ such that $|A \backslash H| \leq m_{B}$. Let $F=F(A), P=P(A), \bar{H}=$ $A \backslash H$. Let $W=H \cap P, V=\bar{H} \cap F, X=F \cap H, Y=\bar{H} \cap P$. See Figure 7 for an illustration. If $|A| \geq 20 m_{B}$, then for any $p, d(p, P) \leq \frac{|W|+|Y|}{|W|-|X|} d(p, H)$.

Proof. Since $d(p, P)=d(p, Y)+d(p, W)$ and $d(p, H)=d(p, X)+d(p, W)$, the lemma is true if $Y=\emptyset$. Otherwise, we need to compare $d(p, X)$ and $d(p, Y)$. By the triangle inequality, we have $d_{\mathrm{a}}(W, X) \leq d_{\mathrm{a}}(W, p)+d_{\mathrm{a}}(p, X)$ and $d_{\mathrm{a}}(p, Y) \leq d_{\mathrm{a}}(p, W)+$ $d_{\mathrm{a}}(W, Y)$. Then

$$
d(p, X) \geq \frac{d(W, X)}{|W|}-\frac{|X|}{|W|} d(p, W), \quad d(p, Y) \leq \frac{d(W, Y)}{|W|}+\frac{|Y|}{|W|} d(p, W)
$$

From these bounds on $d(p, X)$ and $d(p, Y)$, we have $d(p, H) \geq \frac{d(W, X)}{|W|}+\frac{|W|-|X|}{|W|} d(p, W), \quad d(p, P) \leq \frac{d(W, Y)}{|W|}+\frac{|W|+|Y|}{|W|} d(p, W)$.

The lemma then follows from these two inequalities and the following claim.

Claim 4. $d(X, W) \geq d(Y, W+Y)$.

Intuitively, the points in $X \subseteq F$ are among those farthest away from $A$, so $d_{\mathrm{a}}(X, A) \geq d_{\mathrm{a}}(Y, A)$. Since $|A \backslash H| \leq m_{B}$ and $|F|=2 m_{B}$, we have $|X| \geq 2|Y|$. Then the cost of $Y$ cannot be too large compared to that of $X$. The complete proof of the claim is provided in Appendix A.4.

### 2.5.2 Approximating the Optimal Clustering

In this subsection, we design an approximation algorithm and prove our final result Theorem 8 by utilizing the properties of the $(\alpha, \epsilon)$-perturbation resilience. Note that we can generate a list of sufficiently large almost "pure" blobs using Algorithm 3. However, unlike for $(\alpha, \epsilon)$-perturbation resilient $k$-median instances, it is not guaranteed that the robust linkage procedure in [16] can link these blobs into a tree so that a pruning of the tree assigns all but bad points correctly. In Section 2.5.2.1, we design a robust average linkage algorithm to achieve this goal, and in Section 2.5.2.2, we show that this pruning can be found in polynomial time. Although this pruning may not be a good approximation to the optimum, in Section 2.5.2.3 we show that a good approximation can be computed by reassigning points in it, which leads to Theorem 8.

### 2.5.2.1 Constructing A Tree with A Pruning Close to the Optimal Clustering

Here we show that by utilizing the bound on the number of bad points, we can construct a tree with a pruning that assigns all good points correctly. As described in Algorithm 6, we first use Algorithm 3 to generate a list of blobs, and then use a robust version of average linkage to link them into a tree: repeatedly merge the two blobs with the minimum robust average distance defined as follows.

Definition 9. The robust average distance $d_{r a}\left(A_{1}, A_{2}\right)$ between two sets $A_{1}, A_{2}$ is defined as the average distance between their potential good points. That is, $d_{r a}\left(A_{1}, A_{2}\right)=$ $\frac{d\left(P\left(A_{1}\right), P\left(A_{2}\right)\right)}{\left|P\left(A_{1}\right)\right|\left|P\left(A_{2}\right)\right|}$.

We now prove that the tree output by Algorithm 6 has a pruning that correctly assigns all good points.

Lemma 13. The tree output in Algorithm 6 has a pruning $\mathcal{C}^{\prime}$ that assigns all good points correctly.

```
Algorithm 6 Robust Average Linkage
Input: Data set \(P\), distance function \(d(\cdot, \cdot)\) on \(P, \min _{i}\left|C_{i}\right|, \epsilon>0\).
    1: Use Algorithm 3 with \(m_{B}=6 \epsilon n \log n\) and \(\gamma=\frac{4}{5}\) to get a list \(\mathcal{L}_{0}\) of blobs.
    2: Initialize the clustering \(\mathcal{L}\) with each blob being a cluster.
    3: Repeat till only one cluster remains:
    merge clusters \(C, C^{\prime}\) which minimize \(d_{r a}\left(C, C^{\prime}\right)\).
    4: Let \(\mathcal{T}\) be the tree with blobs as leaves and internal nodes corresponding to the
    merges performed.
Output: The tree \(\mathcal{T}\).
```

Proof. To analyze the algorithm, we begin with the following property of good points.
When combined with the property of Algorithm 3 (Lemma 4), it immediately shows that each blob in the list $\mathcal{L}_{0}$ has size at least $\frac{1}{2} \min _{i}\left|C_{i}\right|$, and contains good points from only one optimal cluster.

Claim 5. For any $p \in G_{i}$, all its $\frac{4\left|C_{i}\right|}{5}$ nearest neighbors belong to $C_{i} \cup B$.

Proof. We need to show that for any $j \neq i$ and any good point $q \in G_{j}, d(p, q)$ is large compared to $d_{\mathrm{a}}\left(p, C_{i}\right)$. Intuitively, $p$ is much farther away from $C_{j}$ than from $C_{i}: \beta d\left(p, C_{i}\right) \leq d\left(p, C_{j}\right)$. It suffices to bound $d\left(p, C_{j}\right)$ by $d(p, q)$ and $d\left(p, C_{i}\right)$. By the triangle inequality, we have
$d\left(p, C_{j}\right) \leq\left|C_{j}\right| d(p, q)+d\left(q, C_{j}\right) \quad$ and $\quad d\left(q, C_{j}\right) \leq \frac{1}{\beta} d\left(q, C_{i}\right) \leq \frac{\left|C_{i}\right|}{\beta} d(p, q)+\frac{1}{\beta} d\left(p, C_{i}\right)$.
Combining these inequalities, we have $\left(\beta-\frac{1}{\beta}\right) d\left(p, C_{i}\right) \leq\left(\left|C_{j}\right|+\frac{\left|C_{i}\right|}{\beta}\right) d(p, q)$. When $\alpha>8 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$, we have $5 d_{\mathrm{a}}\left(p, C_{i}\right)<d(p, q)$, which then leads to the conclusion.

It now suffices to prove by induction that the clustering $\mathcal{L} \cap G$ is always laminar to $\mathcal{C} \cap G$. It is true at the beginning by the property of Algorithm 3. Assume for contradiction that the laminarity is first violated after merging $A$ and $D$. There are two cases: (1) $A$ and $D$ are strict subsets of different optimal clusters; (2) $A$ is a strict subset of $G_{i}$ while $D$ is the union of the good points in several optimal clusters. We have the following claims for the two cases respectively.

Claim 6. (a) Suppose $A \in \mathcal{L}, A \cap G \subsetneq G_{i}$, and $D \in \mathcal{L}, D \cap G \subsetneq G_{j}(j \neq i)$. Then there exists $A^{\prime} \neq A$ in $\mathcal{L}$ such that $A^{\prime} \cap G \subsetneq G_{i}$ and $d_{r a}\left(A, A^{\prime}\right)<d_{r a}(A, D)$.
(b) Suppose $A \in \mathcal{L}, A \cap G \subsetneq G_{i}$, and $D \in \mathcal{L}, D \cap G$ is the union of good points in several optimal clusters. Then there exists $A^{\prime} \neq A$ in $\mathcal{L}$ such that $A^{\prime} \cap G \subsetneq G_{i}$ and $d_{r a}\left(A, A^{\prime}\right)<d_{r a}(A, D)$.

The intuition is that since good points in different optimal clusters are far away, we can find a blob $A^{\prime}$ from $C_{i} \cup B$ such that good points in $A^{\prime}$ have smaller average distance to good points in $A$ than those in $D$ have, that is, $d_{\mathrm{a}}\left(A \cap G, A^{\prime} \cap G\right)<$ $\frac{1}{2} d_{\mathrm{a}}(A \cap G, D \cap G)$. Then we show that the bad points in these blobs do not change things much: $d_{r a}\left(A, A^{\prime}\right)$ is approximately less than $d_{\mathrm{a}}\left(A \cap G, A^{\prime} \cap G\right)$, and $d_{r a}(A, D)$ is approximately greater than $d_{\mathrm{a}}(A \cap G, D \cap G)$. These then lead to the claims. Their complete proofs are presented in Appendix A.5.1.

By these two claims, we should first merge $A$ with $A^{\prime}$ rather than with $D$, which is contradictory. So the laminarity is always preserved, which completes the proof.

### 2.5.2.2 Getting A Pruning Close to the Optimal Clustering

We have shown that a tree can be constructed such that there is a pruning, denoted as $\mathcal{C}^{\prime}$, that assigns all good points correctly. Here we show how to find this pruning from the tree. Suppose we can remove the actual bad points and compute the cost between the good points. Since the good points from different clusters are far apart, the good point cost increased by joining different clusters in $\mathcal{C}^{\prime}$ is larger than that saved by splitting clusters in $\mathcal{C}^{\prime}$ (Lemma 11). Then any other pruning has larger cost than $\mathcal{C}^{\prime}$. Unfortunately, we do not know the actual good points. Therefore, we consider the potential good points and compute the robust min-sum cost. It turns out that $\mathcal{C}^{\prime}$ indeed is the pruning with the minimum robust min-sum cost.

Lemma 14. Suppose the pruning $\mathcal{C}^{\prime}=\left\{C_{1}^{\prime}, \ldots, C_{k}^{\prime}\right\}$ in tree $\mathcal{T}$ assigns all good points correctly. Then $\mathcal{C}^{\prime}$ is the minimum robust min-sum cost pruning in the tree.

Proof Sketch. Computing the robust min-sum cost will eliminate the effect of the bad points and work as if we knew the actual good points: the robust min-sum cost saved by splitting a node is at most the good point cost saved (Claim 7), and the robust min-sum cost increased by merging two nodes is in the same order of the good point cost increased (Claim 8).


Figure 8: Notations in Claim 4 and 7.

Claim 7. If $\left|C_{i}^{\prime}\right| \geq 20 m_{B}$, then $d_{\mathrm{rs}}\left(C_{i}^{\prime}\right) \leq d\left(G_{i}, G_{i}\right)$.

Proof. The claim follows from Claim 4 (See Figure 8 for an illustration of the notations) by setting $A=C_{i}^{\prime}$ and $H=G_{i}$. In particular, we have $d_{\mathrm{rs}}(A)=d(P, P) \leq$ $d(W, W)+2 d(Y, W+Y)$ and $d(H, H) \geq d(W, W)+2 d(X, W)$. By Claim 4, $d(Y, W+$ $Y) \leq d(X, W)$, which completes the proof.


Figure 9: Notations in Claim 8.

Claim 8. For $t \in\{i, j\},\left|C_{t}\right| \geq 100 m_{B}$, and $C_{t}^{\prime}$ contains all good points in $C_{t}$ but no good points in other optimal clusters. Then $d_{\mathrm{rs}}\left(C_{i}^{\prime} \cup C_{j}^{\prime}\right)-d\left(G_{i}, G_{i}\right)-d\left(G_{j}, G_{j}\right) \geq$ $\left(\frac{4}{3}-\frac{4}{\beta}\right) d\left(G_{i}, G_{j}\right)$.

Proof. Let $C=C_{i}^{\prime}+C_{j}^{\prime}, F=F(C), P=P(C)$, and $G=G_{i}+G_{j}, \bar{G}=C \backslash G$. Define $W_{i}=G_{i} \cap P, X_{i}=G_{i} \cap F$; define $W_{j}, X_{j}$ similarly. Also, define $Y=P \cap \bar{G}, V=F \cap \bar{G}$. See Figure 9. Then
$d(C \backslash F, C \backslash F)-d\left(G_{i}, G_{i}\right)-d\left(G_{j}, G_{j}\right) \geq 2 d\left(W_{i}, W_{j}\right)-2 d\left(X_{i}, G_{i}\right)-2 d\left(X_{j}, G_{j}\right)$.
We have $d\left(X_{i}, C_{i}\right)+d\left(X_{j}, C_{j}\right) \leq \frac{d\left(X_{i}, C_{j}\right)+d\left(X_{j}, C_{i}\right)}{\beta} \leq \frac{2}{\beta} d\left(C_{i}, C_{j}\right)$ by the definition of good points, and $d\left(C_{i}, C_{j}\right) \leq \frac{3}{2} d\left(W_{i}, W_{j}\right)$ by Lemma 10 . So (10) is at least $\left(\frac{4}{3}-\right.$ $\left.\frac{4}{\beta}\right) d\left(C_{i}, C_{j}\right) \geq\left(\frac{4}{3}-\frac{4}{\beta}\right) d\left(G_{i}, G_{j}\right)$.

By these claims and Lemma 11, when splitting some node in the pruning while merging some other two, the cost increased is larger than the cost saved. Since any pruning of size $k$ can be obtained from $\mathcal{C}^{\prime}$ by splitting some nodes while merging some others, we can show that $\mathcal{C}^{\prime}$ has the minimum robust min-sum cost. See Appendix A.5.2 for the complete proof.

### 2.5.2.3 Getting a Good Approximation

We have showed that the pruning $\mathcal{C}^{\prime}$ that assigns all good points correctly has minimum robust min-sum cost, so we can use dynamic programming on the tree to get the pruning. However, this pruning may not be a good approximation. For example, consider an instance consisting of two unbalanced clusters. Assume that there is only one bad point, belonging to the small cluster. Further assume the distances between the good points in each cluster are negligible, then assigning the bad point incorrectly to the large cluster will lead to an $O\left(\frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}\right)$-approximation. So the pruning $\mathcal{C}^{\prime}$ may not be a constant approximation.

Here we show that a constant factor approximation can be computed by reassigning the points in $\mathcal{C}^{\prime}$. As described in Algorithm 7, we reassign each point $p$ to the

Algorithm 7 Getting a good approximation
Input: A clustering $\mathcal{C}^{\prime}=\left\{C_{1}^{\prime}, \ldots, C_{k}^{\prime}\right\}$, where $G_{i} \subseteq C_{i}^{\prime} \subseteq C_{i} \cup B$.
1: for each point $p$ do
Assign $p$ to the index $i$ such that $d\left(p, P\left(C_{i}^{\prime}\right)\right)$ is minimized.
end for
4: Let $C_{i}^{\prime \prime}$ be the set of points assigned to the index $i$.
Output: The clustering $\mathcal{C}^{\prime \prime}=\left\{C_{1}^{\prime \prime}, \ldots, C_{k}^{\prime \prime}\right\}$.
index $i$ that minimizes the cost between $p$ and the potential good points in the cluster $C_{i}^{\prime}$. To analyze Algorithm 7, we first prove that after reassignment all good points are still assigned correctly (Lemma 15), and then bound the cost.

Lemma 15. For any $p \in G_{i}$, any $j \neq i, d\left(p, P\left(C_{j}^{\prime}\right)\right)>d\left(p, P\left(C_{i}^{\prime}\right)\right)$.

Proof Sketch. By Lemma 12, $d\left(p, P\left(C_{i}^{\prime}\right)\right) \approx d\left(p, C_{i}\right)$. By definition of good points, $\beta d\left(p, C_{i}\right) \leq d\left(p, C_{j}\right)$, so it suffices to show that $d\left(p, P\left(C_{j}^{\prime}\right)\right)$ is not so small compared to $d\left(p, C_{j}\right)$. Let $W_{j}=G_{j} \cap P\left(C_{j}^{\prime}\right)$ denote the good points that are also potential good points, and let $Z_{j}=C_{j} \backslash W_{j}$ denote all other points in $C_{j}$. Since $W_{j} \subseteq P\left(C_{j}^{\prime}\right)$, we only need to prove that $d\left(p, W_{j}\right)$ is large compared to $d\left(p, Z_{j}\right)$. Intuitively, this is true since $p$ is far from good points in $G_{j}$, and $W_{j}$ contains most of $G_{j}$. See Appendix A.5.3 for details.

Theorem 8. Suppose the instance is $(\alpha, \epsilon)$-perturbation resilient to the min-sum objective for $\alpha>\frac{8 \max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$, $\epsilon<\frac{\min _{i}\left|C_{i}\right|}{600 n \log n}$. There exists an algorithm that outputs $a$ clustering which is a $\left(1+\frac{40 \epsilon n \log n}{\min _{i}\left|C_{i}\right|}\right)$-approximation to the optimal clustering in polynomial time. Furthermore, the output clustering is also $(6 \epsilon \log n)$-close to the optimal clustering.

Proof. By Lemma 15, all the good points in $C_{i}$ are assigned correctly to $C_{i}^{\prime \prime}$. Let $A_{i}=C_{i}^{\prime \prime} \backslash G_{i}$ denote all the bad points assigned to $C_{i}^{\prime \prime}$. The cost of the output
clustering $\mathcal{C}^{\prime \prime}$ can be written as follows.

$$
\begin{align*}
\sum_{i} d\left(C_{i}^{\prime \prime}, C_{i}^{\prime \prime}\right) & =\sum_{i} d\left(G_{i}+A_{i}, G_{i}+A_{i}\right) \\
& =\sum_{i} d\left(G_{i}, G_{i}\right)+2 \sum_{i} d\left(G_{i}, A_{i}\right)+\sum_{i} d\left(A_{i}, A_{i}\right) \tag{11}
\end{align*}
$$

We need to bound the last two terms.
Let $r=\frac{\min _{i}\left|C_{i}\right|}{m_{B}}$. By the triangle inequality, we have $d_{\mathrm{a}}\left(A_{i}, A_{i}\right) \leq 2 d_{\mathrm{a}}\left(A_{i}, G_{i}\right)$, leading to

$$
\begin{equation*}
d\left(A_{i}, A_{i}\right) \leq \frac{2\left|A_{i}\right|}{\left|G_{i}\right|} d\left(A_{i}, G_{i}\right) \leq \frac{2 m_{B}}{\left|C_{i}\right|-m_{B}} d\left(A_{i}, G_{i}\right) \leq \frac{2}{r-5} d\left(A_{i}, G_{i}\right) \tag{12}
\end{equation*}
$$

So it suffices to bound $d\left(A_{i}, G_{i}\right)$. In Appendix A.5.3, we prove
Claim 9. $\sum_{i} d\left(A_{i}, G_{i}\right) \leq \frac{r^{2}}{(r-5)^{2}} \sum_{i} d\left(C_{i}, C_{i}\right)-\frac{r^{2}-1}{(r-5)^{2}} \sum_{i} d\left(G_{i}, G_{i}\right)$.
The intuition is as follows. Suppose $p \in A_{i}$ comes from $C_{j}$. Then $d\left(p, G_{i}\right) \approx$ $d\left(p, P\left(C_{i}^{\prime}\right)\right)$, which is smaller than $d\left(p, P\left(C_{j}^{\prime}\right)\right)$ since $p$ is assigned to $i$. By Lemma 12 , we have $d\left(p, P\left(C_{j}^{\prime}\right)\right) \approx d\left(p, G_{j}\right)$. Applying this argument for all $p \in A_{i}$ and all $i$, and noting that $\cup_{i} A_{i}=\cup_{j} B_{j}$, we have $\sum_{i} d\left(A_{i}, G_{i}\right)$ is approximately less than $\sum_{j} d\left(B_{j}, G_{j}\right) \leq \sum_{j}\left[d\left(C_{j}, C_{j}\right)-d\left(G_{j}, G_{j}\right)\right]$.

The proof of correctness is completed by combining this claim and the inequalities (11) and (12).

Running Time Algorithm 3 takes time $O\left(n^{\omega+1}\right)$ (as shown in the proof of Theorem 4), and the rest steps of Algorithm 6 take time $O\left(n^{3}\right)$. Finding the minimum robust min-sum cost pruning in the tree output by Algorithm 6 takes time $O\left(n^{3}\right)$, and Algorithm 7 takes time $O\left(n^{3}\right)$. So the total running time is $O\left(n^{\omega+1}\right)$.

## CHAPTER III

## DISTRIBUTED CLUSTERING

Many modern applications face an explosion of data. Usually, the data is distributed over different locations, such as distributed databases [91, 36], images and videos over networks [87], surveillance [53] and sensor networks [35, 54]. In many of these applications the data are inherently distributed because, as in sensor networks, it is collected at different sites. Since most classic clustering algorithms are designed for the centralized setting, it has become crucial to develop clustering algorithms which are effective in the distributed setting.

Several algorithms for distributed clustering have been proposed and empirically tested. Some of these algorithms [48, 101, 38] are direct adaptations of centralized algorithms which rely on statistics that are easy to compute in a distributed manner. Other algorithms $[60,64]$ generate summaries of local data and transmit them to a central coordinator which then performs the clustering algorithm. They do not provide theoretical guarantees on the clustering quality, or the reduction in communication cost. Additionally, most of these algorithms assume that the distributed nodes can communicate with all other sites or that there is a central coordinator that communicates with all other sites.

In this chapter, we study the problem of distributed $k$-median and $k$-means clustering where the data is distributed across nodes whose communication is restricted to the edges of an arbitrary graph. We provide algorithms with small communication cost and provable guarantees on the clustering quality. Our technique for reducing communication in general graphs is based on the construction of a small set of points called coreset, which act as a proxy for the entire data set. For $k$-means, we further
provide an algorithm for high dimensional data, where we first reduce the dimension of the data by distributed principal component analysis (PCA) and then apply our distributed clustering algorithm to the projected data.

An $\epsilon$-coreset is a weighted set of points whose cost on any set of centers is approximately the cost of the original data on those same centers up to accuracy $\epsilon$. Thus an approximate solution for the coreset is also an approximate solution for the original data. Coresets have previously been studied in the centralized setting ( $[57,45]$ ) but have also recently been used for distributed clustering as in [104] and as implied by [47]. In Section 3.2, we propose a distributed algorithm for $k$-means and $k$-median, by which each node constructs a local portion of a global coreset. The nodes then share the local portions of the coreset, which can be done efficiently in general graphs using a message passing approach. More precisely, each node computes an approximate solution for its local data and communicate the cost of this local solution, and then constructs the local portion of a global coreset using only its local data and the total cost of each node's local solution. For $\epsilon$ constant, this builds a coreset of size $\tilde{O}(k d+s k)$ for $k$-median and $k$-means when the data points have dimension $d$ and are distributed over $s$ sites $^{1}$. If there is a central coordinator among the $s$ sites, then clustering can be performed on the coordinator by collecting the local portions of the coreset with a communication cost equal to the coreset size $\tilde{O}(k d+s k)$. For distributed clustering over general connected topologies, we propose an algorithm based on the distributed coreset construction and a message-passing approach, whose communication cost improves over previous coreset-based algorithms. Experimental results on large scale data sets show that our algorithm performs well in practice. For a fixed amount of communication, our algorithm outperforms other

[^3]coreset construction algorithms.
In the above algorithms, the number of points in the coreset is independent of the number of the original data points, which is useful for large scale applications. However, it is linear in the dimension of the data, leading to high communication cost for high dimensional data. In Section 3.3, we propose a distributed PCA algorithm, and show that its output represents the original data in the sense that any good approximation solution of $k$-means clustering on the projected data is also a good solution on the original data. When combined with the distributed clustering algorithm in Section 3.2, this leads to an algorithm whose communication cost (in terms of the number of points communicated) is independent of the size and the dimension of the original data. Our experiment results demonstrate that this significantly reduces the communication cost while hardly comprising the quality of the $k$-means clustering solutions.

### 3.1 Preliminaries

Let $d(p, q)$ denote the Euclidean distance between any two points $p, q \in \mathbf{R}^{d}$. The goal of $k$-means clustering is to find a set of $k$ centers $\mathbf{x}=\left\{x_{1}, x_{2}, \ldots, x_{k}\right\}$ which minimize the $k$-means cost of data set $P \subseteq \mathbf{R}^{d}$. Here the $k$-means cost is defined as $\operatorname{cost}(P, \mathbf{x})=\sum_{p \in P} d(p, \mathbf{x})^{2}$ where $d(p, \mathbf{x})=\min _{x \in \mathbf{x}} d(p, x)$. If $P$ is a weighted data set with a weighting function $w$, then the $k$-means cost is defined as $\sum_{p \in P} w(p) d(p, \mathbf{x})^{2}$. Similarly, the $k$-median cost is defined as $\sum_{p \in P} d(p, \mathbf{x})$. Both $k$-means and $k$-median cost functions are known to be NP-hard to minimize (see for example [7]). For both objectives, there exist several readily available polynomial time algorithms that achieve constant approximation solutions (see for example $[63,76]$ ).

In the distributed clustering task, we consider a set of $s$ nodes $V=\left\{v_{i}, 1 \leq i \leq s\right\}$ which communicate on an undirected connected graph $G=(V, E)$ with $m=|E|$ edges. More precisely, an edge $\left(v_{i}, v_{j}\right) \in E$ indicates that $v_{i}$ and $v_{j}$ can communicate
with each other. On each node $v_{i}$, there is a local set of data points $P_{i}$, and the global data set is $P=\cup_{i=1}^{s} P_{i}$. The goal is to find a set of $k$ centers $\mathbf{x}$ which optimize $\operatorname{cost}(P, \mathbf{x})$ while keeping the computation efficient and the communication cost as low as possible. Our focus is to reduce the total communication cost while preserving theoretical guarantees for approximating the clustering cost.

Here we measure the communication cost in number of points transmitted. In some cases we measure the communication in number of words transmitted, which will be explicitly pointed out. We also assume for simplicity that there is no latency in the communication.

### 3.1.1 Coresets

For the distributed clustering task, a natural approach to avoid broadcasting raw data is to generate a local summary of the relevant information. If each site computes a summary for their own data set and then communicates this to a central coordinator, a solution can be computed from a much smaller amount of data, drastically reducing the communication.

In the centralized setting, the idea of summarization with respect to the clustering task is captured by the concept of coresets [57, 45]. A coreset is a set of points, together with a weight for each point, such that the cost of this weighted set approximates the cost of the original data for any set of $k$ centers. The formal definition of coresets is:

Definition 10. An $\epsilon$-coreset for a set of points $P$ with respect to a center-based cost function is a set of points $S$ and a set of weights $w: S \rightarrow \mathbf{R}$ such that for any set of centers $\mathbf{x}$,

$$
(1-\epsilon) \operatorname{cost}(P, \mathbf{x}) \leq \sum_{p \in S} w(p) \operatorname{cost}(p, \mathbf{x}) \leq(1+\epsilon) \operatorname{cost}(P, \mathbf{x})
$$

In the centralized setting, many coreset construction algorithms have been proposed for $k$-median, $k$-means and some other cost functions. For example, for points
in $\mathbf{R}^{d}$, algorithms in [45] construct coresets of size $t=\tilde{O}\left(k d / \epsilon^{4}\right)$ for $k$-means and coresets of size $t=\tilde{O}\left(k d / \epsilon^{2}\right)$ for $k$-median. In the distributed setting, it is natural to ask whether there exists an algorithm that constructs a small coreset for the entire point set but still has low communication cost. Note that the union of coresets for multiple data sets is a coreset for the union of the data sets. The immediate construction of combining the local coresets from each node would produce a global coreset whose size was larger by a factor of $s$, greatly increasing the communication complexity. We present a distributed algorithm which constructs a global coreset the same size as the centralized construction and only needs a single value ${ }^{2}$ communicated to each node. This serves as the basis for our distributed clustering algorithm.

### 3.1.2 Principal Component Analysis

PCA is a classical tool for dimension reduction, and has been closely related to $k$ means [41, 70]. In Section 3.3, we first use PCA on high dimensional data and then do distributed clustering on the projected data, which leads to lower communication cost. We introduce the following notations for PCA. View the local data $P_{i}$ as a matrix, whose rows are data points. The global data $P$ is then a concatenation of the local data matrix, i.e. $P^{\top}=\left[P_{1}^{\top}, P_{2}^{\top}, \ldots, P_{s}^{\top}\right]$. For simplicity, we always assume the data is centered, that is, $\sum_{p \in P} p=0$; otherwise, we can first perform a step to center the data, whose communication and computation costs will be dominated by those in the other steps of our algorithms.

For a matrix $X=\left[x_{i j}\right]$, let $\|X\|_{F}^{2}=\sum_{i, j} x_{i, j}^{2}$. We say that $X$ has orthonormal columns if its columns are orthogonal unit vectors. Let $L(X)$ denote the linear subspace spanned by the columns of $X$. For simplicity, for a set of points $P$, we denote $d^{2}(P, L(X)):=\sum_{p \in P} d(p, L(X))^{2}=\sum_{p \in P}\left[\min _{q \in L(X)} d(p, q)\right]^{2}$. For a point $p$, let $\Pi_{X}(p)$ denote its projection to $L(X)$. Note that for an orthogonal matrix $X$, the

[^4]projection of a point $p$ to $L(X)$ will be $p X$ using the coordinates with respect to the column space of $X$, and will be $p X X^{T}$ using the original coordinates.

### 3.2 Distributed Coreset-Based Clustering

Since coresets summarize local information they are a natural tool to use when trying to reduce communication complexity. If each node constructs an $\epsilon$-coreset on its local data, then the union of these coresets is clearly an $\epsilon$-coreset for the entire data set. Unfortunately the size of the coreset in this approach increases greatly with the number of nodes.

Another approach is the one presented in [104]. Its main idea is to approximate the union of local coresets with another coreset. They assume nodes communicate over a rooted tree, with each node passing its coreset to its parent. Because the approximation factor of the constructed coreset depends on the quality of its component coresets, the accuracy a coreset needs (and thus the overall communication complexity) scales with the height of this tree. Although it is possible to find a spanning tree in any communication network, when the graph has large diameter every tree has large height. In particular many natural networks such as grid networks have a large diameter $(\Omega(\sqrt{s})$ for grids) which greatly increases the size of coresets which must be communicated across the lower levels of the tree. We show that it is possible to construct a global coreset with low communication overhead. This is done by distributing the coreset construction procedure rather than combining local coresets. The communication needed to construct this coreset is negligible - just a single value from each data set representing the approximate cost of their local optimal clustering. Since the sampled global $\epsilon$-coreset is the same size as any local $\epsilon$-coreset, this leads to an improvement of the communication cost over the other approaches. See Figure 10 for an illustration. The constructed coreset is smaller by a factor of $s$ in general graphs, and is independent of the communication topology. This method excels in


Figure 10: Illustration of different coreset construction approaches.
sparse networks with large diameters, where the previous approach in [104] requires coresets that are quadratic in the size of the diameter for $k$-median and quartic for $k$-means; see Section 3.2.2 for details. [47] also merge coresets using coreset construction, but they do so in a model of parallel computation and ignore communication costs.

Additional related work. Many empirical algorithms adapt the centralized algorithms to the distributed setting. They generally provide no bound for the clustering quality or the communication cost. For instance, a technique is proposed in [48] to adapt several iterative center-based data clustering algorithms including Lloyd's algorithm for $k$-means to the distributed setting, where sufficient statistics instead of the raw data are sent to a central coordinator. This approach involves transferring data back and forth in each iteration, and thus the communication cost depends on the number of iterations. Similarly, the communication costs of the distributed clustering algorithms proposed in [38] and [101] depend on the number of iterations. Some other algorithms gather local summaries and then perform global clustering on the summaries. The distributed density-based clustering algorithm in [60] clusters and
computes summaries for the local data at each node, and sends the local summaries to a central node where the global clustering is carried out. This algorithm only considers the flat two-tier topology. Some in-network aggregation schemes for computing statistics over distributed data are useful for such distributed clustering algorithms. For example, an algorithm is provided in [35] for approximate duplicate-sensitive aggregates across distributed data sets, such as SUM. An algorithm is proposed in [54] for power-preserving computation of order statistics such as quantile.

Several coreset construction algorithms have been proposed for $k$-median, $k$-means and $k$-line median clustering [57, 30, 56, 74, 45]. For example, the algorithm in [45] constructs a coreset of size $\tilde{O}\left(k d / \epsilon^{2}\right)$ whose cost approximates that of the original data up to accuracy $\epsilon$ with respect to $k$-median in $\mathbf{R}^{d}$. All of these algorithms consider coreset construction in the centralized setting, while our construction algorithm is for the distributed setting.

There has also been work attempting to parallelize clustering algorithms. [47] showed that coresets could be constructed in parallel and then merged together. Bahmani et al. [12] adapted k -means++ to the parallel setting. Their algorithm, kmeans $\|$, essentially builds $O(1)$-coreset of size $O(k \log |P|)$. However, it cannot build $\epsilon$-coreset for $\epsilon=o(1)$, and thus can only guarantee constant approximation solutions.

There is also related work providing approximation solutions for $k$-median based on random sampling [25]. Particularly, they showed that given a sample of size $\tilde{O}\left(\frac{k}{\epsilon^{2}}\right)$ drawn i.i.d. from the data, there exists an algorithm that outputs a solution with an average cost bounded by twice the optimal average cost plus an error bound $\epsilon$. If we convert it to a multiplicative approximation factor, the factor depends on the optimal average cost. When there are outlier points far away from all other points, the optimal average cost can be very small after normalization, then the multiplicative approximation factor is large. The coreset approach provides better guarantees. Additionally, their approach is not applicable to $k$-means.

Balcan et al. [18] and Daume et al. [39] consider fundamental communication complexity questions arising when doing classification in distributed settings. In concurrent and independent work, Vempala et al. [62] study several optimization problems in distributed settings, including $k$-means clustering under an interesting separability assumption.

### 3.2.1 Distributed Coreset Construction

Here we design a distributed coreset construction algorithm for $k$-means and $k$ median. Note that the underlying technique can be extended to other additive clustering objectives such as $k$-line median and sum of distances between points and their centers to the power of $z$.

To gain some intuition on the distributed coreset construction algorithm, we briefly review the coreset construction algorithm in [45] in the centralized setting. The coreset is constructed by computing a constant approximation solution for the entire data set, and then sampling points proportional to their contributions to the cost of this solution. Intuitively, the points close to the nearest centers can be approximately represented by the nearest centers while points far away cannot be well represented. Thus, points should be sampled with probability proportional to their contributions to the cost.

Directly adapting the algorithm to the distributed setting would require computing a constant approximation solution for the entire data set. We show that a global coreset can be constructed in a distributed fashion by estimating the weight of the entire data set with the sum of local approximations. We first compute a local approximation solution for each local data set, and communicate the total costs of these local solutions. Then we sample points proportional to their contributions to the cost of their local solutions. At the end of the algorithm, the coreset consists of the sampled points and the centers in the local solutions. The coreset points are distributed
over the nodes, so we call it distributed coreset. See Algorithm 8 for details.

```
Algorithm 8 Communication aware distributed coreset construction
Input: Local data sets \(\left\{P_{i}\right\}_{i=1}^{s}\), parameter \(t\) (number of points to be sampled).
    for each node \(v_{i} \in V\) do
        Compute a constant approximation \(B_{i}\) for \(P_{i}\).
        Communicate \(\operatorname{cost}\left(P_{i}, B_{i}\right)\) to all other nodes.
    end for
    for each node \(v_{i} \in V\) do
        Set \(t_{i}=\frac{t \operatorname{cost}\left(P_{i}, B_{i}\right)}{\sum_{j=1}^{s} \operatorname{cost}\left(P_{j}, B_{j}\right)}\). Set \(m_{p}=\operatorname{cost}\left(p, B_{i}\right)\) for each \(p \in P_{i}\).
        Pick a non-uniform random sample \(S_{i}\) of \(t_{i}\) points from \(P_{i}\), where for every
        \(q \in S_{i}\) and \(p \in P_{i}\), we have \(q=p\) with probability \(m_{p} / \sum_{z \in P_{i}} m_{z}\).
        Let \(w_{q}=\sum_{i} \sum_{z \in P_{i}} m_{z} /\left(t m_{q}\right)\) for each \(q \in S_{i}\).
        for each \(b \in B_{i}\) do
            Let \(P_{b}=\left\{p \in P_{i}: d(p, b)=d\left(p, B_{i}\right)\right\}, w_{b}=\left|P_{b}\right|-\sum_{q \in P_{b} \cap S} w_{q}\).
        end for
    end for
Output: Distributed coreset: points \(S_{i} \cup B_{i}\) with weights \(\left\{w_{q}: q \in S_{i} \cup B_{i}\right\}_{i=1}^{s}\).
```

Theorem 10. For distributed $k$-means and $k$-median clustering on a graph, there exists an algorithm such that with probability at least $1-\delta$, the union of its output on all nodes is an $\epsilon$-coreset for $P=\cup_{i=1}^{s} P_{i}$. The size of the coreset is $O\left(\frac{1}{\epsilon^{4}}(k d+\right.$ $\left.\left.\log \frac{1}{\delta}\right)+n k \log \frac{n k}{\delta}\right)$ for $k$-means, and $O\left(\frac{1}{\epsilon^{2}}\left(k d+\log \frac{1}{\delta}\right)+n k\right)$ for $k$-median. The total communication cost is $O(m n)$.

As described below, the distributed coreset construction can be achieved by using Algorithm 8 with appropriate $t$, namely $O\left(\frac{1}{\epsilon^{4}}\left(k d+\log \frac{1}{\delta}\right)+n k \log \frac{n k}{\delta}\right)$ for $k$-means and $O\left(\frac{1}{\epsilon^{2}}\left(k d+\log \frac{1}{\delta}\right)\right)$ for $k$-median. The formal proofs are described in the following subsections.

### 3.2.1.1 Proof of Theorem 14: $k$-median

The analysis relies on the definition of the pseudo-dimension of a function space and a sampling lemma.

Definition 11 ([77, 45]). Let $F$ be a finite set of functions from a set $P$ to $\mathbf{R}_{\geq 0}$. For
$f \in F$, let $B(f, r)=\{p: f(p) \leq r\}$. The dimension of the function space $\operatorname{dim}(F, P)$ is the smallest integer $d$ such that for any $G \subseteq P,|\{G \cap B(f, r): f \in F, r \geq 0\}| \leq|G|^{d}$.

Suppose we draw a sample $S$ according to $\left\{m_{p}: p \in P\right\}$, namely for every $q \in S$ and every $p \in P$, we have $q=p$ with probability $\frac{m_{p}}{\sum_{z \in P} m_{z}}$. Set the weights of the points as $w_{p}=\frac{\sum_{z \in P} m_{z}}{m_{p}|S|}$ for $p \in P$. Then for any $f \in F$, the expectation of the weighted cost of $S$ equals the cost of the original data $P$ :

$$
\begin{aligned}
\mathbf{E}\left[\sum_{q \in S} w_{q} f(q)\right] & =\sum_{q \in S} \mathbf{E}\left[w_{q} f(q)\right]=\sum_{q \in S} \sum_{p \in P} \operatorname{Pr}[q=p] w_{p} f(p) \\
& =\sum_{q \in S} \sum_{p \in P} \frac{m_{p}}{\sum_{z \in P} m_{z}} \frac{\sum_{z \in P} m_{z}}{m_{p}|S|} f(p)=\sum_{q \in S} \sum_{p \in P} \frac{1}{|S|} f(p)=\sum_{p \in P} f(p) .
\end{aligned}
$$

The following lemma shows that if the sample size is large enough, then we also have concentration for any $f \in F$. The lemma is implicit in [45] and we include the proof in the appendix for completeness.

Lemma 16. Fix a set $F$ of functions $f: P \rightarrow \mathbf{R}_{\geq 0}$. Let $S$ be a sample drawn i.i.d. from $P$ according to $\left\{m_{p}: p \in P\right\}$, namely, for every $q \in S$ and every $p \in P$, we have $q=p$ with probability $\frac{m_{p}}{\sum_{z \in P} m_{z}}$. Let $w_{p}=\frac{\sum_{z \in P} m_{z}}{m_{p}|S|}$ for $p \in P$. For a sufficiently large $c$, if $|S| \geq \frac{c}{\epsilon^{2}}\left(\operatorname{dim}(F, P)+\log \frac{1}{\delta}\right)$ then with probability at least $1-\delta, \forall f \in F$ : $\left|\sum_{p \in P} f(p)-\sum_{q \in S} w_{q} f(q)\right| \leq \epsilon\left(\sum_{p \in P} m_{p}\right)\left(\max _{p \in P} \frac{f(p)}{m_{p}}\right)$.

To get a small bound on the difference between $\sum_{p \in P} f(p)$ and $\sum_{q \in S} w_{q} f(q)$, we need to choose $m_{p}$ such that $\max _{p \in P} \frac{f(p)}{m_{p}}$ is bounded. More precisely, if we choose $m_{p}=\max _{f \in F} f(p)$, then the difference is bounded by $\epsilon \sum_{p \in P} m_{p}$.

We first consider the centralized setting and review how [45] applied the lemma to construct a coreset for $k$-median as in Definition 10. A natural approach is to apply this lemma directly to the cost, namely, to choose $f_{\mathbf{x}}(p):=\operatorname{cost}(p, \mathbf{x})$. The problem is that a suitable upper bound $m_{p}$ is not available for $\operatorname{cost}(p, \mathbf{x})$. However, we can still apply the lemma to a different set of functions defined as follows. Let $b_{p}$ denote the closest center to $p$ in the approximation solution. Aiming to approximate the
error $\sum_{p}\left[\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)\right]$ rather than to approximate $\sum_{p} \operatorname{cost}(p, \mathbf{x})$ directly, we define $f_{\mathbf{x}}(p):=\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)+\operatorname{cost}\left(p, b_{p}\right)$, where $\operatorname{cost}\left(p, b_{p}\right)$ is added so that $f_{\mathbf{x}}(p) \geq 0$. Since $0 \leq f_{\mathbf{x}}(p) \leq 2 \operatorname{cost}\left(p, b_{p}\right)$, we can apply the lemma to $f_{\mathbf{x}}(p)$ and $m_{p}=$ $2 \operatorname{cost}\left(p, b_{p}\right)$. The lemma then bounds the difference $\left|\sum_{p \in P} f_{\mathbf{x}}(p)-\sum_{q \in S} w_{q} f_{\mathbf{x}}(q)\right|$ by $2 \epsilon \sum_{p \in P} \operatorname{cost}\left(p, b_{p}\right)$, so we have an $O(\epsilon)$-approximation.

Note that $\sum_{p \in P} f_{\mathbf{x}}(p)-\sum_{q \in S} w_{q} f_{\mathbf{x}}(q)$ does not equal $\sum_{p \in P} \operatorname{cost}(p, \mathbf{x})-\sum_{q \in S} w_{q} \operatorname{cost}(q, \mathbf{x})$. However, it equals the difference between $\sum_{p \in P} \operatorname{cost}(p, \mathbf{x})$ and a weighted cost of the sampled points and the centers in the approximation solution. To get a coreset as in Definition 10, we need to add the centers of the approximation solution with specific weights to the coreset. Then when the sample is sufficiently large, the union of the sampled points and the centers is an $\epsilon$-coreset.

Our key contribution in this paper is to show that in the distributed setting, it suffices to choose $b_{p}$ from the local approximation solution for the local data set containing $p$, rather than from an approximation solution for the global data set. Furthermore, the sampling and the weighting of the coreset points can be done in a local manner. In the following, we provide a formal verification of our discussion above. We have the following lemma for $k$-median with $F=\left\{f_{\mathbf{x}}: f_{\mathbf{x}}(p)=d(p, \mathbf{x})-\right.$ $\left.d\left(b_{p}, \mathbf{x}\right)+d\left(p, b_{p}\right), \mathbf{x} \in\left(\mathbf{R}^{d}\right)^{k}\right\}$.

Lemma 17. For $k$-median, the output of Algorithm 8 is an $\epsilon$-coreset with probability at least $1-\delta$, if $t \geq \frac{c}{\epsilon^{2}}\left(\operatorname{dim}(F, P)+\log \frac{1}{\delta}\right)$ for a sufficiently large constant $c$.

Proof. We want to show that for any set of centers $\mathbf{x}$ the true cost for using these centers is well approximated by the cost on the weighted coreset. Note that our coreset has two types of points: sampled points $p \in S=\cup_{i=1}^{s} S_{i}$ with weight $w_{p}:=\frac{\sum_{z \in P} m_{z}}{m_{p}|S|}$ and local solution centers $b \in B=\cup_{i=1}^{s} B_{i}$ with weight $w_{b}:=\left|P_{b}\right|-\sum_{p \in S \cap P_{b}} w_{p}$. We use $b_{p}$ to represent the nearest center to $p$ in the local approximation solution. We use $P_{b}$ to represent the set of points having $b$ as their closest center in the local approximation solution.

As mentioned above, we construct $f_{\mathbf{x}}$ to be the difference between the cost of $p$ and the cost of $b_{p}$ on $\mathbf{x}$ so that Lemma 16 can be applied to $f_{\mathbf{x}}$. Note that $0 \leq f_{\mathbf{x}}(p) \leq$ $2 d\left(p, b_{p}\right)$ by triangle inequality, and $S$ is sufficiently large and chosen according to weights $m_{p}=d\left(p, b_{p}\right)$, so the conditions of Lemma 16 are met. Then we have

$$
\begin{aligned}
D & =\left|\sum_{p \in P} f_{\mathbf{x}}(p)-\sum_{q \in S} w_{q} f_{\mathbf{x}}(q)\right| \leq 2 \epsilon \sum_{p \in P} m_{p}=2 \epsilon \sum_{p \in P} d\left(p, b_{p}\right)=2 \epsilon \sum_{i=1}^{s} d\left(P_{i}, B_{i}\right) \\
& \leq O(\epsilon) \sum_{p \in P} d(p, \mathbf{x})
\end{aligned}
$$

where the last inequality follows from the fact that $B_{i}$ is a constant approximation solution for $P_{i}$.

Next, we show that the coreset is constructed such that $D$ is exactly the difference between the true cost and the weighted cost of the coreset, which then leads to the lemma.

Note that the centers are weighted such that

$$
\begin{equation*}
\sum_{b \in B} w_{b} d(b, \mathbf{x})=\sum_{b \in B}\left|P_{b}\right| d(b, \mathbf{x})-\sum_{b \in B} \sum_{q \in S \cap P_{b}} w_{q} d(b, \mathbf{x})=\sum_{p \in P} d\left(b_{p}, \mathbf{x}\right)-\sum_{q \in S} w_{q} d\left(b_{q}, \mathbf{x}\right) . \tag{13}
\end{equation*}
$$

Also note that $\sum_{p \in P} m_{p}=\sum_{q \in S} w_{q} m_{q}$, so

$$
\begin{align*}
D & =\left|\sum_{p \in P}\left[d(p, \mathbf{x})-d\left(b_{p}, \mathbf{x}\right)+m_{p}\right]-\sum_{q \in S} w_{q}\left[d(q, \mathbf{x})-d\left(b_{q}, \mathbf{x}\right)+m_{q}\right]\right| \\
& =\left|\sum_{p \in P} d(p, \mathbf{x})-\sum_{q \in S} w_{q} d(q, \mathbf{x})-\left[\sum_{p \in P} d\left(b_{p}, \mathbf{x}\right)-\sum_{q \in S} w_{q} d\left(b_{q}, \mathbf{x}\right)\right]\right| \tag{14}
\end{align*}
$$

By plugging (13) into (14), we have

$$
D=\left|\sum_{p \in P} d(p, \mathbf{x})-\sum_{q \in S} w_{q} d(q, \mathbf{x})-\sum_{b \in B} w_{b} d(b, \mathbf{x})\right|=\left|\sum_{p \in P} d(p, \mathbf{x})-\sum_{q \in S \cup B} w_{q} d(q, \mathbf{x})\right|
$$

which implies the lemma.

In [45] it is shown that ${ }^{3} \operatorname{dim}(F, P)=O(k d)$. Therefore, by Lemma 17, when $|S| \geq O\left(\frac{1}{\epsilon^{2}}\left(k d+\log \frac{1}{\delta}\right)\right)$, the weighted cost of $S \cup B$ approximates the $k$-median cost of $P$ for any set of centers, then $(S \cup B, w)$ is an $\epsilon$-coreset for $P$. The total communication cost is bounded by $O(m n)$, since even in the most general case when every node only knows its neighbors, we can broadcast the local costs with $O(m n)$ communication (see Algorithm 10).

### 3.2.1.2 Proof of Theorem 14: $k$-means

We have for $k$-means a similar lemma that when $\left.t=O\left(\frac{1}{\epsilon^{4}}\left(k d+\log \frac{1}{\delta}\right)+n k \log \frac{n k}{\delta}\right)\right)$, the algorithm constructs an $\epsilon$-coreset with probability at least $1-\delta$. The key idea is the same as that for $k$-median: we use centers $b_{p}$ from the local approximation solutions as an approximation to the original data points $p$, and show that the error between the total cost and the weighted sample cost is approximately the error between the cost of $p$ and its sampled cost (compensated by the weighted centers), which is shown to be small by Lemma 16 .

The key difference between $k$-means and $k$-median is that triangle inequality applies directly to the $k$-median cost. In particular, for the $k$-median problem note that $\operatorname{cost}\left(b_{p}, p\right)=d\left(b_{p}, p\right)$ is an upper bound for the error of $b_{p}$ on any set of centers, i.e. $\forall \mathbf{x} \in\left(\mathbf{R}^{d}\right)^{k}, d\left(b_{p}, p\right) \geq\left|d(p, \mathbf{x})-d\left(b_{p}, \mathbf{x}\right)\right|=\left|\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)\right|$ by triangle inequality. Then we can construct $f_{\mathbf{x}}(p):=\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)+d\left(b_{p}, p\right)$ such that $h_{p}(\mathbf{x})$ is bounded. In contrast, for $k$-means, the error $\left|\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)\right|=$ $\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right|$ does not have such an upper bound. The main change to the analysis is that we divide the points into two categories: good points whose costs approximately satisfy the triangle inequality (up to a factor of $1 / \epsilon$ ) and bad points.

[^5]The good points for a fixed set of centers $\mathbf{x}$ are defined as

$$
G(\mathbf{x})=\left\{p \in P:\left|\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)\right| \leq \Delta_{p}\right\}
$$

where the upper bound is $\Delta_{p}=\frac{\operatorname{cost}\left(p, b_{p}\right)}{\epsilon}$. Good points we can bound as before. For bad points we can show that while the difference in cost may be larger than $\operatorname{cost}\left(p, b_{p}\right) / \epsilon$, it must still be small, namely $O\left(\epsilon \min \left\{\operatorname{cost}(p, \mathbf{x}), \operatorname{cost}\left(b_{p}, \mathbf{x}\right)\right\}\right)$.

Formally, the functions $f_{\mathbf{x}}(p)$ are restricted to be defined only over good points:

$$
f_{\mathbf{x}}(p)= \begin{cases}\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)+\Delta_{p} & \text { if } p \in G(\mathbf{x}) \\ 0 & \text { otherwise }\end{cases}
$$

Then $\sum_{p \in P} \operatorname{cost}(p, \mathbf{x})-\sum_{q \in S \cup B} w_{q} \operatorname{cost}(q, \mathbf{x})$ is decomposed into three terms:

$$
\begin{align*}
& \sum_{p \in P} f_{\mathbf{x}}(p)-\sum_{q \in S} w_{q} f_{\mathbf{x}}(q)  \tag{15}\\
& +\sum_{p \in P \backslash G(\mathbf{x})}\left[\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)+\Delta_{p}\right]  \tag{16}\\
& -\sum_{q \in S \backslash G(\mathbf{x})} w_{q}\left[\operatorname{cost}(q, \mathbf{x})-\operatorname{cost}\left(b_{q}, \mathbf{x}\right)+\Delta_{q}\right] \tag{17}
\end{align*}
$$

Lemma 16 bounds (15) by $O(\epsilon) \operatorname{cost}(P, \mathbf{x})$, but we need an accuracy of $\epsilon^{2}$ to compensate for the $1 / \epsilon$ factor in the upper bound, resulting in a $O\left(1 / \epsilon^{4}\right)$ factor in the sample complexity.

We begin by bounding (16). Note that for each term in (16), $\mid \operatorname{cost}(p, \mathbf{x})-$ $\operatorname{cost}\left(b_{p}, \mathbf{x}\right) \mid>\Delta_{p}$ since $p \notin G(\mathbf{x})$. Furthermore, $p \notin G(\mathbf{x})$ only when $p$ and $b_{p}$ are close to each other and far away from $\mathbf{x}$. In Lemma 25 we use this to show that $\left|\operatorname{cost}(p, \mathbf{x})-\operatorname{cost}\left(b_{p}, \mathbf{x}\right)\right| \leq O(\epsilon) \min \left\{\operatorname{cost}(p, \mathbf{x}), \operatorname{cost}\left(b_{p}, \mathbf{x}\right)\right\}$. The details are presented in the appendix.

Using Lemma $25,(16)$ can be bounded by $O(\epsilon) \sum_{p \in P \backslash G(\mathbf{x})} \operatorname{cost}(p, \mathbf{x}) \leq O(\epsilon) \operatorname{cost}(P, \mathbf{x})$.

Similarly, by the definition of $\Delta_{q}$ and Lemma $25,(17)$ is bounded by

$$
\begin{aligned}
(17) & \leq \sum_{q \in S \backslash G(\mathbf{x})} 2 w_{q}\left|\operatorname{cost}(q, \mathbf{x})-\operatorname{cost}\left(b_{q}, \mathbf{x}\right)\right| \leq O(\epsilon) \sum_{q \in S \backslash G(\mathbf{x})} w_{q} \operatorname{cost}\left(b_{q}, \mathbf{x}\right) \\
& \leq O(\epsilon) \sum_{b \in B}\left(\sum_{q \in P_{b} \cap S} w_{q}\right) \operatorname{cost}(b, \mathbf{x}) .
\end{aligned}
$$

Note that the expectation of $\sum_{q \in P_{b} \cap S} w_{q}$ is $\left|P_{b}\right|$. By a sampling argument (Lemma 26), if $t \geq O\left(n k \log \frac{n k}{\delta}\right)$, then $\sum_{q \in P_{b} \cap S} w_{q} \leq 2\left|P_{b}\right|$. Then (17) is bounded by

$$
O(\epsilon) \sum_{b \in B} \operatorname{cost}(b, \mathbf{x})\left|P_{b}\right|=O(\epsilon) \sum_{p \in P} \operatorname{cost}\left(b_{p}, \mathbf{x}\right)
$$

where $\sum_{p \in P} \operatorname{cost}\left(b_{p}, \mathbf{x}\right)$ is at most a constant factor more than the optimum cost.
Since each of (15),(16), and (17) is $O(\epsilon) \operatorname{cost}(P, \mathbf{x})$, we know that their sum is the same magnitude. Combining the above bounds, we have

$$
\left|\operatorname{cost}(P, \mathbf{x})-\sum_{q \in S \cup B} w_{q} \operatorname{cost}(q, \mathbf{x})\right| \leq O(\epsilon) \operatorname{cost}(P, \mathbf{x}) .
$$

The proof is then completed by choosing a suitable $\epsilon$, and bounding $\operatorname{dim}(F, P)=$ $O(k d)$ as in [45].

### 3.2.2 Effect of Network Topology on Communication Cost

In the previous section, we presented a distributed coreset construction algorithm. The coreset constructed can then be used as a proxy for the original data, and we can run any distributed clustering algorithm on it. In this paper, we discuss the approach of simply collecting all local portions of the distributed coreset and run non-distributed clustering algorithm on it. If there is a central coordinator in the communication graph, then we can simply send the local portions of the coreset to the coordinator which can perform the clustering task. The total communication cost is just the size of the coreset.

Here we consider the distributed clustering tasks where the nodes are arranged in some arbitrary connected topology, and can only communicate with their neighbors.

We propose a message passing approach for globally sharing information, and use it for collecting information for coreset construction and sharing the local portions of the coreset. We also consider the special case when the graph is a rooted tree.

```
Algorithm 9 Distributed clustering on a graph
Input: \(\left\{P_{i}, 1 \leq i \leq s\right\}\) : local data sets; \(\left\{N_{i}, 1 \leq i \leq s\right\}\) : the neighbors of \(v_{i} ; \mathcal{A}_{\alpha}\) : an
    \(\alpha\)-approximation algorithm for weighted clustering instances.
    for each node \(v_{i} \in V\) do
        Construct its local portion \(D_{i}\) of an \(\epsilon / 2\)-coreset by Algorithm 8, using Message-
        Passing for communicating the local costs.
    end for
    for each node \(v_{i} \in V\) do
        Call Message-Passing \(\left(D_{i}, N_{i}\right)\). Let \(\mathbf{x}=\mathcal{A}_{\alpha}\left(\cup_{j} D_{j}\right)\).
    end for
Output: x.
```

```
Algorithm 10 Message-Passing \(\left(I_{i}, N_{i}\right)\)
Input: \(I_{i}\) is the message, \(N_{i}\) are the neighbors.
    Let \(R_{i}\) denote the information received.
    Initialize \(R_{i}=\left\{I_{i}\right\}\), and send \(I_{i}\) to all the neighbors.
    while \(R_{i} \neq\left\{I_{j}, 1 \leq j \leq s\right\}\) do
        if receive message \(I_{j} \notin R_{i}\) then
                \(R_{i}=R_{i} \cup\left\{I_{j}\right\}\) and send \(I_{j}\) to all the neighbors.
        end if
    end while
```


### 3.2.2.1 General Graphs

We now present the main result for distributed clustering on graphs.

Theorem 11. Given an $\alpha$-approximation algorithm for weighted $k$-means ( $k$-median respectively) as a subroutine, there exists an algorithm that with probability at least $1-\delta$ outputs a $(1+\epsilon) \alpha$-approximation solution for distributed $k$-means ( $k$-median respectively) clustering. The total communication cost is $O\left(m\left(\frac{1}{\epsilon^{4}}\left(k d+\log \frac{1}{\delta}\right)+n k \log \frac{n k}{\delta}\right)\right)$ for $k$-means, and $O\left(m\left(\frac{1}{\epsilon^{2}}\left(k d+\log \frac{1}{\delta}\right)+n k\right)\right)$ for $k$-median.

Proof. The details are presented in Algorithm 9. By Theorem 14, the output of Algorithm 8 is a coreset. Observe that in Algorithm 10, for any $j, I_{j}$ propagates on
the graph in a breadth-first-search style, so at the end every node receives $I_{j}$. This holds for all $1 \leq j \leq s$, so all nodes has a copy of the coreset at the end, and thus the output is a $(1+\epsilon) \alpha$-approximation solution.

Also observe that in Algorithm 10, for any node $v_{i}$ and $j \in[s], v_{i}$ sends out $I_{j}$ once, so the communication of $v_{i}$ is $\left|N_{i}\right| \times \sum_{j=1}^{s}\left|I_{j}\right|$. The communication cost of Algorithm 10 is $O\left(m \sum_{j=1}^{s}\left|I_{j}\right|\right)$. Then the total communication cost of Algorithm 9 follows from the size of the coreset constructed.

In contrast, an approach where each node constructs an $\epsilon$-coreset for $k$-means and sends it to the other nodes incurs communication cost of $\tilde{O}\left(\frac{m n k d}{\epsilon^{4}}\right)$. Our algorithm significantly reduces this.

### 3.2.2.2 Rooted Trees

Our algorithm can also be applied on a rooted tree, and compares favorably to other approaches involving coresets [104]. We can restrict message passing to operating along this tree, leading to the following theorem.

Theorem 12. Given an $\alpha$-approximation algorithm for weighted $k$-means ( $k$-median respectively) as a subroutine, there exists an algorithm that with probability at least $1-\delta$ outputs a $(1+\epsilon) \alpha$-approximation solution for distributed $k$-means ( $k$-median respectively) clustering on a rooted tree of height $h$. The total communication cost is $O\left(h\left(\frac{1}{\epsilon^{4}}\left(k d+\log \frac{1}{\delta}\right)+n k \log \frac{n k}{\delta}\right)\right)$ for $k$-means, and $O\left(h\left(\frac{1}{\epsilon^{2}}\left(k d+\log \frac{1}{\delta}\right)+n k\right)\right)$ for $k$-median.

Proof. We can construct the distributed coreset using Algorithm 8. In the construction, the costs of the local approximation solutions are sent from every node to the root, and the sum is sent to every node by the root. After the construction, the local portions of the coreset are sent from every node to the root. A local portion $D_{i}$ leads to a communication cost of $O\left(\left|D_{i}\right| h\right)$, so the total communication cost is
$O\left(h \sum_{i=1}^{s}\left|D_{i}\right|\right)$. Once the coreset is constructed at the root, the $\alpha$-approximation algorithm can be applied centrally, and the results can be sent back to all nodes.

Our approach improves the cost of $\tilde{O}\left(\frac{n h^{4} k d}{\epsilon^{4}}\right)$ for $k$-means and the cost of $\tilde{O}\left(\frac{n h^{2} k d}{\epsilon^{2}}\right)$ for $k$-median in [104] ${ }^{4}$. The algorithm in [104] builds on each node a coreset for the union of coresets from its children, and thus needs $O(\epsilon / h)$ accuracy to prevent the accumulation of errors. Since the coreset construction subroutine has quadratic dependence on $1 / \epsilon$ for $k$-median (quartic for $k$-means), the algorithm then has quadratic dependence on $h$ (quartic for $k$-means). Our algorithm does not build coreset on top of coresets, resulting in a better dependence on the height of the tree $h$.

In a general graph, any rooted tree will have its height $h$ at least as large as half the diameter. For sensors in a grid network, this implies $h=\Omega(\sqrt{s})$. In this case, our algorithm gains a significant improvement over existing algorithms.

### 3.3 Distributed $k$-Means Clustering of High Dimensional Data

In Algorithm 8, the number of points in the coreset is independent of the number of the original data points, which is useful for large-scale applications. However, it is linear in the dimension of the data, leading to high communication cost for high dimensional data. Here we propose a distributed PCA algorithm, and show that its output represents the original data in the sense that any good approximation solution of $k$-means clustering on the output projected data is also a good solution on the original data. When combined with the distributed coreset approach, this leads to an algorithm whose communication cost (in terms of the number of points communicated) is independent of the size and the dimension of the original data.

[^6]Related Work on Distributed PCA. Algorithms for (approximate) distributed PCA have been proposed $[93,13,75,83,46]$, but either without theoretical guarantees on the solution quality, or without considerations of the communication cost. Most closely related to our work is [46], which pointed out that the top singular vectors and eigenvalues of the local data set can be viewed as its summary and the union of the local summaries can be viewed as a summary of the global data. It implicitly showed the correctness of the method, but without considering the communication. [93] proposed a variant of the algorithm but provided no theoretical analysis on the tradeoff between communication and approximation guarantees.

In [62] the authors study efficient algorithms in the distributed model. Their model is different, namely, it is an arbitrary partition model in which each server holds a matrix $P_{i}$ and $P=\sum_{i=1}^{s} P_{i}$. Thus, each row of $P$ is additively shared across the $s$ servers, whereas in our model each row of $P$ belongs to a single server, though duplicate rows are allowed. Our model is motivated by applications in which points are indecomposable entities, such as the clustering application. We also focus on how the distributed PCA algorithm affects the quality of the final clustering solution, while [62] focus on getting a low rank approximation.

Other related work includes the recent [50] (see also the references therein), who give a deterministic streaming algorithm for approximate PCA in which each point of $P$ is seen one at a time and uses $O(d k / \epsilon)$ words of communication. Their algorithm naturally gives an $O(s d k / \epsilon)$ word communication algorithm in the distributed model. However, it involves an SVD computation for each point, making the overall computation expensive.

```
Algorithm 11 Distributed PCA
Input: local data \(\left\{P_{i}\right\}_{i=1}^{s}\) and parameter \(t \in \mathbf{N}_{+}\).
    for each node \(v_{i} \in V\) do
        Compute local SVD: \(P_{i}=U_{i} D_{i} V_{i}^{\top}\).
        BROADCAST \(D_{i}^{(t)}, V_{i}^{(t)}\).
    end for
    for each node \(v_{i} \in V\) do
        Set \(Y_{i}=D_{i}^{(t)}\left(V_{i}^{(t)}\right)^{\top}, Y^{\top}=\left[Y_{1}^{\top}, \ldots, Y_{s}^{\top}\right]\).
        Compute global SVD: \(Y=U D V^{\top}\).
        Compute projected data: \(\hat{P}_{i}=P_{i}^{(t)} V^{(t)}\left(V^{(t)}\right)^{\top}\).
    end for
Output: \(\left\{\hat{P}_{i}\right\}_{i=1}^{s}\).
```


### 3.3.1 Distributed PCA

Our distributed PCA algorithm is described in Algorithm 11, where BROADCAST is a shorthand for communicating information to all other nodes. The algorithm performs local PCA on each local data set, and communicates the $t$ largest principal components. These are then concatenated and used to get the $t$ largest global principal components. Finally, all the local data are projected on these $t$ global principal components. See Figure 11 for an illustration.

Figure 11: The key points of the distributed PCA algorithm.

Here we provide a theoretical analysis, which leads to a way to set the algorithm parameters, so that we will not compromise much on the quality of the clustering obtained on the projected data. In the following, we always assume for simplicity that $P$ is normalized, that is, $\sum_{p \in P} p=0$. We now show that the output of Algorithm 11 approximates the original data in the sense that their distances to low dimension subspaces are almost the same:

Theorem 13. Let $X$ be a $d \times j$ matrix whose columns are orthonormal. Let $\epsilon \in(0,1]$ and $t \in \mathbf{N}$ with $d-1 \geq t \geq j+\lceil 8 j / \epsilon\rceil-1$. Then the output of Algorithm 11 satisfies $0 \leq\|P X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2} \leq \epsilon d^{2}(P, L(X)) \quad$ and $\quad 0 \leq\|P X-\hat{P} X\|_{F}^{2} \leq \epsilon d^{2}(P, L(X))$.

Intuitively, it implies that the squared distances to any low dimension subspace $L(X)$ from the projected data and the original data are approximately equal when the number of principal components used is sufficiently large compared to the dimension of $L(X)$. As shown in the next section, this guarantees that the projected data can act as a proxy for the original data in $k$-means clustering.

The key in the proof of the theorem is a property about SVD: the projection $\tilde{A}$ of a $n \times d$ matrix $A$ to the subspace spanned by its first $t$ right singular vectors approximates $A$ in the following sense.

Lemma 18. Let $A \in \mathbf{R}^{n \times d}$ be an $n \times d$ matrix with singular value decomposition $A=U D V^{\top}$. Let $\epsilon \in(0,1]$ and $r, t \in \mathbf{N}_{+}$with $d-1 \geq t \geq r+\lceil r / \epsilon\rceil-1$, and let $\tilde{A}=A V^{(t)}\left(V^{(t)}\right)^{\top}$. Then for any matrix $X$ with $d$ rows and $\|X\|_{F}^{2} \leq r$, we have

$$
\|(A-\tilde{A}) X\|_{F}^{2}=\|A X\|_{F}^{2}-\|\tilde{A} X\|_{F}^{2} \leq \epsilon \sum_{i=r+1}^{d} \sigma_{i}^{2}(A)
$$

A special important case is that $X$ is the orthonormal basis for $r$-dimensional subspace, when the lemma reduces to Lemma 6.1 in [46]. In this case, the lemma means that the projections of $\tilde{A}$ and $A$ on any $r$-dimensional subspace are close, when the projected dimension $t$ is sufficiently large compared to $r$. The proof of the theorem will need a statement in the more general setting where $X$ may not be orthonormal. Proof Sketch of Theorem 13. We first introduce some auxiliary variables for the analysis, which act as intermediate connections between $P$ and $\hat{P}$. Imagine we perform two kinds of projections: first project $P_{i}$ to $\tilde{P}_{i}=P_{i} V_{i}^{(t)}\left(V_{i}^{(t)}\right)^{\top}$, then project $\tilde{P}_{i}$ to $\bar{P}_{i}=\tilde{P}_{i} V^{(t)}\left(V^{(t)}\right)^{\top}$. Let $\tilde{P}$ denote the vertical concatenation of $\tilde{P}_{i}$ and let $\bar{P}$ denote
the vertical concatenation of $\bar{P}_{i}$, i.e.

$$
\tilde{P}=\left[\begin{array}{c}
\tilde{P}_{1} \\
\vdots \\
\tilde{P}_{s}
\end{array}\right] \quad \text { and } \quad \bar{P}=\left[\begin{array}{c}
\bar{P}_{1} \\
\vdots \\
\bar{P}_{s}
\end{array}\right]
$$

These variables are designed such that the difference between $P$ and $\bar{P}$ is easily bounded. Our proof proceeds by first bounding the diference between $P$ and $\bar{P}$, and then bounding that between $\bar{P}$ and $\hat{P}$.

Consider $\|P X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2}$. It can be decomposed as followings:

$$
\begin{aligned}
\|P X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2} & =\left[\|P X\|_{F}^{2}-\|\tilde{P} X\|_{F}^{2}\right]+\left[\|\tilde{P} X\|_{F}^{2}-\|\bar{P} X\|_{F}^{2}\right] \\
& +\left[\|\bar{P} X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2}\right]
\end{aligned}
$$

The first term $\|P X\|_{F}^{2}-\|\tilde{P} X\|_{F}^{2}=\sum_{i=1}^{s}\left[\left\|P_{i} X\right\|_{F}^{2}-\left\|\tilde{P}_{i} X\right\|_{F}^{2}\right]$, each of which can be bounded by Lemma 18 since $\tilde{P}_{i}$ is the SVD truncation of $P$. The second term can be bounded similarly. The more difficult part is to bound the third term. Let $Z=V^{(t)}\left(V^{(t)}\right)^{\top} X$. Then by definition, $\bar{P}_{i}=\tilde{P}_{i} Z, \hat{P}_{i}=P_{i} Z$, and

$$
\|\bar{P} X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2}=\sum_{i=1}^{s}\left[\left\|\tilde{P}_{i} Z\right\|_{F}^{2}-\left\|P_{i} Z\right\|_{F}^{2}\right]
$$

Then Lemma 18 can be applied to show that $\left\|\tilde{P}_{i} Z\right\|_{F}^{2}-\left\|P_{i} Z\right\|_{F}^{2}$ is indeed small.
The bound on $\|P X-\hat{P} X\|_{F}^{2}$ can also be proved by a similar argument. The complete proof can be found in Appendix B.3.

### 3.3.2 Distributed Clustering

In this subsection, we show that any good approximation solution on the projected data constructed by the distributed PCA algorithm is also a good approximation on the original data.

Theorem 14. Let $\mathbf{x}$ be a set of $k$ centers in $\mathbf{R}^{d}$. Let $\epsilon \in(0,1]$ and $t \in \mathbf{N}$ with $d-1 \geq t \geq k+\left\lceil 50 k / \epsilon^{2}\right\rceil$. Then there exists a constant $c_{0} \geq 0$, such that the output of Algorithm 11 satisfies $(1-\epsilon) \operatorname{cost}(P, \mathbf{x}) \leq \operatorname{cost}(\hat{P}, \mathbf{x})+c_{0} \leq(1+\epsilon) \operatorname{cost}(P, \mathbf{x})$.

The analysis follows the ideas in [46]. Let $X \in \mathbf{R}^{d \times k}$ has orthonormal columns that span $\mathbf{x}$. The cost of $P$ can be decomposed into two parts: the squared distances $d^{2}(P, L(X))$ to the subspace spanned by $X$, and the squared distances $\sum_{i} d^{2}\left(\Pi_{X}\left(p_{i}\right), \mathbf{x}\right)$ between the projection of the points on $L(X)$ and $\mathbf{x}$. The cost of $\hat{P}$ can be decomposed similarly. Their difference in the first part (compensated by $c_{0}=\|P\|_{F}^{2}-\|\hat{P}\|_{F}^{2}$ ) can be bounded by $\|P X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2}$. The difference in the second part can be bounded approximately by $\sum_{i} d^{2}\left(\Pi_{X}\left(p_{i}\right), \Pi_{X}\left(\hat{p}_{i}\right)\right) / \epsilon=\|P X-\hat{P} X\|_{F}^{2} / \epsilon$. Then the theorem follows from Theorem 13. The complete proof is provided in the appendix.

By Theorem 14, the distributed coreset construction algorithm in the last section can be applied on the projected data to get a coreset of size independent of the original dimension. Then we get an algorithm with low communication cost for high dimensional data, which is summarized in Theorem 15.

Theorem 15. Given an $\alpha$-approximation algorithm for $k$-means as a subroutine, there exists an algorithm that with probability at least $1-\delta$ outputs a $(1+\epsilon) \alpha$ approximation solution for distributed $k$-means clustering. The total communication cost is $O\left(\frac{m s k}{\epsilon^{2}}\right)$ vectors in $\mathbf{R}^{d}$ plus $O\left(\frac{m}{\epsilon^{4}}\left(\frac{k^{2}}{\epsilon^{2}}+\log \frac{1}{\delta}\right)+m s k \log \frac{s k}{\delta}\right)$ vectors in $\mathbf{R}^{O\left(k / \epsilon^{2}\right)}$.

### 3.4 Experiments

The empirical study consists of two sets of experiments. The first set evaluates the coreset-based algorithm in Section 3.2, and the second set evaluates the approach in Section 3.3 that first reduces data dimension by distributed PCA and then applies the coreset-based algorithm.

### 3.4.1 Experiments on Distributed Coreset-Based Clustering

In these experiments, we seek to determine whether our algorithm is effective for the clustering tasks and how it compares to the other distributed coreset algorithms ${ }^{5}$.

[^7]We present the $k$-means cost of the solution produced by our algorithm with varying communication cost, and compare to those of other algorithms when they use the same amount of communication.

Data sets. Following the setup of [104, 12], for the synthetic data we randomly choose $k=5$ centers from the standard Gaussian distribution in $\mathbf{R}^{10}$, and sample equal number of 20,000 points from the Gaussian distribution around each center. Note that, as in $[104,12]$, we use the cost of the centers as a baseline for comparing the clustering quality. We choose the following real world data sets from [11]: Spam (4601 points in $\mathbf{R}^{58}$ ), Pendigits (10992 points in $\mathbf{R}^{16}$ ), Letter (20000 points in $\mathbf{R}^{16}$ ), and ColorHistogram of the Corel Image data set ( 68040 points in $\mathbf{R}^{32}$ ). We use $k=10$ for these data sets. We further choose YearPredictionMSD ( 515345 points in $\mathbf{R}^{90}$ ) for larger scale experiments, and use $k=50$ for this data set.

Experimental Methodology. To transform the centralized clustering data sets into distributed data sets we first generate a communication graph connecting local sites, and then partition the data into local data sets. To evaluate our algorithm, we consider several network topologies and partition methods.

The algorithms are evaluated on three types of communication graphs: random, grid, and preferential. The random graphs are Erdös-Renyi graphs $G(s, p)$ with $p=0.3$, i.e. they are generated by including each potential edge independently with probability 0.3 . The preferential graphs are generated according to the preferential attachment mechanism in the Barabási-Albert model [3]. For data sets Spam, Pendigits, and Letter, we use random/preferential graphs with 10 sites and $3 \times 3$ grid graphs. For synthetic data set and ColorHistogram, we use random/preferential graphs with 25 sites and $5 \times 5$ grid graphs. For large data set YearPredictionMSD, we use random/preferential graphs with 100 sites and $10 \times 10$ grid graphs.

The data is then distributed over the local sites. When the communication network also empirically outperforms other distributed coreset algorithms.
is a random graph, we consider three partition methods: uniform, similarity-based, and weighted. In the uniform partition, each data point in the global data set is assigned to the local sites with equal probability. In the similarity-based partition, each site has an associated data point randomly selected from the global data. Each data point in the global data is then assigned to the site with probability proportional to its similarity to the associated point of the site, where the similarities are computed by Gaussian kernel function. In the weighted partition, each local site is assigned a weight chosen by $|N(0,1)|$ and then each data point is distributed to the local sites with probability proportional to the site's weight. When the network is a grid graph, we consider the similarity-based and weighted partitions. When the network is a preferential graph, we consider the degree-based partition, where each point is assigned with probability proportional to the site's degree.

To measure the quality of the coreset generated, we run Lloyd's algorithm on the coreset and the global data respectively to get two solutions, and compute the ratio between the costs of the two solutions over the global data. The average ratio over 30 runs is then reported. We compare our algorithm with COMBINE, the method of combining a coreset from each local data set, and with the algorithm of [104] (Zhang et al.). When running the algorithm of Zhang et al., we restrict the general communication network to a spanning tree by picking a root uniformly at random and performing a breadth first search.

Results. Here we focus on the results of the largest data set YearPredictionMSD, and in Appendix B. 2 we present the experimental results for all the data sets.

Figure 12 shows the results over different network topologies and partition methods, where the subtitles indicate the network topology and the partition method (for example, "random, uniform" means the plot is for the random graph and the uniform partition method). We observe that the algorithms perform well with much smaller coreset sizes than predicted by the theoretical bounds. For example, to get 1.1 cost


Figure 12: Normalized $k$-means cost v.s. communication cost over graphs. The subtitles indicate the network topology and the partition method.
ratio, the coreset size and thus the communication needed is only $0.1 \%-1 \%$ of the theoretical bound.

In the uniform partition, our algorithm performs nearly the same as COMBINE. This is not surprising since our algorithm reduces to the COMBINE algorithm when each local site has the same cost and the two algorithms use the same amount of communication. In this case, since in our algorithm the sizes of the local samples are proportional to the costs of the local solutions, it samples the same number of points from each local data set. This is equivalent to the COMBINE algorithm with the same amount of communication. In the similarity-based partition, similar results are observed as it also leads to balanced local costs. However, when the local sites have significantly different costs (as in the weighted and degree-based partitions), our algorithm outperforms COMBINE. As observed in Figure 12, the costs of our solutions consistently improve over those of COMBINE by $2 \%-5 \%$. Our algorithm


Figure 13: Normalized $k$-means cost v.s. communication cost over the spanning trees of the graphs. The subtitles indicate the network topology and the partition method.
then saves $10 \%-20 \%$ communication cost to achieve the same approximation ratio.
Figure 13 shows the results over the spanning trees of the graphs. Our algorithm performs much better than the algorithm of Zhang et al., achieving about $20 \%$ improvement in cost. This is due to the fact that their algorithm needs larger coresets to prevent the accumulation of errors when constructing coresets from component coresets, and thus needs higher communication cost to achieve the same approximation ratio.

Similar results are observed on the other data sets, which are presented in Appendix B.2.

### 3.4.2 Experiments on High Dimensional Data

In these experiments we seek to understand how well the projected data approximates the original data, by measuring the $k$-means costs of the clustering solutions obtained
after dimension reduction.
Data sets. We choose the following real world data sets from [11]: Daily and Sports Activities data ( 9210 points in $\mathbf{R}^{5625}$ ), MNIST handwritten digits (70, 000 points in $\mathbf{R}^{784}$ ). We use $k=10$ for these data sets. We further choose Bag of Words (NYTimes) (300, 000 points in $\mathbf{R}^{102660}$ ) and use $k=20$ for this data set.

Experimental Methodology. Following the setup in the last subsection, we first generate a communication graph, which can be a grid graph, or a random graph that includes each edge independently with probability 0.3 . For Daily and Sports Activities data set, we use random graphs with 10 nodes and $3 \times 3$ grid graphs. For the other data sets, we use random graphs with 100 nodes and $10 \times 10$ grid graphs. Then we distribute the data over the graphs using weighted partition, where each data point is distributed to the local sites with probability proportional to the site's weight chosen from $|N(0,1)|$.

For each projection dimension, we first construct a coreset on the projected data, using the COMBINE or distributed coreset algorithm in the last section. After building the coreset, we then run Lloyd's method on it to get a $k$-means clustering solution. Finally, we compute the ratio of its cost to the $k$-means cost obtained by running Lloyd's method on the original data. The average results over 10 runs are reported. We lower the projection dimension until there is a significant increase in the $k$-means costs.

Results. Figure 14 shows the results of the data sets. The plots show the increase in $k$-means cost ratio upon decreasing the dimension of the data. We can observe that there is a slight increase compared to the huge reduction in dimension and thus communication cost. For example, on Daily and Sports Activities data, the $k$-means cost increases less than $4 \%$ when the dimension is reduced from 5625 to as low as 40. This is even more significant on higher dimensional data: on Bag of Words, the dimension can be reduced from 102660 to around 20. Such reduction then lowers


Figure 14: Normalized $k$-means cost v.s. communication cost for different projection dimensions. Rows: random graphs, grid graphs. Columns: Daily and Sports Activities, MNIST, and Bag of Words data sets. In each subfigure, the $x$-axis represents the communication cost, the $y$-axis represents the $k$-means cost (normalized by baseline), and the number labels are the projection dimensions.
the communication cost by magnitudes. The plots also indicate that the distributed coreset algorithm in the last section performs better than the COMBINE algorithm, when applied with our distributed PCA algorithm.

## CHAPTER IV

## COMMUNITY DETECTION

As we are entering into a network age, there is an increasing interest in analyzing network data in many disciplines, ranging from mathematics and computer science to sociology and biology. A significant amount of recent work in this area has focused on community detection, which is graph clustering viewing the network as a graph. The community structure reflects how entities in a network form meaningful groups such that interactions within the groups are more active compared to those between the groups and the outside world. The discovery of these communities is useful for understanding the structure of the underlying network, or making decisions in the network $[49,52,88,89]$.

Generally, a community should be thought of as a subset whose members have more interactions with each other than with the remainder of the network. This intuition is captured by some recently proposed models [34, 4, 21, 2, 58, 65]. Additionally, recent studies show that networks often exhibit hierarchical organization, in which communities can contain groups of sub-communities, and so forth over multiple scales. For example, this can be observed in ecological niches in food webs, modules in biochemical networks or groups of common interest in social websites [94, 71, 33]. It is also shown empirically and theoretically that hierarchical structures can simultaneously explain and quantitatively reproduce many commonly observed topological properties of networks [32, 97, 51]. This suggests that the hierarchical structure should also be reflected when modeling real world communities, which distinguishes community detection from the classic objective-based clustering approaches discussed in the previous chapters.

Although some heuristic approaches $[51,73]$ have been proposed to detect community hierarchies, few works have formalized this hierarchical property, and there are no theoretical performance guarantees for the algorithms. Inspired by the related work in clustering [16], in this chapter we define a notion of communities that reflects the tight connections within communities and explicitly models the hierarchy of communities. In our model, each member of a community falls into a sub-community, and the sub-communities within this community have active interactions with each other while entities outside this community have fewer interactions with members inside. Given this formalization, we then propose an efficient algorithm that detects all the communities in this model, and prove that all the communities form a hierarchy. Empirical evaluations demonstrate that our formalization successfully models real world communities, and our algorithm compares favorably with existing approaches.

### 4.1 Hierarchical Community Model

A network is typically represented as a graph $G=(V, E)$ on a set of $n=|V|$ points ${ }^{1}$, where the edges could be undirected or directed, unweighted or weighted. The graph implicitly specifies a neighborhood structure on the points, i.e. for each point there is a ranking of all other points according to the level of possible interaction. More precisely, we assume that we have a neighborhood function $N$ which given a point $p$ and a threshold $t$ outputs a list $N_{t}(p)$ containing the $t$ nearest neighbors of $p$ in $V$. Note that we do not assume the pairwise dissimilarity function for the points is a metric as in the previous chapters, but only assume the access to a neighborhood function.

The neighborhood function can be used to formalize a model of hierarchical communities. Using this neighborhood function, the tight connections within communities can be naturally rephrased as follows: for suitable $t$, most points $p$ in the community

[^8]have most of the nearest neighbors $N_{t}(p)$ from the community while points outside have just a few nearest neighbors from the community. Besides this, we also want to formalize the hierarchical structure that sub-communities in a lower, more local level actively interacting with each other form a community in a higher, more global level. The connections between the sub-communities can also be rephrased using the language of neighborhood: a majority of points in each sub-community have most of the nearest neighbors from the sub-communities in the same community.

In the remainder of the section, we specify our model based on the neighborhood function. We begin with the following notion of compact blobs, which will serve as a building block for our model.

Definition 12. A subset $A$ of points is called an $\alpha$-compact blob, if out of the $|A|$ nearest neighbors:

- any point $p \in A$ has at most $\alpha$ n neighbors outside $A$, i.e. $\left|N_{|A|}(p) \backslash A\right| \leq \alpha n$;
- any point $q \notin A$ has at most $\alpha n$ neighbors inside $A$, i.e. $\left|N_{|A|}(q) \cap A\right| \leq \alpha n$.


Figure 15: $\alpha$-compact blob. An edge $(x, y)$ means that $y$ is one of $x$ 's nearest neighbors.

Note that the notion of compact blobs is the same as the clusters that satisfy the $\alpha$-good neighborhood property defined in [16]. The notion captures the desired property of communities to be detected: members in the community have many more
interactions with other members inside the community and have fewer interactions with those outside. However, in practice, the notion may seem somewhat restricted. First, it requires all the members in the community have most interactions with other members inside the community, which may not be the case in real life. For example, some members in the boundary may have more interactions with the outside world, i.e. they have more than $\alpha n$ neighbors from outside. Based on this consideration, we define the $(\alpha, \beta)$-stable property as follows.

Definition 13. $A$ community $C$ is $(\alpha, \beta)$-stable if

- any point $p \in C$ falls into a $\alpha$-compact blob $A_{p} \subseteq C$ of size greater than $6 \alpha n$,
- for any point $p \in C$, at least $\beta$ fraction of points in $A_{p}$ have all but at most $\alpha n$ nearest neighbors from $C$ out of their $|C|$ nearest neighbors,
- any point $q$ outside $C$ has at most an nearest neighbors from $C$ out of their $|C|$ nearest neighbors.


Figure 16: $(\alpha, \beta)$-stable community. An edge $(x, y)$ means that $y$ is one of $x$ 's nearest neighbors. Note that point $w$ lies on the "boundary" of the community. It falls into the compact blob $A_{p}$, but does not have most of its nearest neighbors from the community.

Informally, the first condition means that every point falls into a sufficiently large compact blob in its community. This condition formalizes the local neighborhood structure that each member interacts actively with sufficiently many members in
the community. Note that the compact blob should be large enough so that the membership of the point is clearly established. That is, it should have size comparable to $\alpha n$, the number of connections to points outside. Here we choose a minimum size of $6 \alpha n$ mainly because it guarantees that our algorithm can still identify the blob in the worst case. The second condition means that at least $\beta$ fraction of points in these compact blobs have most of their nearest neighbors from the community. This condition formalizes more global neighborhood structure about how the compact blobs interact with each other to form a community. The third condition formalizes how the community is separated from the outside.

Note that we no longer require all the members in the community have most interactions inside; we only require each member interacts with sufficiently many members and a majority of members in these local groups interact actively. Also note that the definition is hierarchical in nature: sufficiently large compact blobs clearly satisfy the definition of $(\alpha, \beta)$-stable property and thus can be viewed as communities in lower levels. Furthermore, in the next section we will show that all the $(\alpha, \beta)$-stable communities form a hierarchy. We show this by presenting an algorithm and proving that each $(\alpha, \beta)$-stable community is a node in the hierarchy output by the algorithm. So our formulation explicitly models the hierarchical structure of communities observed in networks.

Next we propose a further generalization that considers possible noise in real world data. There may be some abnormal points that do not exhibit clear membership to any community, in the presence of which our definition above does not model the communities well. For example, suppose there is a point that has connections to all other points in the network, then no non-trivial subsets satisfy our definition above. We call such points bad since they do not fit into our community model above. To deal with the noise, we can naturally relax the $(\alpha, \beta)$-stable property to the $(\alpha, \beta, \nu)$-stable property defined as follows. Informally, it requires that the target community satisfies
the $(\alpha, \beta)$-stable property after removing a few bad points $B$. For convenience, we call the other points in $V \backslash B$ good points.

Definition 14. A community $C$ is $(\alpha, \beta, \nu)$-stable if there exist a subset of bad points $B$ of size at most $\nu n$, such that

- any good point $p \in G=C \backslash B$ falls into a compact blob $A_{p} \subseteq C$ of size greater than $6(\alpha+\nu) n$,
- for any point $p \in G$, at least $\beta$ fraction of points in $A_{p}$ have all but at most $\alpha n$ nearest neighbors from $G$ out of their $|G|$ nearest neighbors in $V \backslash B$,
- any good point $q$ outside $C \cup B$ has at most $\alpha$ n nearest neighbors from $G$ out of their $|G|$ nearest neighbors in $V \backslash B$.


Figure 17: $(\alpha, \beta, \nu)$-stable community. An edge $(x, y)$ means that $y$ is one of $x$ 's nearest neighbors. Point $b$ is a bad point and does not exhibit clear membership to any community.

Note 1 The parameters $\alpha, \nu$ are defined globally, that is, they are defined as ratios with respect to the total number of points. So a local change to some community can affect the values of these parameters for the other communities. For example, suppose we add $K n$ new points to some community, with all the new points having neighbors only inside this special community. Since the number of points increases
to $(K+1) n$, the communities outside the modified community are now $(\alpha /(1+$ $K), \beta, \nu /(1+K))$-stable. However, the local change does not affect the identifiability of these communities. Our algorithm described in the next section can still detect these communities, given the value of $(\alpha+\nu) n$.

Note 2 The input of the community detection task is usually a graph representing the network, and there are different ways to lift the graph to a neighborhood function. The simplest one is to directly sort for each point $p$ all the other points $q$ according to the weights of the edges $(p, q)$ and break ties randomly (we assume without loss of generality that the weights are in $[0,1]$ and the weight of an edge not in $E$ is regarded as 0). However, as pointed out in [21], we also have alternative approaches to convert the observed graph into a neighborhood function. More specifically, we assume the observed graph reflects some underlying unobserved set of relations, and thus we can lift the graph to an affinity system based on various beliefs about the connection between the latent relations and the observed graph, and then sort the points according to the affinity system to get the neighborhood function. For example, based on the belief that random walks on the graph can reflect the similarities between entities, we can define the affinity to be the diffusion kernel $\exp \{\lambda A\}$ where $A$ is the adjacent matrix and $\lambda$ is a parameter. Note that the results of appropriate lifting procedures can better reflect the true relationships between entities, and thus the conversion can address the challenging issue of sparsity in the observed graph.

### 4.2 Hierarchical Community Detection Algorithm

In the section, we propose an algorithm for detecting communities satisfying the $(\alpha, \beta, \nu)$-stable property. The goal of our algorithm is to output a set of communities such that each community satisfying the $(\alpha, \beta, \nu)$-stable property is close to one in the output. To be precise, we say that a community $C$ is $\nu$-close to another community $C^{\prime}$ if $\left|C \backslash C^{\prime}\right|+\left|C^{\prime} \backslash C\right| \leq \nu n$. We first describe the details in Algorithm 12, and then

```
Algorithm 12 Hierarchical Community Detection Algorithm
Input: Neighborhood function \(N\) on a set of points \(V, n=|V|, \alpha>0, \nu>0\).
    Initialize \(\mathcal{C}^{\prime}\) to be a set of singleton points, and \(t=6(\alpha+\nu) n+1\).
    while \(\left|\mathcal{C}^{\prime}\right|>1\) do
        Build \(F_{t}\) on \(V\) as follows.
        for any \(x, y \in V\) that satisfy \(\left|N_{t}(x) \cap N_{t}(y)\right| \geq t-2(\alpha+\nu) n\) do
            Connect \(x, y\) in \(F_{t}\).
        end for
        Build \(H_{t}\) on \(\mathcal{C}^{\prime}\) as follows. Let \(N_{F}(x)\) denote the neighbors of \(x\) in \(F_{t}\).
        for any \(U, W \in \mathcal{C}^{\prime}\) do
            if \(U, W\) are singleton subsets, i.e. \(U=\{x\}, W=\{y\}\) then
                Connect \(U, W\) in \(H_{t}\), if \(\left|N_{F}(x) \cap N_{F}(y)\right|>(\alpha+\nu) n\).
            else
                Set \(S_{t}(x, y)=\left|N_{F}(x) \cap N_{F}(y) \cap(U \cup W)\right|, \forall x \in U, y \in W\).
                Connect \(U, W\) in \(H_{t}\), if \(\operatorname{median}_{x \in U, y \in W} S_{t}(x, y)>\frac{|U|+|W|}{4}\).
            end if
        end for
        for any component \(R\) in \(H_{t}\) that satisfies \(\left|\cup_{C \in R} C\right| \geq 4(\alpha+\nu) n\) do
            Update \(\mathcal{C}^{\prime}\) by merging subsets in \(R\) into one subset.
        end for
        \(t=t+1\).
    end while
Output: Hierarchy \(T\) with single points as leaves and internal nodes corresponding to the merges performed.
```

present the analysis in Theorem 16.
Now we prove that the algorithm successfully outputs a hierarchy such that any community satisfying the $(\alpha, \beta, \nu)$-stable property with sufficiently large $\beta$ is close to one of the nodes in the hierarchy. Formally,

Theorem 16. Algorithm 12 outputs a hierarchy such that any community satisfying the $(\alpha, \beta, \nu)$-stable property with $\beta \geq 5 / 6$ is $\nu$-close to a node in the hierarchy. The algorithm runs in time $O\left(n^{\omega+1}\right)$, where $O\left(n^{\omega}\right)$ is the state of the art for matrix multiplication.

The correctness of the theorem follows from Lemma 21 and the running time follows from Lemma 22. In the following analysis, we always assume $\beta \geq 5 / 6$. Before presenting the analysis for the general communities in Lemma 21, we first prove a
lemma for the base case of compact blobs, showing that for any compact blob, a node close to it will be formed.

Lemma 19. For any good point $p$, when $t \leq\left|A_{p}\right|$, good points from $A_{p}$ will not be merged with good points outside $A_{p}$. At the end of the threshold $t=\left|A_{p}\right|$, all points in $A_{p}$ have been merged into a subset.

Proof. We prove this by induction on $t$. The claim is clearly true initially. Now assume for induction that at the beginning of a threshold $t \leq\left|A_{p}\right|$, in $\mathcal{C}^{\prime}$ good points from $A_{p}$ are not merged with good points outside $A_{p}$, i.e. any subset can contain good points from only one of $A_{p}$ and $V \backslash B \backslash A_{p}$. We now analyze the properties of the graphs $F_{t}$ and $H_{t}$, and show that at the end of the current threshold, the claim is still true.

First, as long as $t \leq\left|A_{p}\right|$, the graph $F_{t}$ has the following properties.

- No good point $x$ in $A_{p}$ is connected to a good point $y$ outside $A_{p}$. By the definition of compact blobs, out of the $t$ nearest neighbors, $x$ has at most $(\alpha+\nu) n$ neighbors outside $A_{p}$. For $y \in V \backslash B \backslash A_{p}, y$ has at most $(\alpha+\nu) n$ neighbors in $A_{p}$. Then $x, y$ have at most $2(\alpha+\nu) n<t-2(\alpha+\nu) n$ common neighbors, so they are not connected.
- No bad point $z$ is connected to both a good point $x$ in $A_{p}$ and a good point $y$ outside $A_{p}$. We know that out of the $t$ nearest neighbors, $x$ has at most $(\alpha+\nu) n$ neighbors outside $A_{p}$. So if $z$ is connected to $x$, then $z$ must have more than $t-3(\alpha+\nu) n$ neighbors in $A_{p}$ and less than $3(\alpha+\nu) n$ neighbors outside $A_{p}$. Since $y$ has at most $(\alpha+\nu) n$ neighbors in $A_{p}$, we have that $y, z$ share less than $3(\alpha+\nu) n+(\alpha+\nu) n<t-2(\alpha+\nu) n$ neighbors, so they are not connected.

Based on the properties of $F_{t}$ and the inductive assumption that any subset can contain good points from only one of $A_{p}$ and $V \backslash B \backslash A_{p}$, we show that the graph $H_{t}$ has the following properties.

- No subset $U$ containing good points from $A_{p}$ is connected to a subset $W$ containing good points outside $A_{p}$. This is clearly true if they are singleton subsets. In the other cases, note that the fraction of bad points in $U$ or $W$ is at most $1 / 4$. Then the number of pairs $(x, y)$ with good points $x \in U$ and $y \in W$ is at least $\frac{3}{4}|U| \times \frac{3}{4}|W|>|U||W| / 2$, i.e. more than half of the pairs $(x, y)$ with $x \in U$ and $y \in W$ are pairs of good points. This means there exist good points $x^{*} \in U, y^{*} \in W$ such that $S_{t}\left(x^{*}, y^{*}\right)$ is no less than $\operatorname{median}_{x \in U, y \in W} S_{t}(x, y)$. By the properties of $F_{t}, x^{*}, y^{*}$ have no common neighbors. Therefore, $U$ and $W$ are not connected.
- If a subset $W$ contains only bad points, then it cannot be connected to both a subset containing good points from $A_{p}$ and a subset containing good points outside $A_{p}$. Suppose it is connected to $U$ which contains good points from $A_{p}$. Note that since $W$ contains only bad points, it must contain only a single point $z$. If $U=\{x\}$ is singleton, then $x, z$ share more than $(\alpha+\nu) n$ neighbors in $F_{t}$. Since in $F_{t}, x$ is only connected to good points from $A_{p}$ and bad points, $z$ and $x$ must share some common neighbors from $A_{p}$, then $z$ must be connected to some good points in $A_{p}$. In the other cases, note that the fraction of bad points in $U$ is at most $1 / 4$. So there exists a good point $x^{*} \in U$ such that $S_{t}\left(x^{*}, z\right) \geq \operatorname{median}_{x \in U} S_{t}(x, z)$. Then we have $S_{t}\left(x^{*}, z\right)>(|U|+|W|) / 4>\nu n$, and thus $z$ must also be connected to some good points in $A_{p}$. Similarly, if $W$ is connected to a subset containing good points outside $A_{p}$, then the point in $W$ must connect to some good point outside $A_{p}$. But this is contradictory to the fact that in $F_{t}$ no bad point is connected to both a good point in $A_{p}$ and a good point outside $A_{p}$.

By the properties of $H_{t}$, no connected component contains both good points in $A_{p}$ and good points outside $A_{p}$. So at the end of this threshold $t$, the claim is still
true. Then by induction, we know that when $t \leq\left|A_{p}\right|$, we will not merge good points from $A_{p}$ with good points outside $A_{p}$.

Next we show that at the end of the threshold $t=\left|A_{p}\right|$, we will merge all points in $A_{p}$ into a subset. First, at this threshold, all good points in $A_{p}$ are connected in $F_{t}$. Any good point in $A_{p}$ has at most $(\alpha+\nu) n$ neighbors outside $A_{p}$, so when $t=\left|A_{p}\right|$, any two good points $x, y$ in $A_{p}$ are connected, and thus they share at least $\left|A_{p}\right|$ common neighbors in $F_{t}$. Second, all subsets containing good points in $A_{p}$ are connected in $H_{t}$. If no good points in $A_{p}$ have been merged, then these singleton points will be connected in $H_{t}$ since they share at least $\left|A_{p}\right|$ singleton subsets as common neighbors in $F_{t}$. If some good points in $A_{p}$ have already been merged into non-singleton subsets, we can show that in $H_{t}$ these non-singleton subsets will be connected to each other and connected to singleton subsets containing good points from $A_{p}$. For any such pair of subsets $U$ and $W$, the fraction of bad points in $U$ or $W$ is at most $1 / 4$, so there exist good points $x^{*} \in U, y^{*} \in W$ such that median $\operatorname{meU}_{x \in y \in W} S_{t}(x, y)$ is no less than $S_{t}\left(x^{*}, y^{*}\right)$. Since $x^{*}, y^{*}$ are connected to all good points in $A_{p}$ in $F_{t}, S_{t}\left(x^{*}, y^{*}\right)$ is no less than the number of good points in $U$ and $W$. So median ${ }_{x \in U, y \in W} S_{t}(x, y) \geq$ $S_{t}\left(x^{*}, y^{*}\right)>(|U|+|W|) / 4$, and thus $U, W$ are connected in $H_{t}$. Therefore, all points in $A_{p}$ are merged into a subset.

The following is a consequence of Lemma 19, which will be used in the analysis for the general communities in Lemma 21.

Lemma 20. In Algorithm 12, if a subset $U$ satisfies that for any good point $p \in U$, $A_{p} \subseteq U$, then there exist a subset of good points $P \subseteq U$, such that $\left\{A_{p}: p \in P\right\}$ is a partition of $U \backslash B$.

Proof. We have $U \backslash B=\cup_{p \in U \backslash B} A_{p}$. We only need to show that sets in $\left\{A_{p}: p \in U \backslash B\right\}$ are laminar, i.e. for any $p, q \in U \backslash B$, either $A_{p} \cap A_{q}=\emptyset$ or $A_{p} \subseteq A_{q}$ or $A_{q} \subseteq A_{p}$.

Assume for contradiction that there exist $A_{p}$ and $A_{q}$ such that $A_{p} \backslash A_{q} \neq \emptyset, A_{q} \backslash A_{p} \neq \emptyset$ and $A_{p} \cap A_{q} \neq \emptyset$. Without loss of generality, suppose $\left|A_{p}\right| \leq\left|A_{q}\right|$. Then by Lemma 19, at the end of the threshold $t=\left|A_{p}\right|$, we have merged all good points in $A_{p}$ into a subset. Specifically, this means that we have merged $A_{p} \cap A_{q}$ with $A_{p} \backslash A_{q}$. So for $t \leq\left|A_{q}\right|$, we have merged good points in $A_{q}$ with good points outside $A_{q}$, which is contradictory to Lemma 19.

By the above lemmas, for any good point $p$, the subset $A_{p}$ will be formed before points in it are merged with good points outside. Once these subsets are formed, we can show that subsets in the same target community will be merged together before they are merged with those from other communities, and thus the hierarchy produced has a node close to the target community. Formally, we have the following result.

Lemma 21. For any community $C$ satisfying the ( $\alpha, \beta, \nu$ )-stable property with $\beta \geq$ 5/6, $\mathcal{C}^{\prime} \backslash B$ in Algorithm 12 is always laminar to $C \backslash B$, i.e. for any $C^{\prime} \in \mathcal{C}^{\prime}$, either $\left(C^{\prime} \backslash B\right) \cap(C \backslash B)=\emptyset$ or $\left(C^{\prime} \backslash B\right) \subseteq(C \backslash B)$ or $(C \backslash B) \subseteq\left(C^{\prime} \backslash B\right)$. Furthermore, there is a node $U$ in the hierarchy produced such that $U \backslash B=C \backslash B$.

Proof. we will show by induction on $t$ that: for any community $C$ satisfying the ( $\alpha, \beta, \nu$ )-stable property with $\beta \geq 5 / 6$,

- at the end of threshold $t, \mathcal{C}^{\prime} \backslash B$ is laminar to $C \backslash B$,
- at the end of threshold $t$, for any $C$ such that $|C \backslash B| \leq t$, we have merged all points in $C \backslash B$ into a subset.

These claims are clearly true initially. Assume for induction that they are true for the threshold $t-1$, we now show that they are also true for the threshold $t$.

We first show that the laminarity is preserved. The laminarity is broken only when we connect in $H_{t}$ two subsets $U, W$ such that $U$ is a strict subset of $C$ after removing the bad points, and $W$ is a subset containing good points from outside. If there is a
good point $p \in U$ such that $A_{p} \nsubseteq U$, then by Lemma 19 , they cannot be connected. So we only need to consider the other case when for any good point $p \in U, A_{p} \subseteq U$. For convenience, we call a point great if it is a good point in $C$, and it has less than $\alpha n$ neighbors outside $C \backslash B$ out of the $|C \backslash B|$ nearest neighbors in $V \backslash B$. We now show that $U, W$ are not connected in $H_{t}$. Since $U \backslash B$ is a strict subset of $C \backslash B$, by induction on the second claim, we have $t \leq|C \backslash B|$. Then great points in $U$ and points in $W$ share at most $2(\alpha+\nu) n<t-2(\alpha+\nu) n$ common neighbors, so they are not connected in $F_{t}$. By Lemma 20 and the second condition of the ( $\alpha, \beta, \nu$ )-stable property, we know that at least $5 / 6$ fraction of points in $U \backslash B$ are great points. Then there exist a great point $x^{*} \in U$ and a point $y^{*} \in W$ such that $S_{t}\left(x^{*}, y^{*}\right)$ is no less than median $\operatorname{ref}_{x \in y \in W} S_{t}(x, y)$. Since in $F_{t}$ great points in $U$ are not connected to points in $W$, we have $S_{t}\left(x^{*}, y^{*}\right) \leq(|U|+|W|) / 4$. So median ${ }_{x \in U, y \in W} S_{t}(x, y) \leq(|U|+|W|) / 4$ and $U, W$ are not connected in $H_{t}$. Therefore, the laminarity is preserved.

Next we show that at the end of the threshold $t=|C \backslash B|$, all points in $C \backslash B$ are merged into a subset. By Lemma 19, all good points in $C \backslash B$ are now in sufficiently large subsets. We claim that any two of these subsets $U, W$ are connected in $H_{t}$, and thus will be merged. Again by Lemma 20, we know at least $5 / 6$ fraction of points in $U \backslash B$ or $W \backslash B$ are great points, and thus there exist great points $x^{*} \in U, y^{*} \in W$ such that $S_{t}\left(x^{*}, y^{*}\right)$ is no more than median ${ }_{x \in U, y \in W} S_{t}(x, y)$. Notice that all great points in $U$ are connected to great points in $W$ in $F_{t}$, since they share at least $t-2(\alpha+\nu) n$ neighbors. Then $S_{t}\left(x^{*}, y^{*}\right) \geq 3(|U|+|W|) / 4>(|U|+|W|) / 4$, and thus $\operatorname{median}_{x \in U, y \in W} S_{t}(x, y)>(|U|+|W|) / 4$. Therefore, any two subsets containing good points from $C \backslash B$ are connected in $H_{t}$ and thus are merged.

So the two claims hold for all $t$, specially for $t=n$. Then the algorithm must stop after this threshold, and we have the lemma as desired.

Lemma 22. Algorithm 12 has a running time of $O\left(n^{\omega+1}\right)$.

Proof. To implement the algorithm, we introduce some data structures. For any
$x \in V$, if $y$ is within the $t$ nearest neighbors of $x$, let $I_{t}(x, y)=1$, otherwise $I_{t}(x, y)=$ 0. Initializing $I_{t}$ takes $O\left(n^{2}\right)$ time. Next we compute $C N_{t}(x, y)$, the number of common neighbors between $x$ and $y$. Notice that $C N_{t}(x, y)=\sum_{z \in V} I_{t}(x, z) I_{t}(y, z)$, so $C N_{t}=I_{t} I_{t}^{T}$. Then we can compute the adjacent matrix $F_{t}$ (overloading the notation for the graph $F_{t}$ ) from $C N_{t}$. These take $O\left(n^{\omega}\right)$ time.

To compute the graph $H_{t}$, we introduce the following data structures. Let $F S_{t}(x, y)=$ 1 if $x, y$ are singleton subsets and $F_{t}(x, y)=1$, and let $F S_{t}(x, y)=0$ otherwise. Let $N S_{t}=F S_{t}\left(F S_{t}\right)^{T}$, then for two singleton subsets $x, y, N S_{t}(x, y)$ is the number of singleton subsets they share as neighbors in common in $F_{t}$. Similarly, let $F C_{t}(x, y)=1$ if $x$ and $y$ are in the same subset and $F_{t}(x, y)=1$, and let $F C_{t}(x, y)=0$ otherwise. Let $S_{t}=F_{t}\left(F C_{t}\right)^{T}+F C_{t}\left(F_{t}\right)^{T}$, then for two points $x \in U, y \in W$ where $U, W$ are two non-singleton subsets, $S_{t}(x, y)$ is the number of points in $U \cup W$ they share as neighbors in common in $F_{t}$. Based on $N S_{t}$ and $S_{t}$ we can build the graph $H_{t}$. All these take $O\left(n^{\omega}\right)$ time.

When we perform merge or increase the threshold, we need to update the data structures, which takes $O\left(n^{\omega}\right)$ time. Since there are $O(n)$ merges and $O(n)$ thresholds, Algorithm 12 takes time $O\left(n^{\omega+1}\right)$ in total.

### 4.3 Local Community Detection Algorithm

For networks with millions of nodes, the algorithm in the last section is too computationally expensive. Fortunately, in practical networks, the communities often have size much smaller than that of the whole network. For example, a social network can have millions members, while a member typically has no more than hundreds of close friends. Therefore, it is often sufficient to consider the following problem:

- Given some $p$, which is a good point in a community $C$, and the size of $C$, how to identify $C$ in time independent of the size of the whole network?

```
Algorithm 13 Local Community Detection Algorithm
Input: Neighborhood function \(N\) on a set of points \(V\), a point \(p \in V\), an integer
    \(c>0\), parameters \(\alpha>0, \nu>0\).
    1: Let \(N_{t}(A)\) denote the union of the \(t\) nearest neighbors of points in \(A\).
    Set \(N^{2}=N_{2 c}\left(N_{2 c}(p)\right)\).
    2: Set \(C_{\text {in }}^{\prime}=\left\{q \in N^{2}:\left|N_{c}(q) \cap N_{6(\alpha+\nu) n}(p)\right| \geq 3(\alpha+\nu) n\right\}\).
    3: Set \(C^{\prime}=\left\{q \in N^{2}:\left|N_{c}(q) \cap C_{\text {in }}^{\prime}\right| \geq 3(\alpha+\nu) n\right\}\).
Output: \(C^{\prime}\).
```

Such algorithms are called local algorithms in the literature of community detection. Note that the independence of the size of the whole network can only be achieved under some assumption about the input, since reading the similarities between $p$ and the other points already takes time linear in the size of the network. Here we assume that the algorithm has access to a sorted neighborhood function $N_{t}$. Formally, we assume the following mild condition.
(A) For any $p$ and $t, N_{t}(p)$ can be computed in time $O(t)$.

This assumption can be satisfied by building a sorting list of nearest neighbors for each point, which can be precomputed once in time $O\left(n^{2} \log n\right)$ from the pairwise similarities or dissimilarities.

Here we design a local algorithm (described in Algorithm 13), which solves the aforementioned problem under Assumption (A). To analyze the algorithm, first recall the definition of the stable community $C$, which requires some good points in the community have all but at most $\alpha n$ nearest neighbors from $C \backslash B$ out of their $C \backslash B$ nearest neighbors in $V \backslash B$. Call them inner points of $C$ and denote as $C_{\text {in }}$ for convenience. Intuitively, the algorithm is designed such that $C_{\text {in }}^{\prime}$ is a good estimation of $C_{\text {in }}$. From this, the algorithm then builds $C^{\prime}$ such that $C^{\prime} \approx C$. The construction is localized to an approximate superset $N^{2}$ of $C$, so that the running time is independent of the size of the network.

Theorem 17. Suppose $C$ is an $(\alpha, \beta, \nu)$-stable community with $\beta \geq 5 / 6$. Given
$p \in C \backslash B$ and $c=|C|$, Algorithm 13 outputs $C^{\prime}$ such that $\left|C^{\prime} \backslash C\right|+\left|C \backslash C^{\prime}\right| \leq(\alpha+\nu) n$. The running time is $O\left(|C|^{4}\right)$ under Assumption ( $\boldsymbol{A}$ ).

Proof. First, $N^{2}$ is approximately a superset of $C$. More precisely, we have $\mid C \backslash B \backslash$ $N^{2} \mid \leq \alpha n$. By definition, at most $\alpha n$ points in $A_{p}$ are outside $N_{2 c}(p)$. Since at least $\beta \geq 5 / 6$ fraction of $A_{p}$ are inner points, $N_{2 c}(p)$ contains at least one inner point. Then $N^{2}=N_{2 c}\left(N_{2 c}(p)\right)$ contains all but at most $\alpha n$ good points in $C$. For analysis, denote these missing points as $M=C \backslash B \backslash N^{2}$.

Second, $C_{\text {in }}^{\prime}$ contains all inner points in $C \backslash M$ but contains no good points outside $C$. By definition, $\left|A_{p} \cap N_{6(\alpha+\nu) n}(p)\right| \geq 5(\alpha+\nu) n$. If $q \in C_{\text {in }}$, then $\left|A_{p} \backslash N_{c}(q)\right| \leq$ $(\alpha+\nu) n$. Combining the two leads to $\left|N_{c}(q) \cap N_{6(\alpha+\nu) n}(p)\right|>3(\alpha+\nu) n$ since $A_{p}$ has size at least $6(\alpha+\nu) n$. This means that $C_{\text {in }}^{\prime}$ contains all inner points in $C \backslash M$. To show that $C_{\mathrm{in}}^{\prime}$ contains no good points $q$ outside $C$, note that $\left|N_{c}(q) \cap A_{p}\right| \leq(\alpha+\nu) n$. Also, $\left|N_{6(\alpha+\nu) n}(p) \backslash A_{p}\right| \leq(\alpha+\nu) n$. So we have $\left|N_{c}(q) \cap N_{6(\alpha+\nu) n}(p)\right| \leq 2(\alpha+\nu) n$.

Finally, $C^{\prime}$ contains all good points in $C \backslash M$ but contains no good points outside $C$, which then means $\left|C^{\prime} \backslash C\right|+\left|C \backslash C^{\prime}\right| \leq(\alpha+\nu) n$. If $q \in C \backslash B$, then $\left|A_{q} \backslash N_{c}(q)\right| \leq$ $(\alpha+\nu) n$. Since $C_{\text {in }}^{\prime}$ contains all points in $C_{\text {in }} \backslash M$, and at least $\beta$ fraction of the points $A_{q}$ are in $C_{\text {in }}$, we have that $\left|A_{q} \cap C_{\text {in }}^{\prime}\right| \geq \beta\left|A_{q}\right|-|M| \geq 4(\alpha+\nu) n$. Then $\left|C_{\mathrm{in}}^{\prime} \cap N_{c}(q)\right| \geq 3(\alpha+\nu) n$, and thus $C^{\prime}$ contains all good points in $C \backslash M$. To show $C^{\prime}$ contains no good points $q$ outside $C$, note that $\left|N_{c}(q) \cap C\right| \leq(\alpha+\nu) n$. Since $C_{\text {in }}^{\prime} \subseteq C \cup B$, we have $\left|N_{c}(q) \cap C_{\text {in }}^{\prime}\right| \leq(\alpha+\nu) n+|B|<3(\alpha+\nu) n$.

To bound the running time, first note that $N^{2}$ can be computed in time $|C|^{2}$ under Assumption (A). To build $C_{\mathrm{in}}^{\prime}$, we can search over all $O\left(|C|^{2}\right)$ points in $N^{2}$. For each $q \in N^{2}$, finding $N_{c}(q)$ takes time $O(|C|)$; for each $q^{\prime} \in N_{c}(q)$, testing whether $q^{\prime} \in N_{6(\alpha+\nu) n(p)}$ takes time $O((\alpha+\nu) n)=O(|C|)$. Then the total time to build $C_{\text {in }}^{\prime}$ is $O\left(|C|^{4}\right)$. Similarly, building $C^{\prime}$ takes time $O\left(|C|^{4}\right)$. Therefore, the algorithm takes time $O\left(|C|^{4}\right)$.

### 4.4 Experiments

In this section, we present our experimental results on evaluating our model and algorithm. While our main concern is building theoretical model for communities, empirical study is valuable in verifying the model and providing guidance for further improvement. Therefore, we applied our algorithm on both real world and synthetic data sets.

Note that the networks are represented as graphs, and we need to lift the graphs to get neighborhood functions for our algorithm. We use two lifting approaches for our experiments. The first approach is direct lifting: first, for any $x, y$ set the affinity between $x$ and $y$ to be 1 if $(x, y) \in E$ and 0 otherwise; then for each $x$, sort all the other points according to the affinities; break ties randomly to avoid bias. The second approach is diffusion lifting: first set the affinity matrix $K$ between entities to be $K=\exp \{\lambda A\}$ where $\lambda=0.05$ and $A$ is the adjacent matrix of the graph; then for each $x$, sort all the other points according to the affinities.

For comparison, we implemented two other algorithms: the lazy random walk algorithm (LRW [99]) and the Girvan-Newman algorithm (GN [51]). The lazy random walk algorithm performs truncated random walk from a seed point in the network and outputs selected communities where the selection is guided by the walk distribution and conductance. The conductance has been widely used as a criterion for quantifying the tight connections within communities, and thus the comparison to the lazy random walk algorithm provides an evaluation on how well our model and algorithm capture this intuition. The GN algorithm repeatedly removes the edge with the maximum edge-betweenness and regards the created connected components as communities. Although no theoretical model of hierarchical communities is targeted, the algorithm builds a hierarchy during its execution. It has been shown that the algorithm performs remarkably well on modeling communities in real-world data sets $[51,89]$. We use the code from [28] for fast computation of edge betweenness in
the algorithm.
For algorithms with parameters, we run them multiple times with different values of parameters, and report the best result. More specifically, we run our algorithm using parameters $(\alpha+\nu)=\frac{i}{5 n}(i=1,2, \ldots, 5)$. For the lazy random walk algorithm, we enumerate the parameters $\theta_{0}=0.05 i(i=1, \ldots, 4)$ and $b=1,2, \ldots,\lceil\log m\rceil$. In each run, 100 seed points are generated uniformly at random, each of which leads to a community. Since not all communities are meaningful (e.g. a singleton subset or the entire set of points), communities containing less than 10 points or containing more than $n-10$ points are removed, and the rest communities are regarded as the output communities. We then evaluate the average error of the output communities. The error for a ground-truth community $C$ with respect to a set $\mathcal{C}$ of output communities is defined as

$$
\operatorname{error}(C, \mathcal{C})=\min _{C^{\prime} \in \mathcal{C}} \frac{\left|C \backslash C^{\prime}\right|+\left|C^{\prime} \backslash C\right|}{n}
$$

This criterion measures how well the ground-truth communities are recovered by the algorithm. We further note that our algorithm outputs fewer communities than the other algorithms in all the conducted experiments, and thus has advantage when they achieve similar performance.

### 4.4.1 Evaluation on Real-World Networks

To evaluate the accuracy of the proposed method, we conduct experiments on the following real world data sets ${ }^{2}$ : karate [103], dolphins [82], polbooks [68], and football [51].

Figure 18 shows the average error and running time of the algorithms. We observe that our algorithm with diffusion lifting achieves the best performance on 3 out of 4 data sets, and achieves performance comparable to the GN algorithm on the football data set. It recovers the ground truth communities remarkably well over all the data

[^9]

Figure 18: The average error and running time using our hierarchical community detection algorithm with direct lifting (HCD+direct) or diffusion lifting neighborhood function (HCD+diffusion), Lazy Random Walk (LRW [99]) and the Girvan-Newman algorithm (GN [51]). Note that the running time is in log scale.
sets. Our algorithm with direct lifting does not achieve good results. Note that this is due to the fact that diffusion lifting reflects the true neighborhood structure more accurately than direct lifting. More precisely, when we sort neighbors for a point $p$ in direct lifting, all points not adjacent to $p$ are ranked randomly. In fact some of them can be reached by a few steps and thus should be ranked as close neighbors, while others are actually far away from the point $p$. On the other hand, diffusion lifting leads to a neighborhood function that more accurately reflects the neighborhood information. The LRW algorithm has the worst performance, though it is the fastest. Our algorithm, especially with the diffusion lifting, runs 10-100 faster than the GN algorithm. Therefore, our algorithm with suitable neighborhood functions is the most favorable for detecting real world communities.

Table 1: The parameters of the synthetic data sets for community detection. $n / m$ : number of nodes/edges; $k /$ maxk: average/maximum degree of the nodes; $\operatorname{minc} / \operatorname{maxc}$ : minimum/maximum size of the lower level communities; $\min C / \max C$ : minimum/maximum size of the higher level communities.

| Data set | $n$ | $m$ | $k$ | $\operatorname{maxk}$ | $\operatorname{minc}$ | $\operatorname{maxc}$ | $\min C$ | $\max C$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| LF50 | 50 | $\approx 500$ | 10 | 15 | 10 | 15 | 20 | 30 |
| LF100 | 100 | $\approx 1500$ | 15 | 20 | 15 | 20 | 30 | 40 |
| LF150 | 150 | $\approx 3000$ | 20 | 30 | 20 | 30 | 40 | 60 |
| LF200 | 200 | $\approx 6000$ | 30 | 40 | 30 | 40 | 60 | 80 |



Figure 19: The average error on the synthetic networks

### 4.4.2 Evaluation on Synthetic Networks

Besides real-world networks, we further use the Lancichinetti-Fortunato (LF) bench$m^{3} \mathrm{k}^{3}$ graphs [72] to evaluate the performance of the algorithms. By varying the

[^10]

Figure 20: The running time on synthetic networks
parameters of the networks, we can analyze the behavior of the algorithms in detail. We generate four unweighted undirected benchmark networks with two level community hierarchies. The numbers of nodes are $50,100,150$ and 200 respectively, and some important parameters of the networks are given in Table 1. For each type of data set, we range the mixing parameter $\mu$ from 0.1 to 0.5 with a span of 0.1 , and set the low-level mixing parameter $\mu_{1}=\mu / 4$ and the high-level mixing parameter $\mu_{2}=\mu-\mu_{1}$, resulting in five networks. Generally, the higher the mixing parameter of a network is, the more difficult it is to reveal the community structure.

Figure 19 shows the average errors of the algorithms and Figure 20 shows the running time. Our algorithm with direct or diffusion lifting and the GN algorithm achieve similar results on all the benchmark networks. The errors of these algorithms are below $5 \%$, and hardly increase with the mixing parameter. This suggests that they recover the ground truth communities remarkably well even in the hard case when the members of the communities have significant connections with the outside. In contrast, the LRW algorithm does not recover the communities well, even though
it runs much faster than the other algorithms. Our algorithm runs about 50 times faster than the GN algorithm over all the data sets. These results are consistent with those observed on real world data sets, and again demonstrate the advantage of our algorithm.

## APPENDIX A

## CLUSTERING UNDER PERTURBATION RESILIENCE

## A. 1 Dynamic Programming to Find the Minimum Cost $k$ Cluster Pruning

The idea of using dynamic programming to find the optimal $k$-clustering in a tree of clusters is proposed in [9]. We can find the optimal clustering by examining the entire tree of clusters produced. Denote the cost of the optimal $m$-clustering of a tree node $R$ as $\operatorname{cost}(R, m)$. The optimal $m$-clustering of a tree node $R$ is either the entire subtree as one cluster ( $m=1$ ), or the minimum over all choices of $m_{1}$-clustering over its left subtree and $m_{2}$-clustering over its right subtree $(1<m \leq k)$, where $m_{1}, m_{2}$ are positive integers such that $m_{1}+m_{2}=m$. Therefore, we can traverse the tree bottom up, recursively solving the $m$-clustering problem for $1 \leq m \leq k$ for each tree node. The algorithm is presented in Algorithm 14. Suppose that computing the cost of a cluster takes time $t\left(t=O\left(n^{2}\right)\right.$ for $k$-median, $k$-means and min-sum). Since there are $O(n)$ nodes, and on each node $R$, computing $\operatorname{cost}(R, 1)$ takes time $t$, computing $\operatorname{cost}(R, m)(1<m \leq k)$ takes $O\left(k^{2}\right)$, in total the algorithm takes time $O\left(n t+n k^{2}\right)$.

Notice when $\mathcal{T}$ is a multi-branch tree and not suitable for dynamic programming, we need to turn it into a 2-branch tree $\mathcal{T}^{\prime}$ as follows. For each node with more than 2 children, for example, the node $R$ with children $R_{1}, R_{2}, \ldots, R_{t}(t>2)$, we first merge $R_{1}$ and $R_{2}$ into one node, then merge this node with $R_{3}$; repeat until we merge all nodes $R_{1}, R_{2}, \cdots, R_{t}$ into $R$. In this way, we get a 2-branch tree $\mathcal{T}^{\prime}$ and can run dynamic programming on it. Note that each pruning in $\mathcal{T}$ has a corresponding pruning in $\mathcal{T}^{\prime}$, so the minimum cost pruning of $\mathcal{T}^{\prime}$ has no greater cost than the minimum cost pruning of $\mathcal{T}$.

Also note that when the cost function is center-based, such as $k$-median, the algorithm essentially computes a center for the node $R$ when computing $\operatorname{cost}(R, 1)$. So it can output the centers together with the pruning.

```
Algorithm 14 Dynamic Programming in Tree of Clusters
Input: A tree of clusters \(\mathcal{T}\) on a data set \(P\), distance function \(d(\cdot, \cdot)\) on \(P, k\).
    Traverse \(\mathcal{T}\) bottom up.
    for each node \(R \in \mathcal{T}\) do
        Calculate \(\operatorname{cost}(R, 1)\). For \(1<m \leq k\), calculate \(\operatorname{cost}(R, m)\) as follows.
        if \(R\) is a leaf then
            \(\operatorname{cost}(R, m)=\operatorname{cost}(R, 1)\).
        else
            \(\operatorname{cost}(R, m)=\min \left[\operatorname{cost}\left(R_{1}, m_{1}\right)+\operatorname{cost}\left(R_{2}, m_{2}\right)\right]\), where \(R_{1}, R_{2}\) are \(R\) 's children
            and \(m_{1}+m_{2}=m\).
        end if
    end for
    Traverse backwards to get the \(k\)-clustering \(\mathcal{C}\) that achieves \(\operatorname{cost}(r, k)\) where \(r\) is
    the root.
Output: The \(k\)-clustering \(\mathcal{C}\).
```


## A.2 An Efficient Implementation of Algorithm 1

Here we show an efficient implementation of Algorithm 1, namely Algorithm 15. The Phase 1 of this implementation takes time only $O\left(n^{3}\right)$.

Notice at each merge step in Algorithm 1, we only need to find the two clusters with the minimum closure distance. So we hope to compute the minimum closure distance without computing all the distances between any two current clusters. First we notice the following facts.

Fact 3. In the execution of Algorithm 1, if d is the minimum closure distance for the current clustering, then
(1) there exist $c, p \in P$ such that $d=d(c, p)$;
(2) $d$ is no less than the minimum closure distances in previous clusterings.

Proof. For the first claim, let $c$ be the center of the ball in the definition of closure distance, and $p$ be the farthest point from the center in the ball, then $d=d(c, p)$.

```
Algorithm 15 Efficient Implementation of Algorithm 1
Input: Data set \(P\), distance function \(d(\cdot, \cdot)\) on \(P\).
    Sort all the pairwise distances in ascending order.
    for each \(p \in P\) and \(1 \leq i \leq n\) do
        Compute \(L^{p}, \chi(p, i)\). Then compute \(\chi^{*}(p, i)\) by Equation (18).
    end for
    Let the current clustering be \(n\) singleton clusters.
    for \(d(p, q)\) in ascending order do
        Suppose \(q=L_{i}^{p}\). Check if \(d(p, q)\) satisfies the three claims in Fact 4, where the
        third claim can be checked by verifying if \(\chi^{*}(p, i)=-1\).
        If so, merge all the clusters covered by \(\mathbb{B}(p, d(p, q))\).
    end for
10: Construct the tree \(\mathcal{T}\) with points as leaves and internal nodes corresponding to
    the merges performed.
11: Run dynamic programming on \(\mathcal{T}\) to get the minimum cost pruning \(\tilde{\mathcal{C}}\).
Output: The clustering \(\tilde{\mathcal{C}}\).
```

The second claim comes from the fact that the clusters in the current clustering are supersets of those in previous clusterings.

Fact 3 implies that we can check in ascending order the pairwise distances no less than the minimum closure distance in the last clustering, and determine if the checked pairwise distance is the minimum closure distance in the current clustering. More specifically, suppose we have some black-box method for checking if a pairwise distance is the minimum closure distance in the current clustering, we can perform the closure linkage as follows: sort the pairwise distances in a list in ascending order; start from the first distance in the list; check if the current distance is the minimum closure distance in the current clustering; if it is, merge clusters covered by the ball defined by the checked distance; continue to check the next distance in the list. So it is sufficient to design a method to determine if a pairwise distance is the minimum closure distance in the current clustering. Our method is based on the following facts.

Fact 4. In Algorithm 1, if $d(c, p)$ is the minimum closure distance for the current clustering, then
(1) at least 2 clusters intersects $\mathbb{B}(c, d(c, p))$;
(2) all the clusters intersecting $\mathbb{B}(c, d(c, p))$ are covered by $\mathbb{B}(c, d(c, p))$;
(3) for any $p^{\prime} \in \mathbb{B}(c, d(c, p)), q \notin \mathbb{B}(c, d(c, p)), d\left(c, p^{\prime}\right)<d\left(p^{\prime}, q\right)$.

Proof. The first claim and the third claim follow from the definition. We can prove the second claim by induction. This is trivial at the beginning. Suppose it is true up to any previous clustering, we prove it for the current clustering $\mathcal{C}^{\prime}$. We need to show that for any $C^{\prime} \in \mathcal{C}^{\prime}$ such that $C^{\prime} \cap \mathbb{B}(c, d(c, p)) \neq \varnothing, C^{\prime} \subseteq \mathbb{B}(c, d(c, p))$. If $c \in C^{\prime}$, then by definition, $C^{\prime} \subseteq \mathbb{B}(c, d(c, p))$. If $C^{\prime}$ is a single point set $\left\{c_{1}\right\}$, then trivially $C^{\prime} \subseteq \mathbb{B}(c, d(c, p))$. What is left is the case when $c \notin C^{\prime}$ and $C^{\prime}$ is generated by merging clusters in a previous step. Suppose when $C^{\prime}$ is formed, the closure distance between those clusters is defined by $c_{1} \in C^{\prime}$ and $p_{1}$. By induction, if $c \in \mathbb{B}\left(c_{1}, d\left(c_{1}, p_{1}\right)\right), c$ would have been merged into $C^{\prime}$ when $C^{\prime}$ is merged, which is contradictory to $c \notin C^{\prime}$. So we have $c \notin \mathbb{B}\left(c_{1}, d\left(c_{1}, p_{1}\right)\right)$, i.e. $d\left(c, c_{1}\right)>d\left(c_{1}, p_{1}\right)$. Then by the margin requirement of $\mathbb{B}\left(c_{1}, d\left(c_{1}, p_{1}\right)\right), d(c, q)>d\left(c_{1}, q\right)$ for any $q \in \mathbb{B}(c, d(c, p)) \cap C^{\prime}$. This further leads to $c_{1} \in \mathbb{B}(c, d(c, p))$, since otherwise by the margin requirement of $\mathbb{B}(c, d(c, p))$ and $q \in \mathbb{B}(c, d(c, p))$, we would have $d(c, q)<d\left(c_{1}, q\right)$. So for any point $q^{\prime} \in C^{\prime}$, since $d\left(c_{1}, q^{\prime}\right) \leq d\left(c_{1}, p_{1}\right)<d\left(c, c_{1}\right)$, we have $q^{\prime} \in \mathbb{B}(c, d(c, p))$ from the margin requirement, so $C^{\prime} \subseteq \mathbb{B}(c, d(c, q))$.

Notice if a pairwise distance satisfies the three claims, then it defines a closure distance for the clusters covered. So if we check the pairwise distances in ascending order, then the first one that satisfies the three claims must be the minimum closure distance in the current clustering. So we have a method to determine if a pairwise distance is the minimum closure distance.

However, naively checking the third claim in Fact 4 takes $O\left(n^{2}\right)$, which is still not good enough. We can refine this step since intuitively, for every $c$, if $d(c, q)$ comes after $d(c, p)$ in the distance list, then when checking $d(c, q)$, we can utilize the information obtained from checking $d(c, p)$. Specifically, for every $p \in P$, define $L^{p}=\left(L_{1}^{p}, \ldots, L_{n}^{p}\right)$ to be a sorted list of points in $P$, according to their distances to $p$ in ascending order.

Let $\chi^{*}(p, i)$ denote the index of the farthest point in $L^{p}$, which makes $d\left(p, L_{i}^{p}\right)$ fail the third claim in Fact 4. Formally, define $\chi^{*}(p, i)$ to be the maximum $j>i$ such that there exits $s \leq i$ satisfying $d\left(p, L_{s}^{p}\right) \geq d\left(L_{s}^{p}, L_{j}^{p}\right)$. If no such point $L_{j}^{p}$ exists, let $\chi^{*}(p, i)=-1$. Then $d\left(p, L_{i}^{p}\right)$ satisfies the third claim if and only if $\chi^{*}(p, i)=-1$, thus we turn the task of checking the claim into computing $\chi^{*}(p, i)$. In order to use the information obtained when previously checking $d\left(p, L_{i-1}^{p}\right)$, we compute $\chi^{*}(p, i)$ from $\chi^{*}(p, i-1)$. By the definition of $\chi^{*}$, what is new of $\chi^{*}(p, i)$ compared to $\chi^{*}(p, i-1)$ is just the maximum $j>i$ such that $d\left(p, L_{i}^{p}\right) \geq d\left(L_{i}^{p}, L_{j}^{p}\right)$. Define $\chi(p, i)$ to be the maximum $j>i$ such that $d\left(p, L_{i}^{p}\right) \geq d\left(L_{i}^{p}, L_{j}^{p}\right)$; if no such $j$ exists, let $\chi(p, i)=-1$. Then it is easy to verify that

$$
\begin{equation*}
\chi^{*}(p, i)=\max \left\{\chi^{*}(p, i-1), \chi(p, i)\right\} . \tag{18}
\end{equation*}
$$

It takes $O(n)$ time to compute $\chi(p, i)$, thus we can compute $\chi^{*}(p, i)$ for all $p \in P, 1 \leq$ $i \leq n$ in $O\left(n^{3}\right)$ time. The implementation is finally summarized in Algorithm 15.

## A. 3 Proofs for Bounding the Number of Bad Points for $k$-Median

We now present the proof of Fact 1, which is used in Claims 1 and 2 for translating $d^{\prime}$ to $d$. For notations used, see Figure 21 in Section 2.3 .1 for illustration. We begin with the following technical fact:

Fact 5. Suppose the clustering instance is $(\alpha, \epsilon)$-perturbation resilient. If $\min _{i}\left|C_{i}\right|>$ $\left(\frac{2}{\alpha-1}+3\right) \epsilon n+1$, then $c_{i}^{\prime} \neq c_{j}(\forall j \neq i)$.

Proof. The intuition is that if $c_{i}^{\prime}=c_{j}$, then under $d^{\prime}$, points in $W_{j}$ should be closer to $c_{j}^{\prime}$ than to $c_{i}^{\prime}=c_{j}$. So under $d$, these points are $\alpha$ time closer to $c_{j}^{\prime}$ than to $c_{j}$. This means that the distance between $c_{j}$ and $c_{j}^{\prime}$ is no so large compared to the average distance between $c_{j}$ and $W_{j}$ by the triangle inequality. On the other hand, this also means that $c_{j}$ has $(1-1 / \alpha) d\left(c_{j}, W_{j}\right)$ more cost than $c_{j}^{\prime}$ on $W_{j}$. Then $c_{j}$ should save this
cost on other parts of $C_{j}$. In other words, $c_{j}^{\prime}$ has at least $(1-1 / \alpha) d\left(c_{j}, W_{j}\right)$ more cost on these points than $c_{j}$ has. By the triangle inequality, the distance between $c_{j}$ and $c_{j}^{\prime}$ is much larger than the average distance between $c_{j}$ and $W_{j}$, which is contradictory.

Formally, assume for contradiction that $c_{i}^{\prime}=c_{j}$. To apply the intuition described above, we first need to show $c_{j}^{\prime} \neq c_{l}(\forall l)$. Clearly, $c_{j}^{\prime} \neq c_{j}$, since otherwise, moving all the points in $C_{j}^{\prime}$ to $C_{i}^{\prime}$ will not increase the cost, which violates $(\alpha, \epsilon)$-perturbation resilience. We also know that $c_{j}^{\prime} \neq c_{l}(l \neq j)$ since otherwise, there is $p \in W_{j}$, $d\left(c_{l}, p\right)=d\left(c_{j}^{\prime}, p\right) \leq d^{\prime}\left(c_{j}^{\prime}, p\right)<d^{\prime}\left(c_{i}^{\prime}, p\right)=d\left(c_{j}, p\right)$, which contradicts the fact that $p \in C_{j}$.

Now we can apply the intuition described above to show that $c_{i}^{\prime}=c_{j}$ and $c_{j}^{\prime} \neq$ $c_{l}(\forall l)$ lead to an contradiction. Note that points in $W_{j} \cup V_{j}=C_{j} \cap C_{j}^{\prime}$ are closer to $c_{j}^{\prime}$ than to $c_{i}^{\prime}=c_{j}$ under $d^{\prime}$. Then back to $d$, for any $p \in W_{j}$, since $c_{j}^{\prime} \neq c_{l}(\forall l), \alpha d\left(c_{j}^{\prime}, p\right)=$ $d^{\prime}\left(c_{j}^{\prime}, p\right) \leq d^{\prime}\left(c_{i}^{\prime}, p\right)=d\left(c_{j}, p\right)$, resulting in $d\left(c_{j}^{\prime}, W_{j}\right) \leq d\left(c_{j}, W_{j}\right) / \alpha$. Similarly, for any $p \in V_{j}, \alpha d\left(c_{j}^{\prime}, p\right)=d^{\prime}\left(c_{j}^{\prime}, p\right) \leq d^{\prime}\left(c_{i}^{\prime}, p\right)=\alpha d\left(c_{j}, p\right)$, resulting in $d\left(c_{j}^{\prime}, V_{j}\right) \leq d\left(c_{j}, V_{j}\right)$. These facts have two consequences.

First, since points in $W_{j}$ are $\alpha$ time closer to $c_{j}^{\prime}$ than to $c_{j}$, the distance between $c_{j}^{\prime}$ and $c_{j}$ is small:

$$
\begin{equation*}
d\left(c_{j}^{\prime}, c_{j}\right) \leq \frac{d\left(c_{j}^{\prime}, W_{j}\right)}{\left|W_{j}\right|}+\frac{d\left(c_{j}, W_{j}\right)}{\left|W_{j}\right|} \leq\left(1+\frac{1}{\alpha}\right) d\left(c_{j}, W_{j}\right) \tag{19}
\end{equation*}
$$

Second, since $c_{j}$ is the optimal center for $C_{j}=W_{j} \cup V_{j} \cup M_{j}$, it should save a lot of cost on $M_{j}$ compared to $c_{j}^{\prime}$, which suggests $c_{j}$ and $c_{j}^{\prime}$ would be far apart. Formally,

$$
d\left(c_{j}^{\prime}, C_{j}\right)=d\left(c_{j}^{\prime}, W_{j} \cup V_{j} \cup M_{j}\right) \geq d\left(c_{j}, C_{j}\right)=d\left(c_{j}, W_{j} \cup V_{j} \cup M_{j}\right)
$$

Since $d\left(c_{j}^{\prime}, W_{j}\right) \leq d\left(c_{j}, W_{j}\right) / \alpha$ and $d\left(c_{j}^{\prime}, V_{j}\right) \leq d\left(c_{j}, V_{j}\right)$, we have

$$
\begin{align*}
d\left(c_{j}^{\prime}, M_{j}\right)-d\left(c_{j}, M_{j}\right) & \geq d\left(c_{j}, W_{j}\right)-\frac{1}{\alpha} d\left(c_{j}, W_{j}\right), \\
\left|M_{i}\right| d\left(c_{j}^{\prime}, c_{j}\right) & \geq\left(1-\frac{1}{\alpha}\right) d\left(c_{j}, W_{j}\right) \tag{20}
\end{align*}
$$

When $\left|C_{j}\right|>\left(\frac{2}{\alpha-1}+3\right) \epsilon n+1$, we have $(1-1 / \alpha)\left|W_{j}\right|>(1+1 / \alpha)\left|M_{j}\right|$. Then Inequalities 20 and 19 lead to $d\left(c_{j}, c_{j}^{\prime}\right)=0$. This means $c_{j}=c_{j}^{\prime}$ which is a contradiction to the assumptions.

Fact 1. Suppose the clustering instance is $(\alpha, \epsilon)$-perturbation resilient and $\min _{i}\left|C_{i}\right|>$ $\left(\frac{2}{\alpha-1}+3\right) \epsilon n+1$. If $c_{i} \neq c_{i}^{\prime}$, then we have
$d^{\prime}\left(c_{i}^{\prime}, W_{i}\right) \geq \alpha d\left(c_{i}^{\prime}, W_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right), \quad d^{\prime}\left(c_{i}, W_{i}\right)=d\left(c_{i}, W_{i}\right)$,
$d^{\prime}\left(c_{i}^{\prime}, V_{i}\right)=\alpha d\left(c_{i}^{\prime}, V_{i}\right), \quad d^{\prime}\left(c_{i}, V_{i}\right)=\alpha d\left(c_{i}, V_{i}\right)$,
$d^{\prime}\left(c_{i}^{\prime}, A_{i}\right) \geq \alpha d\left(c_{i}^{\prime}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right), \quad d^{\prime}\left(c_{i}, A_{i}\right) \leq \alpha d\left(c_{i}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)+\alpha(1+\alpha) d\left(c_{i}^{\prime}, c_{i}\right)$.
Proof. These translations can be verified by the definition of $d^{\prime}$. In most cases, $d^{\prime}(\cdot, \cdot)=\alpha d(\cdot, \cdot)$; the only exceptions are the distances between $p$ and $c(p)$. The detailed verification is presented below.

Since $c_{i}^{\prime} \neq c_{i}$, and by Fact 5 , we know $c_{i}^{\prime} \neq c_{j}(\forall j)$. So when translating $d^{\prime}\left(c_{i}^{\prime}, C\right)(C$ is $W_{i}, V_{i}$ or $\left.A_{i}\right)$, we only need to check if $c\left(c_{i}^{\prime}\right) \in C$. For $W_{i}, d^{\prime}\left(c_{i}^{\prime}, W_{i}\right) \geq d^{\prime}\left(c_{i}^{\prime}, W_{i} \backslash\right.$ $\left.\left\{c\left(c_{i}^{\prime}\right)\right\}\right)=\alpha d\left(c_{i}^{\prime}, W_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)$. For $V_{i}$, since there is no center in $V_{i}, d^{\prime}\left(c_{i}^{\prime}, V_{i}\right)=$ $\alpha d\left(c_{i}^{\prime}, V_{i}\right)$. For $A_{i}, d^{\prime}\left(c_{i}^{\prime}, A_{i}\right) \geq d^{\prime}\left(c_{i}^{\prime}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)=\alpha d\left(c_{i}^{\prime}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)$.

Now consider the sum of distances concerning $c_{i}$. For $d^{\prime}\left(c_{i}, W_{i}\right)$ and $d^{\prime}\left(c_{i}, V_{i}\right)$, the equations follow from the definition of $d^{\prime}$. For $A_{i}$, if $c\left(c_{i}^{\prime}\right) \notin A_{i}$, then the inequality is trivial. If $c\left(c_{i}^{\prime}\right) \in A_{i}$, then

$$
d^{\prime}\left(c_{i}, A_{i}\right)=d^{\prime}\left(c_{i}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)+d^{\prime}\left(c_{i}, c\left(c_{i}^{\prime}\right)\right) \leq \alpha d\left(c_{i}, A_{i} \backslash\left\{c\left(c_{i}^{\prime}\right)\right\}\right)+\alpha d\left(c_{i}, c\left(c_{i}^{\prime}\right)\right)
$$

We have $d\left(c_{i}, c\left(c_{i}^{\prime}\right)\right) \leq d\left(c_{i}, c_{i}^{\prime}\right)+d\left(c_{i}^{\prime}, c\left(c_{i}^{\prime}\right)\right)$. If $c_{i}^{\prime}$ is a selected bad point, then $d\left(c_{i}^{\prime}, c\left(c_{i}^{\prime}\right)\right) \leq \alpha d\left(c_{i}^{\prime}, c_{i}\right)$. Otherwise, $c\left(c_{i}^{\prime}\right)$ is the nearest center for $c_{i}^{\prime}$, then $d\left(c_{i}^{\prime}, c\left(c_{i}^{\prime}\right)\right) \leq$ $d\left(c_{i}^{\prime}, c_{i}\right)$. In any case, the inequality for $d^{\prime}\left(c_{i}, A_{i}\right)$ follows.

## A. 4 Structural Properties of $(\alpha, \epsilon)$-Perturbation Resilient Min-Sum Instances

## A.4.1 Proofs for Bounding the Number of Bad Points for Min-Sum

Recall the definitions of the bad points and the perturbation constructed to bound the number of bad points in Section 2.5.1.1. Suppose an intervals $[r, 2 r]$ contains the costs of more than $2 \epsilon n$ bad points. Let $\hat{B}$ denote a subset of $2 \epsilon n$ bad points in this interval, and let $\hat{B}_{i}=\hat{B} \cap C_{i}, K_{i}=C_{i} \backslash \hat{B}_{i}$. Suppose for a bad point $p \in \hat{B}_{i}, C_{j}$ is its second nearest cluster, that is, $j=\arg \min _{\ell \neq i} d\left(p, C_{\ell}\right)$. Denote this as $p \in D_{j}$. Let $\tilde{C}_{i}=K_{i} \cup D_{i}$. The perturbation is constructed as follows: blow up all distances by a factor of $\alpha$ except those within $\tilde{C}_{i}, 1 \leq i \leq k$. Let $\left\{C_{i}^{\prime}\right\}$ denote the optimal clustering after perturbation. Recall the definitions of $U_{i}, V_{i}, W_{i}$ and $\tilde{U}_{i}, \tilde{V}_{i}, \tilde{W}_{i}$, and see Figure 6 for an illustration. The following facts come from their definitions.

Fact 6. We have $\cup_{i} U_{i}=\cup_{i} \tilde{U}_{i}, \cup_{i} V_{i}=\cup_{i} \tilde{V}_{i}$ and $\cup_{i} W_{i}=\cup_{i} \tilde{W}_{i}$. Furthermore,

$$
\begin{aligned}
\sum_{i} d\left(\tilde{U}_{i}, C_{i}\right) & \leq \beta \sum_{i} d\left(U_{i}, C_{i}\right) \\
\sum_{i} d\left(\tilde{V}_{i}, C_{i}\right) & \leq \sum_{i} d\left(V_{i}, C_{i}\right) \\
\sum_{i} d\left(\tilde{W}_{i}, C_{i}\right) & \leq \sum_{i} d\left(W_{i}, C_{i}\right)
\end{aligned}
$$

We are ready to prove the claim needed for bounding the number of bad points.

Claim 3.(a). The costs saved and added by moving $\left\{U_{i}, 1 \leq i \leq k\right\}$ satisfy

$$
\begin{aligned}
& 2 \sum_{i} d^{\prime}\left(U_{i}, C_{i}^{\prime} \cap K_{i}\right)-2 \sum_{i} d^{\prime}\left(\tilde{U}_{i}, \tilde{C}_{i}\right) \\
\geq & \frac{3}{10} \alpha \sum_{i} d\left(U_{i}, C_{i}\right)-\frac{2 \alpha}{100} \sum_{i} d\left(W_{i}, C_{i}\right)-\frac{8 \alpha+16}{100} r \epsilon n .
\end{aligned}
$$

Proof. Intuitively, $\sum_{i} d^{\prime}\left(U_{i}, C_{i}^{\prime} \cap K_{i}\right) \approx \alpha \sum_{i} d\left(U_{i}, C_{i}\right)$ and $\sum_{i} d^{\prime}\left(\tilde{U}_{i}, \tilde{C}_{i}\right) \approx \sum_{i} d\left(\tilde{U}_{i}, C_{i}\right)$. Their difference is then roughly $(\alpha-\beta) \sum_{i} d\left(U_{i}, C_{i}\right)$, since $\sum_{i} d\left(\tilde{U}_{i}, C_{i}\right) \leq \beta \sum_{i} d\left(U_{i}, C_{i}\right)$.

Formally, we have

$$
\begin{align*}
d^{\prime}\left(U_{i}, C_{i}^{\prime} \cap K_{i}\right) & =\alpha d\left(U_{i}, C_{i}^{\prime} \cap K_{i}\right)=\alpha d\left(U_{i}, C_{i}\right)-\alpha d\left(U_{i}, \tilde{W}_{i}+\hat{B}_{i}\right),  \tag{21}\\
d^{\prime}\left(\tilde{U}_{i}, \tilde{C}_{i}\right) & =d\left(\tilde{U}_{i}, \tilde{C}_{i}\right)=d\left(\tilde{U}_{i}, K_{i}\right)+d\left(\tilde{U}_{i}, D_{i}\right) \leq d\left(\tilde{U}_{i}, C_{i}\right)+d\left(\tilde{U}_{i}, D_{i}\right) \tag{22}
\end{align*}
$$

Then it suffices to bound the approximation error $d\left(U_{i}, \tilde{W}_{i}+\hat{B}_{i}\right)$ and $d\left(\tilde{U}_{i}, D_{i}\right)$. First,

$$
\begin{aligned}
d\left(U_{i}, \tilde{W}_{i}+\hat{B}_{i}\right) & \leq \frac{\left|\tilde{W}_{i}+\hat{B}_{i}\right|}{\left|C_{i}\right|} d\left(U_{i}, C_{i}\right)+\frac{\left|U_{i}\right|}{\left|C_{i}\right|} d\left(C_{i}, \tilde{W}_{i}+\hat{B}_{i}\right) \\
& \leq \frac{3}{100} d\left(U_{i}, C_{i}\right)+\frac{1}{100} d\left(C_{i}, \tilde{W}_{i}+\hat{B}_{i}\right)
\end{aligned}
$$

where the first inequality is by Fact 7 , and the second is from the fact that $\left|\tilde{W}_{i}\right| \leq$ $\epsilon n,\left|\hat{B}_{i}\right| \leq 2 \epsilon n,\left|U_{i}\right| \leq \epsilon n$ and $\left|C_{i}\right| \geq 100 \epsilon n$. For the second term on the right hand side, we have $\sum_{i} d\left(C_{i}, \tilde{W}_{i}\right) \leq \sum_{i} d\left(W_{i}, C_{i}\right)$, and points in $\hat{B}_{i}$ has cost at most $2 r$ : $d\left(\hat{B}_{i}, C_{i}\right) \leq 2 r\left|\hat{B}_{i}\right|$. So

$$
\sum_{i} d\left(U_{i}, \tilde{W}_{i}+\hat{B}_{i}\right) \leq \frac{3}{100} \sum_{i} d\left(U_{i}, C_{i}\right)+\frac{1}{100} \sum_{i} d\left(C_{i}, W_{i}\right)+\frac{4 r \epsilon n}{100}
$$

Similarly, for $d\left(\tilde{U}_{i}, D_{i}\right)$ we have

$$
\sum_{i} d\left(\tilde{U}_{i}, D_{i}\right) \leq \sum_{i}\left[\frac{\left|D_{i}\right|}{\left|C_{i}\right|} d\left(\tilde{U}_{i}, C_{i}\right)+\frac{\left|\tilde{U}_{i}\right|}{\left|C_{i}\right|} d\left(C_{i}, D_{i}\right)\right] \leq \frac{2 \beta}{100} \sum_{i} d\left(U_{i}, C_{i}\right)+\frac{8 r \epsilon n}{100}
$$

The claim follows by summing (21) and (22) over $1 \leq i \leq k$ and plugging the last two inequalities.

Claim 3.(b). The costs saved and added by moving $\left\{V_{i}, 1 \leq i \leq k\right\}$ satisfy

$$
\begin{aligned}
& 2 \sum_{i} d^{\prime}\left(V_{i}, C_{i}^{\prime} \cap C_{i}\right)-2 \sum_{i} d^{\prime}\left(\tilde{V}_{i}, \tilde{C}_{i}\right) \\
\geq & \frac{99}{50}(\alpha-2) \sum_{i} d\left(V_{i}, C_{i}\right)-\frac{2 \alpha}{100} \sum_{i} d\left(W_{i}, C_{i}\right)-\frac{8 \alpha+16 \beta}{100} r \epsilon n .
\end{aligned}
$$

Proof. The intuition is similar to that of Claim 3.(a): $\sum_{i} d^{\prime}\left(V_{i}, C_{i}^{\prime} \cap C_{i}\right) \approx \alpha \sum_{i} d\left(V_{i}, C_{i}\right)$ and $\sum_{i} d^{\prime}\left(\tilde{V}_{i}, \tilde{C}_{i}\right) \approx \sum_{i} d\left(\tilde{V}_{i}, C_{i}\right)$. Since $\sum_{i} d\left(V_{i}, C_{i}\right) \geq \sum_{i} d\left(\tilde{V}_{i}, C_{i}\right)$, their difference is roughly $(\alpha-1) \sum_{i} d\left(V_{i}, C_{i}\right)$.

Formally, we have

$$
\begin{align*}
d^{\prime}\left(V_{i}, C_{i}^{\prime} \cap C_{i}\right) & =\alpha d\left(V_{i}, C_{i}^{\prime} \cap C_{i}\right)=\alpha d\left(V_{i}, C_{i}\right)-\alpha d\left(V_{i}, C_{i} \backslash C_{i}^{\prime}\right)  \tag{23}\\
d^{\prime}\left(\tilde{V}_{i}, \tilde{C}_{i}\right) & =d\left(\tilde{V}_{i}, \tilde{C}_{i}\right) \leq d\left(\tilde{V}_{i}, C_{i}\right)+d\left(\tilde{V}_{i}, D_{i}\right) \tag{24}
\end{align*}
$$

Then it suffices to bound the approximation error $d\left(V_{i}, C_{i} \backslash C_{i}^{\prime}\right)$ and $d\left(\tilde{V}_{i}, D_{i}\right)$. First,

$$
\begin{aligned}
d\left(V_{i}, C_{i} \backslash C_{i}^{\prime}\right) & \leq \frac{\left|C_{i} \backslash C_{i}^{\prime}\right|}{\left|C_{i}\right|} d\left(V_{i}, C_{i}\right)+\frac{\left|V_{i}\right|}{\left|C_{i}\right|} d\left(C_{i} \backslash C_{i}^{\prime}, C_{i}\right) \\
& \leq \frac{1}{100} d\left(V_{i}, C_{i}\right)+\frac{1}{100} d\left(C_{i} \backslash C_{i}^{\prime}, C_{i}\right)
\end{aligned}
$$

where the first inequality is by Fact 7 and the second is from the fact that $\left|C_{i} \backslash C_{i}^{\prime}\right| \leq$ $\epsilon n,\left|V_{i}\right| \leq \epsilon n$ and $\left|C_{i}\right| \geq 100 \epsilon n$. For the second term on the right hand side, we have $d\left(C_{i} \backslash C_{i}^{\prime}, C_{i}\right) \leq d\left(\tilde{W}_{i}, C_{i}\right)+d\left(\hat{B}_{i} \backslash C_{i}^{\prime}, C_{i}\right)$. Noting that $\sum_{i} d\left(\tilde{W}_{i}, C_{i}\right) \leq \sum_{i} d\left(W_{i}, C_{i}\right)$ and that the points in $\hat{B}_{i}$ have cost at most $2 r$, we have

$$
\sum_{i} d\left(V_{i}, C_{i} \backslash C_{i}^{\prime}\right) \leq \frac{1}{100} \sum_{i} d\left(V_{i}, C_{i}\right)+\frac{1}{100} \sum_{i} d\left(W_{i}, C_{i}\right)+\frac{4 r \epsilon n}{100}
$$

Similarly, for $d\left(\tilde{V}_{i}, D_{i}\right)$ we have

$$
\sum_{i} d\left(\tilde{V}_{i}, D_{i}\right) \leq \sum_{i}\left[\frac{\left|D_{i}\right|}{\left|C_{i}\right|} d\left(\tilde{V}_{i}, C_{i}\right)+\frac{\left|\tilde{V}_{i}\right|}{\left|C_{i}\right|} d\left(C_{i}, D_{i}\right)\right] \leq \frac{2}{100} \sum_{i} d\left(V_{i}, C_{i}\right)+\frac{8 \beta r \epsilon n}{100}
$$

The claim follows by summing (23) and (24) over $1 \leq i \leq k$ and plugging the last two inequalities.

Claim 3.(c). The costs saved and added by moving $\left\{W_{i}, 1 \leq i \leq k\right\}$ satisfy

$$
\begin{aligned}
& 2 \sum_{i} d^{\prime}\left(W_{i}, C_{i}^{\prime} \cap C_{i}\right)-2 \sum_{i} d^{\prime}\left(\tilde{W}_{i}, \tilde{W}_{i}+C_{i}^{\prime} \cap \tilde{C}_{i}\right) \\
\geq & \frac{98}{50}(\alpha-2) \sum_{i} d\left(W_{i}, C_{i}\right)-\frac{8 \alpha+8 \beta}{100} r \epsilon n .
\end{aligned}
$$

Proof. The intuition is similar to that of Claim 3.(a): $\sum_{i} d^{\prime}\left(W_{i}, C_{i}^{\prime} \cap C_{i}\right) \approx \alpha \sum_{i} d\left(W_{i}, C_{i}\right)$ and $\sum_{i} d^{\prime}\left(\tilde{W}_{i}, \tilde{W}_{i}+C_{i}^{\prime} \cap \tilde{C}_{i}\right) \approx \sum_{i} d\left(\tilde{W}_{i}, C_{i}\right)$. Since $\sum_{i} d\left(W_{i}, C_{i}\right) \geq \sum_{i} d\left(\tilde{W}_{i}, C_{i}\right)$, their difference is roughly $(\alpha-1) \sum_{i} d\left(W_{i}, C_{i}\right)$.

Formally, we have

$$
\begin{gather*}
d^{\prime}\left(W_{i}, C_{i}^{\prime} \cap C_{i}\right)=\alpha d\left(W_{i}, C_{i}^{\prime} \cap C_{i}\right)=\alpha d\left(W_{i}, C_{i}\right)-\alpha d\left(W_{i}, C_{i} \backslash C_{i}^{\prime}\right) .  \tag{25}\\
d^{\prime}\left(\tilde{W}_{i}, \tilde{W}_{i}+C_{i}^{\prime} \cap \tilde{C}_{i}\right)=d\left(\tilde{W}_{i}, C_{i}^{\prime} \cap D_{i}+C_{i} \cap \tilde{C}_{i}\right) \leq d\left(\tilde{W}_{i}, C_{i}^{\prime} \cap D_{i}\right)+d\left(\tilde{W}_{i}, C_{i}\right) . \tag{26}
\end{gather*}
$$

Then it suffices to bound the approximation error $d\left(W_{i}, C_{i} \backslash C_{i}^{\prime}\right)$ and $d\left(\tilde{W}_{i}, C_{i}^{\prime} \cap D_{i}\right)$.
First,

$$
\begin{aligned}
d\left(W_{i}, C_{i} \backslash C_{i}^{\prime}\right) & \leq \frac{\left|C_{i} \backslash C_{i}^{\prime}\right|}{\left|C_{i}\right|} d\left(W_{i}, C_{i}\right)+\frac{\left|W_{i}\right|}{\left|C_{i}\right|} d\left(C_{i} \backslash C_{i}^{\prime}, C_{i}\right) \\
& \leq \frac{1}{100} d\left(W_{i}, C_{i}\right)+\frac{1}{100} d\left(C_{i} \backslash C_{i}^{\prime}, C_{i}\right)
\end{aligned}
$$

where the first inequality is by Fact 7 and the second from the fact that $\left|C_{i} \backslash C_{i}^{\prime}\right| \leq$ $\epsilon n,\left|W_{i}\right| \leq \epsilon n$ and $\left|C_{i}\right| \geq 100 \epsilon n$. For the second term on the right hand side, we have $d\left(C_{i} \backslash C_{i}^{\prime}, C_{i}\right)=d\left(\tilde{W}_{i}, C_{i}\right)+d\left(\hat{B}_{i} \backslash C_{i}^{\prime}, C_{i}\right)$. Noting that $\sum_{i} d\left(\tilde{W}_{i}, C_{i}\right) \leq \sum_{i} d\left(W_{i}, C_{i}\right)$ and that points in $\hat{B}_{i}$ have cost at most $2 r$, we have

$$
\sum_{i} d\left(W_{i}, C_{i} \backslash C_{i}^{\prime}\right) \leq \frac{2}{100} \sum_{i} d\left(W_{i}, C_{i}\right)+\frac{4 r \epsilon n}{100}
$$

Similarly, for $d\left(\tilde{W}_{i}, C_{i}^{\prime} \cap D_{i}\right)$ we have

$$
\begin{aligned}
\sum_{i} d\left(\tilde{W}_{i}, C_{i}^{\prime} \cap D_{i}\right) & \leq \sum_{i}\left[\frac{\left|C_{i}^{\prime} \cap D_{i}\right|}{\left|C_{i}\right|} d\left(\tilde{W}_{i}, C_{i}\right)+\frac{\left|\tilde{W}_{i}\right|}{\left|C_{i}\right|} d\left(C_{i}^{\prime} \cap D_{i}, C_{i}\right)\right] \\
& \leq \frac{1}{100} \sum_{i} d\left(W_{i}, C_{i}\right)+\frac{4 \beta r \epsilon n}{100}
\end{aligned}
$$

The claim follows by summing (25) and (26) over $1 \leq i \leq k$ and plugging the last two inequalities.

## A.4.2 Properties of Good Points in Min-Sum

The first property is that the cost between two optimal clusters is roughly that between sufficiently large subsets of their good points (Lemma 10). To prove this, we begin with the following useful fact.

Fact 7. For any non-empty sets $A, B$ and $C$, we have $d_{\mathrm{a}}(A, B) \leq d_{\mathrm{a}}(A, C)+d_{\mathrm{a}}(C, B)$, and thus $d(A, B) \leq \frac{|B|}{|C|} d(A, C)+\frac{|A|}{|C|} d(C, B)$.

Then we have the following useful property for good points, which shows that good points are much closer to its own cluster than to good points in any other cluster. This property then leads to Lemma 10.

Lemma 23. For $G_{A} \subseteq G_{i}, G_{B} \subseteq G_{j}, j \neq i$, we have

$$
d_{\mathrm{a}}\left(G_{A}, C_{i}\right) \leq \gamma_{j i} d_{\mathrm{a}}\left(G_{A}, G_{B}\right), \text { where } \gamma_{j i}=\frac{\left|C_{j}\right|}{(\beta-1 / \beta)\left|C_{i}\right|}+\frac{1}{\beta^{2}-1}
$$

Consequently, if $\alpha>\frac{8 \max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$, we have $d_{\mathrm{a}}\left(G_{A}, C_{i}\right) \leq \frac{11}{50} d_{\mathrm{a}}\left(G_{A}, G_{B}\right)$.

Proof. For any $p \in G_{A}$, we have $\beta d\left(p, C_{i}\right)<d\left(p, C_{j}\right)$. It follows from Fact 7 that

$$
\begin{aligned}
& \beta d\left(G_{A}, C_{i}\right)<d\left(G_{A}, C_{j}\right) \leq \frac{\left|C_{j}\right|}{\left|G_{B}\right|} d\left(G_{A}, G_{B}\right)+\frac{\left|G_{A}\right|}{\left|G_{B}\right|} d\left(C_{j}, G_{B}\right) \\
& \beta d\left(G_{B}, C_{j}\right)<d\left(G_{B}, C_{i}\right) \leq \frac{\left|C_{i}\right|}{\left|G_{A}\right|} d\left(G_{B}, G_{A}\right)+\frac{\left|G_{B}\right|}{\left|G_{A}\right|} d\left(C_{i}, G_{A}\right)
\end{aligned}
$$

Plug the second inequality into the first inequality, then the lemma follows.
Lemma 10. Suppose $\alpha>\frac{8 \max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$ and $W_{i} \subseteq G_{i}, W_{j} \subseteq G_{j}$. When $\left|C_{i}\right| \geq 50\left|C_{i} \backslash W_{i}\right|$ and $\left|C_{j}\right| \geq 50\left|C_{j} \backslash W_{j}\right|$, we have $d\left(C_{i}, C_{j}\right) \leq \frac{3}{2} d\left(W_{i}, W_{j}\right)$.

Proof. By Fact 7 and Lemma 23, we have

$$
d_{\mathrm{a}}\left(C_{i}, C_{j}\right) \leq d_{\mathrm{a}}\left(C_{i}, W_{i}\right)+d_{\mathrm{a}}\left(W_{i}, W_{j}\right)+d_{\mathrm{a}}\left(W_{j}, C_{j}\right) \leq\left(\frac{11}{50}+1+\frac{11}{50}\right) d_{\mathrm{a}}\left(W_{i}, W_{j}\right)
$$

which leads to $d\left(C_{i}, C_{j}\right) \leq \frac{36}{25} \frac{\left|C_{i}\right|\left|C_{j}\right|}{\left|W_{i}\right|\left|W_{j}\right|} d\left(W_{i}, W_{j}\right) \leq \frac{3}{2} d\left(W_{i}, W_{j}\right)$.
Another property of good points is that the good points from two different clusters have cost much larger than those in a third cluster have (Lemma 11). To prove this, we need to prove Lemma 24, which bounds the cost of the optimal clustering $\mathcal{C}^{\prime}=\left\{C_{t}^{\prime}\right\}$ under the perturbed distance function. Recall the definitions of the perturbation and $\mathcal{C}^{\prime}$ in Section 2.5.1.2. The perturbation blows up all pairwise distances by a factor of $\alpha$ except the intra-cluster distances in $\tilde{\mathcal{C}}$, where $\tilde{\mathcal{C}}$ is the clustering obtained from the optimal clustering by splitting $C_{i}$ into $A$ and $C_{i} \backslash A$ and merging $C_{j}$ and $C_{l}$. Let
$\mathcal{C}^{\prime}=\left\{C_{i}^{\prime}\right\}$ denote the optimal clustering under the perturbed distance function $d^{\prime}$, where the clusters are indexed so that $C_{i}^{\prime}$ corresponds to $C_{i}$ and the distance between the two clustering is $\sum_{i}\left|C_{i} \backslash C_{i}^{\prime}\right|$.

To bound the cost of $\mathcal{C}^{\prime}=\left\{C_{t}^{\prime}\right\}$, we compare it to the cost of the optimal clustering $\mathcal{C}=\left\{C_{t}\right\}$ before perturbation. If $\mathcal{C}^{\prime}=\mathcal{C}$, then the cost is only increased by blowing up the distances between $A$ and $C_{i} \backslash A$ (Claim 10). However, the optimal clustering may change after the perturbation, so we need to consider how much cost is saved by the change (Claim 11).

Intuitively, the cost saved should be small. To see this, consider a point $p$ moved from $C_{s}$ to $C_{t}^{\prime}$. Then we need to pay $d^{\prime}\left(p, C_{t}^{\prime}\right)$ instead of $d\left(p, C_{s}\right)$. Note that $p$ is in $C_{s}$ but not $C_{t}$, so $d\left(p, C_{s}\right) \leq d\left(p, C_{t}\right)$. Also, $C_{t}^{\prime}$ and $C_{t}$ differ only on at most $\epsilon n$ points, then $d^{\prime}\left(p, C_{t}^{\prime}\right)$ is larger or comparable to $d\left(p, C_{t}\right)$ and thus $d\left(p, C_{s}\right)$.

There are two technical details in the above description. The first is to translate $d^{\prime}\left(p, C_{t}^{\prime}\right)$ to $d\left(p, C_{t}^{\prime}\right)$. We consider two cases (as in the proof of Claim 11). If $p$ is moved between $C_{j}$ and $C_{l}$, then $d^{\prime}\left(p, C_{t}^{\prime}\right)$ is roughly $d\left(p, C_{t}^{\prime}\right)$ since the distances between $C_{j}, C_{l}$ are not blown up. Otherwise, $d^{\prime}\left(p, C_{t}^{\prime}\right)$ is roughly $\alpha d\left(p, C_{t}^{\prime}\right)$. Another technical detail is to show that $d\left(p, C_{t}^{\prime}\right)$ roughly equals $d\left(p, C_{t}\right)$. Since $d\left(p, C_{t}^{\prime}\right) \geq d\left(p, C_{t}^{\prime} \cap C_{t}\right)$, it suffices to show that $d\left(p, C_{t}^{\prime} \cap C_{t}\right)$ is comparable to $d\left(p, C_{t}\right)$, where Fact 2 turns out to be useful.

Lemma 24. Suppose $\alpha>\frac{6 \max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$ and $\min _{i}\left|C_{i}\right| \geq 100 m_{B}$. We have

$$
\sum_{t=1}^{k} d^{\prime}\left(C_{t}^{\prime}, C_{t}^{\prime}\right)-\sum_{t=1}^{k} d\left(C_{t}, C_{t}\right) \geq 2(\alpha-1) d\left(A, G_{i} \backslash A\right)-\frac{4 \alpha+8}{100} d\left(C_{j}, C_{l}\right)
$$

Proof. Let $K_{t}=C_{t} \cap C_{t}^{\prime}, A_{t}=C_{t}^{\prime} \backslash C_{t}, M_{t}=C_{t} \backslash C_{t}^{\prime}$. See Figure 21 for an illustration.


Figure 21: Notations in Lemma 24.

We have

$$
\begin{aligned}
& \sum_{t=1}^{k} d^{\prime}\left(C_{t}^{\prime}, C_{t}^{\prime}\right)-\sum_{t=1}^{k} d\left(C_{t}, C_{t}\right) \\
\geq & \sum_{t=1}^{k}\left[d^{\prime}\left(K_{t}, K_{t}\right)+2 d^{\prime}\left(A_{t}, K_{t}\right)\right]-\sum_{t=1}^{k}\left[d\left(K_{t}, K_{t}\right)+2 d\left(M_{t}, C_{t}\right)\right] \\
= & {\left[\sum_{t=1}^{k} d^{\prime}\left(K_{t}, K_{t}\right)-\sum_{t=1}^{k} d\left(K_{t}, K_{t}\right)\right]+2\left[\sum_{t=1}^{k} d^{\prime}\left(A_{t}, K_{t}\right)-\sum_{t=1}^{k} d\left(M_{t}, C_{t}\right)\right] . }
\end{aligned}
$$

The first term on the right hand side corresponds to the cost increased by blowing up the distances within the clusters, the second term corresponds to the cost increased by moving points away. We will bound the two terms respectively in the following two claims, which then lead to the lemma.

Let $l(p)$ denote the index of the optimal cluster in $\mathcal{C}$ that $p$ falls in: if $p \in C_{t}$, then $l(p)=t$. Similarly, let $l^{\prime}(p)$ denote the optimal cluster in $\mathcal{C}^{\prime}$ that $p$ falls in after perturbation: if $p \in C_{t}^{\prime}$, then $l^{\prime}(p)=t$.

The first term is roughly the cost increased by blowing the distances between $A_{i}$ and $C_{i} \backslash A_{i}$, which is about $2(\alpha-1) d\left(A, C_{i} \backslash A\right)$. However, some points in $C_{i}$ may move away, so we need to exclude the cost of these points. More precisely, we only consider good points, and also exclude the cost of the good points moved away ( $G_{i} \cap M_{i}$ ).

## Claim 10.

$$
\begin{equation*}
\sum_{t=1}^{k} d^{\prime}\left(K_{t}, K_{t}\right)-\sum_{t=1}^{k} d\left(K_{t}, K_{t}\right) \geq 2(\alpha-1) d\left(A, G_{i} \backslash A\right)-\frac{2(\alpha-1)}{\beta} \sum_{p \in G_{i} \cap M_{i}} d\left(p, C_{l^{\prime}(p)}\right) \tag{27}
\end{equation*}
$$

Proof. By the definition of the perturbation, we have

$$
\begin{aligned}
& \sum_{t=1}^{k} d^{\prime}\left(K_{t}, K_{t}\right)-\sum_{t=1}^{k} d\left(K_{t}, K_{t}\right) \\
\geq & 2 d^{\prime}\left(A \cap K_{i},\left(G_{i} \backslash A\right) \cap K_{i}\right)-2 d\left(A \cap K_{i},\left(G_{i} \backslash A\right) \cap K_{i}\right) \\
\geq & 2(\alpha-1) d\left(A \cap K_{i},\left(G_{i} \backslash A\right) \cap K_{i}\right) \\
\geq & 2(\alpha-1)\left[d\left(A, G_{i} \backslash A\right)-d\left(A \cap M_{i}, G_{i} \backslash A\right)-d\left(\left(G_{i} \backslash A\right) \cap M_{i}, A\right)\right] \\
\geq & 2(\alpha-1)\left[d\left(A, G_{i} \backslash A\right)-d\left(G_{i} \cap M_{i}, C_{i}\right)\right] .
\end{aligned}
$$

The proof is completed by noting that for any point $p \in G_{i} \cap M_{i}, d\left(p, C_{i}\right) \leq$ $\frac{1}{\beta} d\left(p, C_{l^{\prime}(p)}\right)$.

The second term is roughly the cost increased by moving points away. Consider a point $p \in C_{1}$ that moves to $C_{2}^{\prime}$. The new cost is $d^{\prime}\left(p, C_{2}^{\prime}\right) \approx d^{\prime}\left(p, C_{2}\right)=\alpha d\left(p, C_{2}\right)$, and the old cost is $d\left(p, C_{1}\right) \leq d\left(p, C_{2}\right)$, so the cost increased is roughly $(\alpha-1) d\left(p, C_{2}\right)$. Note that $C_{2}^{\prime}$ only approximately equals $C_{2}$. Also, the above intuition does not hold for points that move between $C_{j}$ and $C_{l}$ since the distances between them are not blown up. These facts only decrease the bound slightly, as shown in the following claim.

Claim 11. Let $X=\left(\cup_{t} A_{t}\right) \backslash\left(A_{l} \cap C_{j}\right) \backslash\left(A_{j} \cap C_{l}\right)$.

$$
\begin{equation*}
\sum_{t=1}^{k} d^{\prime}\left(A_{t}, K_{t}\right)-\sum_{t=1}^{k} d\left(M_{t}, C_{t}\right) \geq\left(\frac{98 \alpha}{100}-1\right) \sum_{p \in X} d\left(p, C_{l^{\prime}(p)}\right)-\frac{2 \alpha+4}{100} d\left(C_{j}, C_{l}\right) \tag{28}
\end{equation*}
$$

Proof. We have

$$
\begin{equation*}
\sum_{t=1}^{k} d^{\prime}\left(A_{t}, K_{t}\right)-\sum_{t=1}^{k} d\left(M_{t}, C_{t}\right) \geq \sum_{t=1}^{k} \sum_{p \in A_{t}}\left[d^{\prime}\left(p, K_{t}\right)-d\left(p, C_{l(p)}\right)\right] \tag{29}
\end{equation*}
$$

Intuitively, $d^{\prime}\left(p, K_{l^{\prime}(p)}\right)$ should be larger or comparable to $d\left(p, C_{l(p)}\right)$. On one hand, $d\left(p, C_{l(p)}\right) \leq d\left(p, C_{l^{\prime}(p)}\right)$ since $p$ is assigned to $C_{l(p)}$ instead of $C_{l^{\prime}(p)}$ in the optimal clustering under $d$. On the other hand, we also know that $d\left(p, K_{l^{\prime}(p)}\right)$ is comparable to $d\left(p, C_{l^{\prime}(p)}\right)$ by Fact 2 .

Before using this intuition, we first need to translate $d^{\prime}\left(p, K_{l^{\prime}(p)}\right)$ to $d\left(p, K_{l^{\prime}(p)}\right)$. Since the distances between $C_{j}$ and $C_{l}$ is not blown up, we need to consider separately the case when $p$ is moved between $C_{j}$ and $C_{l}$. Equivalently, we divide $\cup_{t} A_{t}$ into two parts: $V=\left(A_{j} \cap C_{l}\right) \cup\left(A_{l} \cap C_{j}\right)$ and $X=\left(\cup_{t} A_{t}\right) \backslash\left(A_{l} \cap C_{j}\right) \backslash\left(A_{j} \cap C_{l}\right)$. Now we consider the two parts respectively.

Case 1: Suppose $p \in A_{j} \cap C_{l}$. By Fact 2, we have $d^{\prime}\left(p, K_{l^{\prime}(p)}\right)=d\left(p, K_{j}\right) \geq$ $\frac{\left|K_{j}\right|}{\left|C_{j}\right|} d\left(p, C_{j}\right)-\frac{1}{\left|C_{j}\right|} d\left(M_{j}, C_{j}\right)$. Since $d\left(p, C_{l(p)}\right)=d\left(p, C_{l}\right) \leq d\left(p, C_{j}\right)$, and $d\left(M_{j}, C_{j}\right) \leq$ $d\left(M_{j}, C_{l}\right) \leq d\left(C_{j}, C_{l}\right)$,

$$
\begin{aligned}
d^{\prime}\left(p, K_{l^{\prime}(p)}\right)-d\left(p, C_{l(p)}\right) & \geq-\frac{\left|M_{j}\right|}{\left|C_{j}\right|} d\left(p, C_{j}\right)-\frac{1}{\left|C_{j}\right|} d\left(C_{j}, C_{l}\right), \\
\sum_{p \in A_{j} \cap C_{l}}\left[d^{\prime}\left(p, K_{l^{\prime}(p)}\right)-d\left(p, C_{l(p)}\right)\right] & \geq-\left[\frac{\left|M_{j}\right|}{\left|C_{j}\right|}+\frac{\left|A_{j} \cap C_{l}\right|}{\left|C_{j}\right|}\right] d\left(C_{j}, C_{l}\right)
\end{aligned}
$$

Since $\left|M_{j}\right| \leq \epsilon n,\left|A_{j}\right| \leq \epsilon n$, this is bounded by $-\frac{2}{100} d\left(C_{j}, C_{l}\right)$. A similar argument holds for $A_{j} \cap C_{l}$. So

$$
\begin{equation*}
\sum_{p \in V}\left[d^{\prime}\left(p, K_{l^{\prime}(p)}\right)-d\left(p, C_{l(p)}\right)\right] \geq-\frac{4}{100} d\left(C_{j}, C_{l}\right) \tag{30}
\end{equation*}
$$

Case 2: For $p \in X$, we have by Fact 2

$$
d^{\prime}\left(p, K_{l^{\prime}(p)}\right)=\alpha d\left(p, K_{l^{\prime}(p)}\right) \geq \alpha\left[\frac{\left|K_{l^{\prime}(p)}\right|}{\left|C_{l^{\prime}(p)}\right|} d\left(p, C_{l^{\prime}(p)}\right)-\frac{1}{\left|C_{l^{\prime}(p)}\right|} d\left(M_{l^{\prime}(p)}, C_{l^{\prime}(p)}\right)\right] .
$$

Then for $X$, since $d\left(p, C_{l(p)}\right) \leq d\left(p, C_{l^{\prime}(p)}\right)$ and $\frac{\left|K_{l^{\prime}(p)}\right|}{\left|C_{l(p)}\right|} \geq \frac{99}{100}$, we have

$$
\begin{equation*}
\sum_{p \in X}\left[d^{\prime}\left(p, K_{l^{\prime}(p)}\right)-d\left(p, C_{l(p)}\right)\right] \geq\left(\frac{99 \alpha}{100}-1\right) \sum_{p \in X} d\left(p, C_{l^{\prime}(p)}\right)-\sum_{p \in X} \frac{\alpha}{\left|C_{l^{\prime}(p)}\right|} d\left(M_{l^{\prime}(p)}, C_{l^{\prime}(p)}\right) \tag{31}
\end{equation*}
$$

Since $X \subseteq \cup_{t} A_{t}$, and $\left|C_{t}\right| \geq 100\left|A_{t}\right|$, the second term on the right hand side is
bounded by

$$
\begin{align*}
\sum_{t} \frac{\alpha\left|A_{t}\right|}{\left|C_{t}\right|} d\left(M_{t}, C_{t}\right) & \leq \frac{\alpha}{100} \sum_{t} d\left(M_{t}, C_{t}\right)=\frac{\alpha}{100} \sum_{p \in \cup_{t} A_{t}} d\left(p, C_{l(p)}\right) \\
& =\frac{\alpha}{100}\left[\sum_{p \in V} d\left(p, C_{l(p)}\right)+\sum_{p \in V} d\left(p, C_{l(p)}\right)\right] \\
& \leq \frac{\alpha}{100}\left[\sum_{p \in V} d\left(p, C_{l^{\prime}(p)}\right)+2 d\left(C_{i}, C_{j}\right)\right] \tag{32}
\end{align*}
$$

The claim follows from the inequalities (30), (31), and (32).

The proof is completed by combining the two claims.
Lemma 11. Suppose $\alpha>\frac{8 \max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$ and $\epsilon<\frac{\min _{i}\left|C_{i}\right|}{600 n}$. For any three different optimal clusters $C_{i}, C_{j}$, and $C_{l}$, and any $A \subset G_{i}, \frac{18}{5} d\left(A, G_{i} \backslash A\right)<d\left(G_{j}, G_{l}\right)$. Consequently, $\frac{9}{5} d\left(G_{i}, G_{i}\right)<d\left(G_{j}, G_{l}\right)$.

Proof. The key idea is as follows. Let $\tilde{\mathcal{C}}$ denote the clustering obtained from the optimal clustering by splitting $C_{i}$ into $A$ and $C_{i} \backslash A$ and merging $C_{j}$ and $C_{l}$, i.e. $\tilde{\mathcal{C}}=$ $\left\{A, C_{i} \backslash A, C_{j} \cup C_{l}\right\} \cup\left\{C_{t}, t \neq i, j, l\right\}$. Suppose we construct a perturbation that favors the clustering $\tilde{\mathcal{C}}$ : blow up all pairwise distances by a factor of $\alpha$ except the intra-cluster distances in $\tilde{\mathcal{C}}$. Let $\mathcal{C}^{\prime}=\left\{C_{i}^{\prime}\right\}$ denote the optimal clustering under the perturbed distance function $d^{\prime}$, where the clusters are indexed so that $C_{i}^{\prime}$ corresponds to $C_{i}$ and the distance between the two clustering is $\sum_{i}\left|C_{i} \backslash C_{i}^{\prime}\right|$. By $(\alpha, \epsilon)$-perturbation resilience, we know that $\mathcal{C}^{\prime}$ is different from $\tilde{\mathcal{C}}$ and has no greater cost than $\tilde{\mathcal{C}}$. We then show that compared to the optimal cost under the original distances, the cost of $\tilde{\mathcal{C}}$ under perturbed distances $d^{\prime}$ is larger by at most $O\left(d\left(C_{j}, C_{l}\right)\right)=O\left(d\left(G_{j}, G_{l}\right)\right)$, while the cost of $\mathcal{C}^{\prime}$ under perturbed distances $d^{\prime}$ is larger by roughly $O(\alpha) d\left(A, G_{i} \backslash A\right)$. These then leads to the first statement.

More precisely, the cost of $\tilde{\mathcal{C}}$ under $d^{\prime}$ is larger than that of $\mathcal{C}$ under $d$ by at most $2 d\left(C_{j}, C_{l}\right)$. For $\mathcal{C}^{\prime}$, we have $\sum_{t=1}^{k} d^{\prime}\left(C_{t}^{\prime}, C_{t}^{\prime}\right)-\sum_{t=1}^{k} d\left(C_{t}, C_{t}\right) \geq 2(\alpha-1) d\left(A, G_{i} \backslash A\right)-$


Figure 22: Notations in Claim 4.
$\frac{4 \alpha+8}{100} d\left(C_{j}, C_{l}\right)$ by Lemma 24 . Since $\mathcal{C}^{\prime}$ has smaller cost than $\tilde{\mathcal{C}}$, we have

$$
2(\alpha-1) d\left(A, G_{i} \backslash A\right)-\frac{4 \alpha+8}{100} d\left(C_{j}, C_{l}\right) \leq 2 d\left(C_{j}, C_{l}\right)
$$

When $\alpha>\frac{8 \max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}$, we have $\frac{27}{5} d\left(A, G_{i} \backslash A\right) \leq d\left(C_{j}, C_{l}\right)$. By Lemma 10, $d\left(C_{j}, C_{l}\right) \leq$ $\frac{3}{2} d\left(G_{j}, G_{l}\right)$, which then leads to the first part of the lemma.

The second part of the lemma follows from the fact that $\sum_{A \subseteq G_{i}} d\left(A, G_{i} \backslash A\right)=$ $\frac{2^{\left|G_{i}\right|}}{2} d\left(G_{i}, G_{i}\right)$.

## A.4.3 Properties of Potential Good Points in Min-Sum

A key property of the potential good points is that the cost between any point $p$ and the potential good points in a sufficiently large set $A$ is roughly the cost between $p$ and any sufficiently large subset $H$ of $A$ (Lemma 12). In its proof in Section 2.5.1.2, we need the following claim.

Claim 4. Suppose $H \subseteq A$ such that $|A \backslash H| \leq m_{B}$. Let $F=F(A), P=P(A), \bar{H}=$ $A \backslash H$. Let $W=H \cap P, V=\bar{H} \cap F, X=F \cap H, Y=\bar{H} \cap P$. See Figure 22 for illustration. If $|A| \geq 20 m_{B}$, then $d(W, X) \geq d(Y, W+Y)$.

Proof. The claim is true if $Y=\emptyset$. Otherwise, by definition of the potential bad points $F=F(A)$, we have $d_{\mathrm{a}}(X, A) \geq d_{\mathrm{a}}(Y, A)$. By definition,

$$
\begin{align*}
|A| d_{\mathrm{a}}(X, A) & =|W| d_{\mathrm{a}}(X, W)+|V| d_{\mathrm{a}}(X, V)+|Y| d_{\mathrm{a}}(X, Y)+|X| d_{\mathrm{a}}(X, X)  \tag{33}\\
|A| d_{\mathrm{a}}(Y, A) & =|W+Y| d_{\mathrm{a}}(Y, W+Y)+|V| d_{\mathrm{a}}(Y, V)+|X| d_{\mathrm{a}}(Y, X) \tag{34}
\end{align*}
$$

To compare $d(X, W)$ and $d(Y, W+Y)$, we need to bound the other terms in (33) and (34). By Fact 7,

$$
\begin{gathered}
d_{\mathrm{a}}(X, V) \leq d_{\mathrm{a}}(X, W)+d_{\mathrm{a}}(W, Y)+d_{\mathrm{a}}(Y, V), \\
d_{\mathrm{a}}(X, Y) \leq d_{\mathrm{a}}(X, W)+d_{\mathrm{a}}(Y, W), \quad d_{\mathrm{a}}(X, X) \leq 2 d_{\mathrm{a}}(X, W) .
\end{gathered}
$$

Now we plug these into (33), and then plug (33) and (34) into $d_{\mathrm{a}}(X, A) \geq d_{\mathrm{a}}(Y, A)$. Since $d(W, Y) \leq d(Y, W+Y)$ and $d_{\mathrm{a}}(Y, X) \geq 0$, we have

$$
[|W|-|Y|-|V|] d(Y, W+Y) \leq[|W|+2|X|+|Y|+|V|] \frac{|Y|}{|X|} d(X, W)
$$

Since $|X+V|=|F|=2 m_{B}$ and $|Y+V|=|A \backslash H| \leq m_{B}$, we have $\frac{|Y|}{|X|} \leq 1 / 2$. Then the lemma follows from the fact that $|A| \geq 20 m_{B},|F|=2 m_{B}$ and $|A \backslash H| \leq m_{B}$.

## A. 5 Approximation Algorithm for ( $\alpha, \epsilon$ )-Perturbation Resilient Min-Sum

## A.5.1 Proofs for Tree Construction for Min-Sum

To show that the robust average linkage algorithm keeps the laminarity of the clustering $\mathcal{L}$ with respect to the optimal clustering (after removing the bad points), we need to show the following claim.

Claim 6. (a) Suppose $A \in \mathcal{L}, A \cap G \subsetneq G_{i}$, and $D \in \mathcal{L}, D \cap G \subsetneq G_{j}(j \neq i)$. Then there exists $A^{\prime} \neq A$ in $\mathcal{L}$ such that $A^{\prime} \cap G \subsetneq G_{i}$ and $d_{r a}\left(A, A^{\prime}\right)<d_{r a}(A, D)$.
(b) Suppose $A \in \mathcal{L}, A \cap G \subsetneq G_{i}$, and $D \in \mathcal{L}, D \cap G$ is the union of good points in several optimal clusters. Then there exists $A^{\prime} \neq A$ in $\mathcal{L}$ such that $A^{\prime} \cap G \subsetneq G_{i}$ and $d_{r a}\left(A, A^{\prime}\right)<d_{r a}(A, D)$.

Proof. (a) The claim follows from the following three statements:

1. $d_{\mathrm{a}}\left(A \cap G, A^{\prime} \cap G\right)<\frac{1}{2} d_{\mathrm{a}}(A \cap G, D \cap G)$;
2. $d_{r a}\left(A, A^{\prime}\right) \leq \frac{7}{5} d_{\mathrm{a}}\left(A \cap G, A^{\prime} \cap G\right)$;
3. $\frac{9}{10} d_{\mathrm{a}}(A \cap G, D \cap G) \leq d_{r a}(A, D)$.
4. For simplicity, let $G_{A}=A \cap G, G_{D}=D \cap G$. From Lemma 23, we have

$$
d_{\mathrm{a}}\left(G_{A}, C_{i}\right) \leq \gamma_{j i} d_{\mathrm{a}}\left(G_{A}, G_{D}\right) \text {, where } \gamma_{j i}=\frac{\left|C_{j}\right|}{(\beta-1 / \beta)\left|C_{i}\right|}+\frac{1}{\beta^{2}-1} .
$$

Since $d\left(G_{A}, G_{i} \backslash G_{A}\right) \leq d\left(G_{A}, C_{i}\right)$, we have
$d_{\mathrm{a}}\left(G_{A}, G_{i} \backslash G_{A}\right) \leq \frac{\left|C_{i}\right|}{\left|G_{i} \backslash G_{A}\right|} d_{\mathrm{a}}\left(G_{A}, C_{i}\right) \leq \gamma_{j i} \frac{\left|C_{i}\right|}{\left|G_{i} \backslash G_{A}\right|} d_{\mathrm{a}}\left(G_{A}, G_{D}\right) \leq \frac{1}{2} d_{\mathrm{a}}\left(G_{A}, G_{D}\right)$
where the last step follows from $\alpha \geq 6 \frac{\max _{i}\left|C_{i}\right|}{\min _{i}\left|C_{i}\right|}+2,\left|G_{i} \backslash A\right|$ is at least $\frac{1}{2} \min _{i}\left|C_{i}\right|-m_{B}$.
2. By Lemma 12 and the fact that $|A| \geq \frac{1}{2} \min _{i}\left|C_{i}\right|,\left|A^{\prime}\right| \geq \frac{1}{2} \min _{i}\left|C_{i}\right|$ and $\min _{i}\left|C_{i}\right|>$ $100 m_{B}$, we have

$$
d\left(P(A), P\left(A^{\prime}\right)\right) \leq \frac{10}{9} d\left(P(A), A^{\prime} \cap G\right) \leq \frac{100}{81} d\left(A \cap G, A^{\prime} \cap G\right) .
$$

Then the claim follows from the fact that $|P(A)| \geq \frac{48}{50}|A|,\left|P\left(A^{\prime}\right)\right| \geq \frac{48}{50}\left|A^{\prime}\right|$.
3. For simplicity, let $G_{A}=A \cap G, G_{D}=D \cap G$. Divide $G_{A}$ into two parts: $W_{A}=$ $G_{A} \cap P(A)$ and $X_{A}=G_{A} \cap F(A)$. Define $W_{D}$ and $X_{D}$ similarly. See Figure 23 for illustration.


Figure 23: Notations in Claim 6.

To show $d\left(G_{A}, G_{D}\right) \leq O(1) d(P(A), P(D))$, it suffices to show $d\left(G_{A}, G_{D}\right) \leq O(1) d\left(W_{A}, W_{D}\right)$. Since $d\left(G_{A}, G_{D}\right) \leq d\left(W_{A}, W_{D}\right)+d\left(G_{A}, X_{D}\right)+d\left(G_{D}, X_{A}\right)$, we only need to bound $d\left(G_{A}, X_{D}\right)$ and $d\left(G_{D}, X_{A}\right)$. By Fact 7 and Lemma 23 we have

$$
d_{\mathrm{a}}\left(G_{A}, X_{D}\right) \leq d_{\mathrm{a}}\left(G_{A}, G_{D}\right)+d_{\mathrm{a}}\left(G_{D}, X_{D}\right) \leq\left(1+\frac{11}{50}\right) d_{\mathrm{a}}\left(G_{A}, G_{D}\right) .
$$

Since $\left|X_{D}\right| \leq 2 m_{B}$ and $|D| \geq \frac{1}{2} \min _{i}\left|C_{i}\right| \geq 50 m_{B}$, we have $d\left(G_{A}, X_{D}\right) \leq \frac{61}{50} \frac{\left|X_{D}\right|}{\left|G_{D}\right|} d\left(G_{A}, G_{D}\right) \leq$ $\frac{1}{20} d\left(G_{A}, G_{D}\right)$. Similarly, $d\left(G_{D}, X_{A}\right) \leq \frac{1}{20} d\left(G_{A}, G_{D}\right)$. Therefore, $d\left(G_{A}, G_{D}\right) \leq d\left(W_{A}, W_{D}\right)++d\left(G_{A}, X_{D}\right)+d\left(G_{D}, X_{A}\right) \leq d\left(W_{A}, W_{D}\right)+\frac{1}{10} d\left(G_{A}, G_{D}\right)$
which leads to $\frac{9}{10} d\left(G_{A}, G_{D}\right) \leq d\left(W_{A}, W_{D}\right) \leq d(P(A), P(D))$. Then the claim follows from the fact that $\left|G_{A}\right| \geq|P(A)|,\left|G_{D}\right| \geq|P(D)|$.
(b) The proof idea is similar to that for (a). The only difference is the proof for $d_{\mathrm{a}}\left(A \cap G, A^{\prime} \cap G\right)<\frac{1}{2} d_{\mathrm{a}}(A \cap G, D \cap G)$. Since $D \cap G=\cup_{j \in I_{D}} G_{j}$, it suffices to show that $d_{\mathrm{a}}\left(A \cap G, A^{\prime} \cap G\right)<\frac{1}{2} d_{\mathrm{a}}\left(A \cap G, G_{j}\right)$ for any $j \in I_{D}$, which can be proved by the same argument in Claim 6.

## A.5.2 Proofs for Getting the Pruning Close to the Optimal Clustering

The same argument as that for Claim 8 leads to a corollary for the general case when multiple clusters are merged.

Corollary 1. Let $I \subseteq[k]$. Suppose for any $t \in I,\left|C_{t}\right| \geq 100 m_{B}$, and $C_{t}^{\prime}$ contains all good points in $C_{t}$ but no good points in other optimal clusters. Then

$$
d_{\mathrm{rs}}\left(\cup_{t \in I} C_{t}^{\prime}\right)-\sum_{t \in I} d\left(G_{t}, G_{t}\right) \geq\left(\frac{4}{3}-\frac{4}{\beta}\right) \sum_{s \neq t \in I} d\left(G_{t}, G_{s}\right) .
$$

Lemma 14. Suppose the pruning $\mathcal{C}^{\prime}=\left\{C_{1}^{\prime}, \ldots, C_{k}^{\prime}\right\}$ in the tree $\mathcal{T}$ assigns all good points correctly. Then $\mathcal{C}^{\prime}$ is the minimum robust min-sum cost pruning in the tree.

Proof. First, by Lemma 11, good points from different clusters are far apart while good points in the same cluster are close. Second, by Claim 7 and Corollary 1, the cost of good points can be approximated by the cost of the potential good points (the robust min-sum cost). We now use the above lemmas to show that $\mathcal{C}^{\prime}$ has minimum robust min-sum cost, so that we can use dynamic programming on the tree to get the pruning.

Suppose a pruning $\mathcal{P}$ is obtained by splitting $h$ clusters in $\mathcal{C}^{\prime}$ and at the same time joining some other clusters into $g$ unions. Specifically, for $1 \leq i \leq h$, split $C_{i}^{\prime}$ into $m_{i} \geq 2$ clusters $P_{i, 1}, \ldots, P_{i, m_{i}}$; after that, merge $C_{h+1}^{\prime}, \ldots, C_{h+l_{g}}^{\prime}$ into $g$ unions, i.e. for $1 \leq j \leq g, l_{0}=0$, merge $l_{j}-l_{j-1} \geq 2$ clusters $C_{h+l_{j-1}+1}^{\prime}, \ldots, C_{h+l_{j}}^{\prime}$ into a union $U_{j} ;$ the other clusters in $\mathcal{C}^{\prime}$ remain the same in $\mathcal{P}$. Since the number of clusters is still $k$, we have $\sum_{i} m_{i}-h=l_{g}-g$.

By Claim 7, the cost saved by splitting the $h$ clusters is

$$
\begin{equation*}
\sum_{1 \leq i \leq h} d_{\mathrm{rs}}\left(C_{i}^{\prime}\right)-\sum_{1 \leq i \leq h} \sum_{1 \leq p \leq m_{i}} d_{\mathrm{rs}}\left(P_{i, p}\right) \leq \sum_{1 \leq i \leq h} d_{\mathrm{rs}}\left(C_{i}^{\prime}\right) \leq \sum_{1 \leq i \leq h} d\left(G_{i}, G_{i}\right) . \tag{35}
\end{equation*}
$$

The cost increased by joining clusters is

$$
\begin{align*}
& \sum_{1 \leq j \leq g}\left[d_{\mathrm{rs}}\left(U_{j}\right)-\sum_{h+l_{j-1}<t \leq h+l_{j}} d_{\mathrm{rs}}\left(C_{t}^{\prime}\right)\right] \\
\geq & \sum_{1 \leq j \leq g}\left[d_{\mathrm{rs}}\left(U_{j}\right)-\sum_{h+l_{j-1}<t \leq h+l_{j}} d\left(G_{t}, G_{t}\right)\right] \\
\geq & \sum_{1 \leq j \leq g}\left[\sum_{h+l_{j-1}<t \neq s \leq h+l_{j}}\left(\frac{4}{3}-\frac{4}{\beta}\right) d\left(G_{t}, G_{s}\right)\right] \tag{36}
\end{align*}
$$

where the first inequality follows from Claim 7, and the second inequality follows from Corollary 1. To prove $\mathcal{C}^{\prime}$ is the minimum cost pruning, we need to show that the saved cost (35) is less than the increased cost (36). Since by Lemma 11, each term in (36) is larger than any term in (35), it is sufficient to show that the number of the terms in (36) is no less than the number of the terms in (35), that is $\sum_{1 \leq j \leq g}\binom{l_{j}-l_{j-1}}{2} \geq h$. We have $\sum_{j}\binom{l_{j}-l_{j-1}}{2}=\frac{1}{2} \sum_{j}\left(l_{j}-l_{j-1}\right)\left(l_{j}-l_{j-1}-1\right) \geq \sum_{j}\left(l_{j}-l_{j-1}-1\right)=l_{g}-g$, where the inequality is from $l_{j}-l_{j-1} \geq 2$. Since $m_{i} \geq 2, l_{g}-g=\sum_{i-1}^{h} m_{i}-h \geq h$, which completes the proof.

## A.5.3 Getting a Good Approximation

To get a good approximation from a clustering $\mathcal{C}^{\prime}$ that assigns all good points correctly, Algorithm 7 reassigns each point $p$ to the index $i$ that minimizes the cost between


Figure 24: Notations in Lemma 15 and Claim 9.
$p$ and the potential good points in $C_{i}^{\prime}$. The following lemma shows that after this reassignment, all good points are still assigned correctly.

Lemma 15. For any $p \in G_{i}$, any $j \neq i, d\left(p, P\left(C_{j}^{\prime}\right)\right)>d\left(p, P\left(C_{i}^{\prime}\right)\right)$.
Proof. Let $W_{i}=G_{i} \cap P\left(C_{i}^{\prime}\right)$ denote the good points that are also potential good points, and let $Z_{i}=C_{i} \backslash W_{i}$ denote all other points in $C_{i}$. See Figure 24 for an illustration. By Lemma $12, d\left(p, P\left(C_{i}^{\prime}\right)\right) \approx d\left(p, C_{i}\right)$. By definition of good points, $\beta d\left(p, C_{i}\right) \leq d\left(p, C_{j}\right)$. So it suffices to show that $d\left(p, P\left(C_{j}^{\prime}\right)\right)$ is not so small compared to $d\left(p, C_{j}\right)$. Since $W_{j} \subseteq P\left(C_{j}^{\prime}\right)$, it suffices to prove that $d\left(p, W_{j}\right)$ is large compared to $d\left(p, Z_{j}\right)$.

First, by triangle inequality, $d\left(p, Z_{j}\right) \leq \frac{\left|Z_{j}\right|}{\left|W_{j}\right|} d\left(p, W_{j}\right)+\frac{1}{\left|W_{j}\right|} d\left(Z_{j}, W_{j}\right)$. Also, $d\left(Z_{j}, W_{j}\right) \leq$ $d\left(C_{j}, W_{j}\right) \leq \frac{1}{\beta} d\left(C_{i}, W_{j}\right)$ by definition of good points. Furthermore, $d\left(C_{i}, W_{j}\right) \leq$ $\left|W_{j}\right| d\left(p, C_{i}\right)+\left|C_{i}\right| d\left(p, W_{j}\right)$. So

$$
\begin{aligned}
d\left(p, Z_{j}\right) & \leq\left(\frac{\left|Z_{j}\right|}{\left|W_{j}\right|}+\frac{\left|C_{i}\right|}{\beta\left|W_{j}\right|}\right) d\left(p, W_{j}\right)+\frac{1}{\beta} d\left(p, C_{i}\right) \\
& \leq\left(\frac{\left|Z_{j}\right|}{\left|W_{j}\right|}+\frac{\left|C_{i}\right|}{\beta\left|W_{j}\right|}\right) d\left(p, W_{j}\right)+\frac{1}{\beta^{2}}\left[d\left(p, Z_{j}\right)+d\left(p, W_{j}\right)\right]
\end{aligned}
$$

Therefore, we have $d\left(p, Z_{j}\right) \leq \frac{1}{3} d\left(p, W_{j}\right)$, since $\left|Z_{j}\right| \leq 4 m_{B},\left|W_{j}\right| \geq \frac{95}{100}\left|C_{j}\right| \geq 95 m_{B}$. This leads to $d\left(p, W_{j}\right) \geq \frac{3}{4} d\left(p, C_{j}\right)$. Then the lemma follows from

$$
d\left(p, P\left(C_{j}^{\prime}\right)\right) \geq d\left(p, W_{j}\right) \geq \frac{3}{4} d\left(p, C_{j}\right) \geq \frac{3}{4} \beta d\left(p, C_{i}\right) \geq \frac{3}{4} \beta d\left(p, G_{i}\right) \geq \frac{30}{44} \beta d\left(p, P\left(C_{i}^{\prime}\right)\right)
$$

where the last step follows from Lemma 12.

Recall that $A_{i}$ are all the bad points reassigned to index $i$ by Algorithm 7. To bound the cost of the clustering after reassignment, we need to bound $d\left(A_{i}, G_{i}\right)$ as follows.

Claim 9. $\sum_{i} d\left(A_{i}, G_{i}\right) \leq \frac{r^{2}}{(r-5)^{2}} \sum_{i} d\left(C_{i}, C_{i}\right)-\frac{r^{2}-1}{(r-5)^{2}} \sum_{i} d\left(G_{i}, G_{i}\right)$, where $r=\frac{\min _{i}\left|C_{i}\right|}{m_{B}}$.
Proof. Let $W_{i}=P\left(C_{i}^{\prime}\right) \cap G_{i}$. See Figure 24 for an illustration. By Fact 6 , we have

$$
\begin{equation*}
d\left(A_{i}, G_{i}\right) \leq \frac{\left|G_{i}\right|}{\left|W_{i}\right|} d\left(A_{i}, W_{i}\right)+\frac{\left|A_{i}\right|}{\left|W_{i}\right|} d\left(G_{i}, W_{i}\right) \leq \frac{\left|G_{i}\right|}{\left|W_{i}\right|} d\left(A_{i}, P\left(C_{i}^{\prime}\right)\right)+\frac{\left|A_{i}\right|}{\left|W_{i}\right|} d\left(G_{i}, G_{i}\right) \tag{37}
\end{equation*}
$$

So it suffices to bound $d\left(A_{i}, P\left(C_{i}^{\prime}\right)\right)$. Fix $p \in A_{i}$, and suppose $p \in C_{j}$. We have

$$
\begin{aligned}
d\left(p, P\left(C_{i}^{\prime}\right)\right) & \leq d\left(p, P\left(C_{j}^{\prime}\right)\right) \leq \frac{\left|W_{j}\right|+\left|Y_{j}\right|}{\left|W_{j}\right|-\left|X_{j}\right|} d\left(p, G_{j}\right) \\
& \leq \frac{\left|C_{j}\right|}{\left|C_{j}\right|-2\left|X_{j}\right|-\left|Y_{j}\right|} d\left(p, G_{j}\right)=\frac{r}{r-5} d\left(p, G_{j}\right)
\end{aligned}
$$

where the second step follows from Lemma 12 and the last from $\left|X_{j}\right| \leq 2 m_{B}$ and $\left|Y_{j}\right| \leq m_{B}$. Then

$$
\begin{align*}
\sum_{i=1}^{k} d\left(A_{i}, P\left(C_{i}^{\prime}\right)\right) & \leq \frac{r}{r-5} \sum_{j} \sum_{p \in\left(\cup_{i} A_{i}\right) \cap C_{j}} d\left(p, G_{j}\right)=\frac{r}{r-5} \sum_{j=1}^{k} d\left(B_{j}, G_{j}\right) \\
& \leq \frac{r}{r-5} \sum_{j=1}^{k}\left[d\left(C_{j}, C_{j}\right)-d\left(G_{j}, G_{j}\right)\right] \tag{38}
\end{align*}
$$

The claim follows from the inequalities (37) and (38) and $\left|X_{i}\right| \leq 2 m_{B},\left|A_{i}\right| \leq m_{B}$.

## APPENDIX B

## DISTRIBUTED CLUSTERING

## B. 1 Proofs for Section 3.2.1

The proof of Lemma 16 follows from the analysis in [45], although not explicitly stated there. We begin with the following theorem for uniform sampling on a function space. The theorem is from [45] but rephrased for convenience.

Theorem 18 (Theorem 6.9 in [45]). Let $F$ be a set of functions from $P$ to $\mathbf{R}_{\geq 0}$, and let $\epsilon \in(0,1)$. Let $S$ be a sample of

$$
|S|=\frac{c}{\epsilon^{2}}\left(\operatorname{dim}(F, P)+\log \frac{1}{\delta}\right)
$$

i.i.d items from $P$, where $c$ is a sufficiently large constant. Then, with probability at least $1-\delta$, for any $f \in F$ and any $r \geq 0$,

$$
\left|\frac{\sum_{p \in P, f(p) \leq r} f(p)}{|P|}-\frac{\sum_{q \in S, f(q) \leq r} f(q)}{|S|}\right| \leq \epsilon r .
$$

Proof of Lemma 16. Without loss of generality, assume $m_{p} \in \mathbf{N}^{+}$. Define $G$ as follows: for each $p \in P$, include $m_{p}$ copies $\left\{p_{i}\right\}_{i=1}^{m_{p}}$ of $p$ in $G$ and define $f\left(p_{i}\right)=f(p) / m_{p}$. Then $S$ is equivalent to a sample draw i.i.d. and uniformly at random from $G$. We now apply Theorem 18 on $G$ and $r=\max _{f \in F, p^{\prime} \in G} f\left(p^{\prime}\right)$. By Theorem 18, we know that for any $f \in F$,

$$
\begin{equation*}
\left|\frac{\sum_{p^{\prime} \in G} f\left(p^{\prime}\right)}{|G|}-\frac{\sum_{q^{\prime} \in S} f\left(q^{\prime}\right)}{|S|}\right| \leq \epsilon \max _{p^{\prime} \in G} f\left(p^{\prime}\right) . \tag{39}
\end{equation*}
$$

The lemma then follows from multiplying both sides of (39) by $|G|=\sum_{p \in P} m_{p}$. Also note that the dimension $\operatorname{dim}(F, G)$ is the same as that of $\operatorname{dim}(F, P)$ as pointed out by [45].

Lemma 25. If $d\left(p, b_{p}\right)^{2} / \epsilon \leq\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right|$, then

$$
\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right| \leq 8 \epsilon \min \left\{d(p, \mathbf{x})^{2}, d\left(b_{p}, \mathbf{x}\right)^{2}\right\} .
$$

Proof. We first have by triangle inequality

$$
\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right| \leq d\left(p, b_{p}\right)\left[d(p, \mathbf{x})+d\left(b_{p}, \mathbf{x}\right)\right] .
$$

Then by $d\left(p, b_{p}\right)^{2} / \epsilon \leq\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right|$,

$$
d\left(p, b_{p}\right) \leq \epsilon\left[d(p, \mathbf{x})+d\left(b_{p}, \mathbf{x}\right)\right] .
$$

Therefore, we have

$$
\begin{aligned}
\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right| & \leq d\left(p, b_{p}\right)\left[d(p, \mathbf{x})+d\left(b_{p}, \mathbf{x}\right)\right] \leq \epsilon\left[d(p, \mathbf{x})+d\left(b_{p}, \mathbf{x}\right)\right]^{2} \\
& \leq 2 \epsilon\left[d(p, \mathbf{x})^{2}+d\left(b_{p}, \mathbf{x}\right)^{2}\right] \leq 2 \epsilon\left[d(p, \mathbf{x})^{2}+\left(d(p, \mathbf{x})+d\left(p, b_{p}\right)\right)^{2}\right] \\
& \leq 2 \epsilon\left[d(p, \mathbf{x})^{2}+2 d(p, \mathbf{x})^{2}+2 d\left(p, b_{p}\right)^{2}\right] \leq 6 \epsilon d(p, \mathbf{x})^{2}+4 \epsilon d\left(p, b_{p}\right)^{2} \\
& \leq 6 \epsilon d(p, \mathbf{x})^{2}+4 \epsilon^{2}\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right|
\end{aligned}
$$

for sufficiently small $\epsilon$. Then

$$
\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right| \leq \frac{6 \epsilon}{1-4 \epsilon^{2}} d(p, \mathbf{x})^{2} \leq 8 \epsilon d(p, \mathbf{x})^{2}
$$

Similarly, $\left|d(p, \mathbf{x})^{2}-d\left(b_{p}, \mathbf{x}\right)^{2}\right| \leq 8 \epsilon d\left(b_{p}, \mathbf{x}\right)^{2}$. The lemma follows from the last two inequalities.

Lemma 26 (Corollary 15.4 in [45]). Let $0<\delta<1 / 2$, and $t \geq c|B| \log \frac{|B|}{\delta}$ for a sufficiently large $c$. Then with probability at least $1-\delta, \forall b \in B_{i}, \sum_{q \in P_{b} \cap S} w_{q} \leq 2\left|P_{b}\right|$.

## B. 2 Complete Experimental Results

Here we present the results of all the data sets over different network topologies and data partition methods.

Figure 25 shows the results of all the data sets on random graphs. The first column of Figure 25 shows that our algorithm and COMBINE perform nearly the same in
the uniform data partition. This is not surprising since our algorithm reduces to the COMBINE algorithm when each local site has the same cost and the two algorithms use the same amount of communication. In this case, since in our algorithm the sizes of the local samples are proportional to the costs of the local solutions, it samples the same number of points from each local data set. This is equivalent to the COMBINE algorithm with the same amount of communication. In the similarity-based partition, similar results are observed as this partition method also leads to balanced local costs. However, in the weighted partition where local sites have significantly different contributions to the total cost, our algorithm outperforms COMBINE. It improves the $k$-means cost by $2 \%-5 \%$, and thus saves $10 \%-30 \%$ communication cost to achieve the same approximation ratio.

Figure 26 shows the results of all the data sets on grid and preferential graphs. Similar to the results on random graphs, our algorithm performs nearly the same as COMBINE in the similarity-based partition and outperforms COMBINE in the weighted partition and degree-based partition. Furthermore, Figure 25 and 26 also show that the performance of our algorithm merely changes over different network topologies and partition methods.

Figure 27 shows the results of all the data sets on the spanning trees of the random graphs and Figure 28 shows those on the spanning trees of the grid and preferential graphs. Compared to the algorithm of Zhang et al., our algorithm consistently shows much better performance on all the data sets in different settings. It improves the $k$-means cost by $10 \%-30 \%$, and thus can achieve even better approximation ratio with only $10 \%$ communication cost. This is because the algorithm of Zhang et al. constructs coresets from component coresets and needs larger coresets to prevent the accumulation of errors. Figure 27 also shows that although their costs decrease with the increase of the communication, the decrease is slower on larger graphs (e.g., as in the experiments for YearPredictionMSD). This is due to the fact that the spanning
tree of a larger graph has larger height, leading to more accumulation of errors. In this case, more communication is needed to prevent the accumulation.


Figure 25: $k$-means cost on random graphs. Columns: random graph with uniform partition, random graph with similarity-based partition, and random graph with weighted partition. Rows: Spam, Pendigits, Letter, synthetic, ColorHistogram, and YearPredictionMSD.


Figure 26: $k$-means cost on grid and preferential graphs. Columns: grid graph with similarity-based partition, grid graph with weighted partition, and preferential graph with degree-based partition. Rows: Spam, Pendigits, Letter, synthetic, ColorHistogram, and YearPredictionMSD.


Figure 27: $k$-means cost on the spanning trees of the random graphs. Columns: random graph with uniform partition, random graph with similarity-based partition, and random graph with weighted partition. Rows: Spam, Pendigits, Letter, synthetic, ColorHistogram, and YearPredictionMSD.


Figure 28: $k$-means cost on the spanning trees of the grid and preferential graphs. Columns: grid graph with similarity-based partition, grid graph with weighted partition, and preferential graph with degree-based partition. Rows: Spam, Pendigits, Letter, synthetic, ColorHistogram, and YearPredictionMSD.

## B. 3 Distributed $k$-Means Clustering of High Dimensional Data

## B.3.1 Proof of Lemma 18

Lemma 18. Let $A \in \mathbf{R}^{n \times d}$ be an $n \times d$ matrix with singular value decomposition $A=U D V^{\top}$. Let $\epsilon \in(0,1]$ and $r, t \in \mathbf{N}_{+}$with $d-1 \geq t \geq r+\lceil r / \epsilon\rceil-1$, and let $\tilde{A}=A V^{(t)}\left(V^{(t)}\right)^{\top}$. Then for any matrix $X$ with $d$ rows and $\|X\|_{F}^{2} \leq r$, we have

$$
\|(A-\tilde{A}) X\|_{F}^{2}=\|A X\|_{F}^{2}-\|\tilde{A} X\|_{F}^{2} \leq \epsilon \sum_{i=r+1}^{d} \sigma_{i}^{2}(A)
$$

Proof. The proof follows the idea in the proof of Lemma 6.1 in [46].
For convenience, let $\overline{D^{(t)}}$ denote the diagonal matrix that contains the first $t$ diagonal entries in $D$ and is 0 otherwise. Then $\tilde{A}=U \overline{D^{(t)}} V^{\top}$ We first have

$$
\begin{aligned}
\|A X\|_{F}^{2}-\|\tilde{A} X\|_{F}^{2} & =\left\|U D V^{\top} X\right\|_{F}^{2}-\left\|U \overline{D^{(t)}} V^{\top} X\right\|_{F}^{2} \\
& =\left\|D V^{\top} X\right\|_{F}^{2}-\left\|\overline{D^{(t)}} V^{\top} X\right\|_{F}^{2} \\
& =\left\|\left(D-\overline{D^{(t)}}\right) V^{\top} X\right\|_{F}^{2} \\
& =\left\|U\left(D-\overline{D^{(t)}}\right) V^{\top} X\right\|_{F}^{2} \\
& =\|A X-\tilde{A} X\|_{F}^{2} .
\end{aligned}
$$

where the second and fourth equalities follow since $U$ has orthonormal columns, and the third equality follows since for $M=V^{\top} X$ we have

$$
\begin{aligned}
\|D M\|_{F}^{2}-\left\|\overline{D^{(t)}} M\right\|_{F}^{2} & =\sum_{i=1}^{d} \sum_{j=1}^{d} \sigma_{i}^{2}(A) m_{i j}^{2}-\sum_{i=1}^{t} \sum_{j=1}^{d} \sigma_{i}^{2}(A) m_{i j}^{2} \\
& =\sum_{i=t+1}^{d} \sum_{j=1}^{d} \sigma_{i}^{2}(A) m_{i j}^{2}=\left\|\left(D-\overline{D^{(t)}}\right) M\right\|_{F}^{2} .
\end{aligned}
$$

Next, we bound $\|A X-\tilde{A} X\|_{F}^{2}$. We have

$$
\|A X-\tilde{A} X\|_{F}^{2}=\left\|\left(D-\overline{D^{(t)}}\right) V^{\top} X\right\|_{F}^{2} \leq\left\|\left(D-\overline{D^{(t)}}\right)\right\|_{S}^{2}\|X\|_{F}^{2}=r \sigma_{t+1}^{2}(A)
$$

where the inequality follows because the spectral norm is consistent with the Euclidean norm. This implies the lemma since

$$
\begin{equation*}
r \sigma_{t+1}^{2}(A) \leq \epsilon(t-r+1) \sigma_{t+1}^{2}(A) \leq \epsilon \sum_{i=r+1}^{t+1} \sigma_{i}^{2}(A) \leq \epsilon \sum_{i=r+1}^{d} \sigma_{i}^{2}(A) \tag{40}
\end{equation*}
$$

where the first inequality follows for our choice of $t$.

## B.3.2 Proof of Theorem 13

We first introduce some intermediate variables for analysis. Imagine we perform two projections: first project $P_{i}$ to $\tilde{P}_{i}=P_{i} V_{i}{ }^{(t)}\left(V_{i}^{(t)}\right)^{\top}$, then project $\tilde{P}_{i}$ to $\bar{P}_{i}=$ $\tilde{P}_{i} V^{(t)}\left(V^{(t)}\right)^{\top}$. Let $\tilde{P}$ denote the vertical concatenation of $\tilde{P}_{i}$ and let $\bar{P}$ denote the vertical concatenation of $\bar{P}_{i}$, i.e.

$$
\tilde{P}=\left[\begin{array}{c}
\tilde{P}_{1} \\
\vdots \\
\tilde{P}_{s}
\end{array}\right] \quad \text { and } \quad \bar{P}=\left[\begin{array}{c}
\bar{P}_{1} \\
\vdots \\
\bar{P}_{s}
\end{array}\right]
$$

The following is a technical lemma that will be used in the proof of Theorem 13.

## Lemma 27.

$$
d^{2}(\tilde{P}, L(X)) \leq(1+\epsilon) d^{2}(P, L(X))
$$

Proof. We have

$$
\begin{aligned}
d^{2}(\tilde{P}, L(X))-d^{2}(P, L(X)) & =\left\|\tilde{P}-\tilde{P} X X^{\top}\right\|_{F}^{2}-\left\|P-P X X^{\top}\right\|_{F}^{2} \\
& =\|\tilde{P}\|_{F}^{2}-\left\|\tilde{P} X X^{\top}\right\|_{F}^{2}-\left(\|P\|_{F}^{2}-\left\|P X X^{\top}\right\|_{F}^{2}\right) \\
& =\sum_{i=1}^{s}\left[\left\|\tilde{P}_{i}\right\|_{F}^{2}-\left\|P_{i}\right\|_{F}^{2}\right]+\sum_{i=1}^{s}\left[\left\|P_{i} X X^{\top}\right\|_{F}^{2}-\left\|\tilde{P}_{i} X X^{\top}\right\|_{F}^{2}\right] .
\end{aligned}
$$

By the Pythagorean Theorem, $\left\|\tilde{P}_{i}\right\|_{F}^{2} \leq\left\|P_{i}\right\|_{F}^{2}$. Also, since $X$ is orthonormal, $\left\|P_{i} X X^{\top}\right\|_{F}^{2}=$ $\left\|P_{i} X\right\|_{F}^{2}$ and $\left\|\tilde{P}_{i} X X^{\top}\right\|_{F}^{2}=\left\|\tilde{P}_{i} X\right\|_{F}^{2}$. Then

$$
\begin{align*}
d^{2}(\tilde{P}, L(X))-d^{2}(P, L(X)) & \leq \sum_{i=1}^{s}\left[\left\|P_{i} X\right\|_{F}^{2}-\left\|\tilde{P}_{i} X\right\|_{F}^{2}\right] \\
& \leq \sum_{i=1}^{s} \epsilon d^{2}\left(P_{i}, L(X)\right)=\epsilon d^{2}(P, L(X)) \tag{41}
\end{align*}
$$

where the second inequality follows from Lemma 18.

Theorem 13. Let $X$ be a $d \times j$ matrix whose columns are orthonormal. Let $\epsilon \in(0,1]$ and $t \in \mathbf{N}$ with $d-1 \geq t \geq j+\lceil 8 j / \epsilon\rceil-1$. Then the output of Algorithm 11 satisfies $0 \leq\|P X-\hat{P} X\|_{F}^{2} \leq \epsilon d^{2}(P, L(X)) \quad$ and $\quad 0 \leq\|P X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2} \leq \epsilon d^{2}(P, L(X))$.

Proof. For the first statement, we have

$$
\begin{align*}
\|P X-\hat{P} X\|_{F}^{2} & \leq 2\|P X-\tilde{P} X\|_{F}^{2}  \tag{42}\\
& +2\|\tilde{P} X-\bar{P} X\|_{F}^{2}  \tag{43}\\
& +2\|\bar{P} X-\hat{P} X\|_{F}^{2} \tag{44}
\end{align*}
$$

For (42), we have by Lemma 18

$$
\begin{equation*}
\|P X-\tilde{P} X\|_{F}^{2}=\sum_{i=1}^{s}\left\|P_{i} X-\tilde{P}_{i} X\right\|_{F}^{2} \leq \sum_{i=1}^{s} \frac{\epsilon}{4} d^{2}\left(P_{i}, L(X)\right)=\frac{\epsilon}{8} d^{2}(P, L(X)) \tag{45}
\end{equation*}
$$

Similarly, for (43) we have by Lemma 18

$$
\begin{equation*}
\|\tilde{P} X-\bar{P} X\|_{F}^{2} \leq \frac{\epsilon}{8} d^{2}(\tilde{P}, L(X)) \tag{46}
\end{equation*}
$$

To bound (44), let $Y=V^{(t)}\left(V^{(t)}\right)^{\top} X$. Then by definition, $\bar{P}_{i} X=\tilde{P}_{i} Y$ and $\hat{P}_{i} X=$ $P_{i} Y$. By Lemma 18, we have

$$
\begin{align*}
\|\bar{P} X-\hat{P} X\|_{F}^{2} & =\sum_{i=1}^{s}\left\|\tilde{P}_{i} Y-P_{i} Y\right\|_{F}^{2} \leq \sum_{i=1}^{s} \frac{\epsilon}{8} \sum_{i=r+1}^{s} \sigma_{i}^{2}\left(P_{i}\right) \\
& \leq \frac{\epsilon}{8} \sum_{i=1}^{s} d^{2}\left(P_{i}, L(X)\right)=\frac{\epsilon}{8} d^{2}(P, L(X)) \tag{47}
\end{align*}
$$

Combining (45)(46) and (47) leads to

$$
\begin{equation*}
\|P X-\hat{P} X\|_{F}^{2} \leq \frac{\epsilon}{2} d^{2}(P, L(X))+\frac{\epsilon}{4} d^{2}(\tilde{P}, L(X)) \tag{48}
\end{equation*}
$$

The first statement then follows from (48) and Lemma 27.
For the second statement, we have a similar argument.

$$
\begin{align*}
\|P X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2} & =\|P X\|_{F}^{2}-\|\tilde{P} X\|_{F}^{2}  \tag{49}\\
& +\|\tilde{P} X\|_{F}^{2}-\|\bar{P} X\|_{F}^{2}  \tag{50}\\
& +\|\bar{P} X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2} \tag{51}
\end{align*}
$$

For (49), we have by Lemma 18

$$
\begin{align*}
\|P X\|_{F}^{2}-\|\tilde{P} X\|_{F}^{2} & =\sum_{i=1}^{s}\left[\left\|P_{i} X\right\|_{F}^{2}-\left\|\tilde{P}_{i} X\right\|_{F}^{2}\right]  \tag{52}\\
& \leq \sum_{i=1}^{s} \frac{\epsilon}{4} d^{2}\left(P_{i}, L(X)\right)=\frac{\epsilon}{4} d^{2}(P, L(X)) . \tag{53}
\end{align*}
$$

Similarly, for (50) we have by Lemma 18

$$
\begin{equation*}
\|\tilde{P} X\|_{F}^{2}-\|\bar{P} X\|_{F}^{2} \leq \frac{\epsilon}{4} d^{2}(\tilde{P}, L(X)) \tag{54}
\end{equation*}
$$

By Lemma 18, we have

$$
\begin{align*}
\|\bar{P} X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2} & =\sum_{i=1}^{s}\left[\left\|\tilde{P}_{i} Y\right\|_{F}^{2}-\left\|P_{i} Y\right\|_{F}^{2}\right] \leq \sum_{i=1}^{s} \frac{\epsilon}{4} \sum_{i=r+1}^{s} \sigma_{i}^{2}\left(P_{i}\right) \\
& \leq \frac{\epsilon}{4} \sum_{i=1}^{s} d^{2}\left(P_{i}, L(X)\right)=\frac{\epsilon}{4} d^{2}(P, L(X)) . \tag{55}
\end{align*}
$$

Combining (53)(54) and (55) leads to

$$
\begin{equation*}
\|P X\|_{F}^{2}-\|\hat{P} X\|_{F}^{2} \leq \frac{\epsilon}{2} d^{2}(P, L(X))+\frac{\epsilon}{4} d^{2}(\tilde{P}, L(X)) . \tag{56}
\end{equation*}
$$

The second statement then follows from (56) and Lemma 27.

## B.3.3 Proof of Theorem 14

The analysis follows the ideas in [46], but is tailored for the distributed setting. We first begin with the following lemma, showing that the cost of the projected data to any low dimension subspace approximates that of the original data, compensated by a positive constant.

Lemma 28. Let $X$ be a $d \times j$ matrix whose columns are orthonormal. Let $\epsilon \in(0,1]$ and $t \in \mathbf{N}$ with $d-1 \geq t \geq j+\lceil 4 j / \epsilon\rceil-1$. Then there exists $c_{1} \geq 0$ such that

$$
0 \leq d^{2}(\hat{P}, L(X))+c_{1}-d^{2}(P, L(X)) \leq \epsilon d^{2}(P, L(X))
$$

Proof. We have from Pythagorean Theorem

$$
\begin{aligned}
d^{2}(\hat{P}, L(X))-d^{2}(P, L(X)) & =\|\hat{P}\|_{2}^{2}-\|\hat{P} X\|_{2}^{2}-\left(\|P\|_{2}^{2}-\|P X\|_{2}^{2}\right) \\
& =\|P X\|_{2}^{2}-\|\hat{P} X\|_{2}^{2}-c_{1}
\end{aligned}
$$

where $c_{1}=\|P\|_{2}^{2}-\|\hat{P}\|_{2}^{2}$. Note that by Lemma 18 ,

$$
c_{1}=\sum_{i=1}^{n}\left[\left\|P_{i}\right\|_{2}^{2}-\left\|P_{i}^{(t)}\right\|_{2}^{2}\right]+\left\|P^{(t)}\right\|_{2}^{2}-\|\hat{P}\|_{2}^{2} \geq 0
$$

Then the lemma follows from Theorem 13.

The next lemma shows that the projection of the projected data to any low dimension subspace approximates the projection of the projected data in the sense that their distances are small.

Let $p_{i}$ denote the $i$ th row of the data $P$, and let $\hat{p}_{i}$ denote the $i$ th row of $\hat{P}$.
Lemma 29. Let $X$ be a $d \times j$ matrix whose columns are orthonormal. Let $\epsilon \in(0,1]$ and $t \in \mathbf{N}$ with $d-1 \geq t \geq j+\lceil 4 j / \epsilon\rceil-1$. Then

$$
\sum_{i=1}^{|P|} d\left(\Pi_{X}\left(p_{i}\right), \Pi_{X}\left(\hat{p}_{i}\right)\right)^{2} \leq \epsilon d^{2}(P, L(X))
$$

Proof. Since $X$ is orthogonal, $\Pi_{X}(p)=p X X^{T}$. Then

$$
\sum_{i=1}^{|P|} d\left(\Pi_{X}\left(p_{i}\right), \Pi_{X}\left(\hat{p}_{i}\right)\right)^{2}=\sum_{i=1}^{|P|}\left\|p_{i} X X^{T}-\hat{p}_{i} X X^{T}\right\|_{2}^{2}=\left\|P X X^{T}-\hat{P} X X^{T}\right\|_{2}^{2}
$$

This can be simplified to $\|(P-\hat{P}) X\|_{2}^{2}$ since

$$
\begin{aligned}
\left\|P X X^{T}-\hat{P} X X^{T}\right\|_{2}^{2} & \left.=\left\|(P-\hat{P}) X X^{T}\right\|_{2}^{2}=\operatorname{trace}\left[(P-\hat{P}) X X^{T} X X^{T}(P-\hat{P})^{T}\right)\right] \\
& \left.=\operatorname{trace}\left[(P-\hat{P}) X X^{T}(P-\hat{P})^{T}\right)\right]=\|(P-\hat{P}) X\|_{2}^{2}
\end{aligned}
$$

The lemma then follows from Theorem 13.

The above two lemmas are the key elements needed to show our final theorem. Before proving the theorem, we further need the following "weak triangle inequality", which is well known in the coreset literature. The proof is included in the appendix for completeness.

Lemma 30 (Lemma 7.1 in [46]). For any $0 \leq \epsilon \leq 1$, a compact set $C \subseteq \mathbf{R}^{d}$, and $p, q \in \mathbf{R}^{d}$,

$$
\left|d(p, C)^{2}-d(q, C)^{2}\right| \leq \frac{12 d(p, q)^{2}}{\epsilon}+\frac{\epsilon}{2} d(p, C)^{2}
$$

Proof. Using the triangle inequality,

$$
\begin{align*}
d(p, C)^{2}-d(q, C)^{2} \mid & =|d(p, C)-d(q, C)| \cdot(d(p, C)+d(q, C)) \\
& \leq d(p, q)(2 d(p, C)+d(p, q)) \\
& \leq d(p, q)^{2}+2 d(p, C) d(p, q) \tag{57}
\end{align*}
$$

Either $d(p, C) \leq d(p, q) / \epsilon$ or $d(p, q)<\epsilon d(p, C)$. Hence,

$$
d(p, C) d(p, q) \leq \frac{d(p, q)^{2}}{\epsilon}+\epsilon d(p, C)^{2}
$$

Combining the last inequality with (57) yields

$$
\left|d(p, C)^{2}-d(q, C)^{2}\right| \leq d(p, q)^{2}+\frac{2 d(p, q)^{2}}{\epsilon}+2 \epsilon d(p, C)^{2} \leq \frac{3 d(p, q)^{2}}{\epsilon}+2 \epsilon d(p, C)^{2}
$$

Finally, the lemma follows by replacing $\epsilon$ with $\epsilon / 4$.

We are now ready to prove the theorem, which guarantees that a coreset for the output of the distributed PCA algorithm is also a coreset for the original data.

Theorem 14. Let $\mathbf{x}$ be a set of $k$ centers in $\mathbf{R}^{d}$. Let $\epsilon \in(0,1]$ and $t \in \mathbf{N}$ with $d-1 \geq t \geq k+\left\lceil 50 k / \epsilon^{2}\right\rceil$. Then there exists a constant $c_{0} \geq 0$, such that the output of Algorithm 11 satisfies

$$
(1-\epsilon) d^{2}(P, \mathbf{x}) \leq d^{2}(\hat{P}, \mathbf{x})+c_{0} \leq(1+\epsilon) d^{2}(P, \mathbf{x})
$$

Proof. Let $X \in \mathbf{R}^{d \times k}$ has orthonormal columns that span $\mathbf{x}$. Let $c_{0}$ be the constant $c_{1}$ in Lemma 28. Then by Pythagorean theorem we have

$$
\begin{aligned}
& \left|d^{2}(\hat{P}, \mathbf{x})+c_{0}-d^{2}(P, \mathbf{x})\right| \\
= & \left|d^{2}(\hat{P}, L(X))+c_{0}-d^{2}(P, L(X))+\sum_{i=1}^{|P|}\left[d\left(\Pi_{X}\left(p_{i}\right), \mathbf{x}\right)^{2}-d\left(\Pi_{X}\left(\hat{p}_{i}\right), \mathbf{x}\right)^{2}\right]\right|
\end{aligned}
$$

By Lemma 28 we have

$$
\begin{equation*}
\left|d^{2}(\hat{P}, L(X))+c_{0}-d^{2}(P, L(X))\right| \leq \frac{\epsilon^{2}}{4} d^{2}(P, L(X)) \tag{58}
\end{equation*}
$$

By Lemma 29 and Lemma 30 we have

$$
\begin{align*}
\sum_{i=1}^{|P|}\left|d\left(\Pi_{X}\left(p_{i}\right), \mathbf{x}\right)^{2}-d\left(\Pi_{X}\left(\hat{p}_{i}\right), \mathbf{x}\right)^{2}\right| & \leq \sum_{i=1}^{|P|}\left[\frac{12 d\left(\Pi_{X}\left(p_{i}\right), \Pi_{X}\left(\hat{p}_{i}\right)\right)^{2}}{\epsilon}+\frac{\epsilon}{2} d\left(\Pi_{X}\left(p_{i}\right), \mathbf{x}\right)^{2}\right] \\
& \leq \frac{\epsilon}{4} d^{2}(P, \mathbf{x})+\frac{\epsilon}{2} \sum_{i=1}^{|P|} d\left(\Pi_{X}\left(p_{i}\right), \mathbf{x}\right)^{2} \tag{59}
\end{align*}
$$

Since $d^{2}(P, L(X)) \leq d^{2}(P, \mathbf{x})$ and $d\left(\Pi_{X}\left(p_{i}\right), \mathbf{x}\right) \leq d\left(p_{i}, \mathbf{x}\right)$, the theorem follows from (58) and (59).

## REFERENCES

[1] Ackerman, M. and Ben-David, S., "Measures of clustering quality: A working set of axioms for clustering," in Advances in Neural Information Processing Systems, 2008.
[2] Airoldi, E. M., Blei, D. M., Fienberg, S. E., and Xing, E. P., "Mixed membership stochastic blockmodels," The Journal of Machine Learning Research, 2008.
[3] Albert, R. and Barabási, A.-L., "Statistical mechanics of complex networks," Reviews of Modern Physics, 2002.
[4] Arora, S., Ge, R., Sachdeva, S., and Schoenebeck, G., "Finding overlapping communities in social networks: toward a rigorous approach," in Proceedings of the ACM Conference on Electronic Commerce, 2012.
[5] Arthur, D. and Vassilvitskii, S., "k-means++: The advantages of careful seeding," in Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2007.
[6] Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., and Pandit, V., "Local search heuristics for k-median and facility location problems," SIAM Journal of Computing, 2004.
[7] Awasthi, P. and Balcan, M., "Center based clustering: A foundational perspective." Survey Chapter in Handbook of Cluster Analysis (Manuscript), 2013.
[8] Awasthi, P., Blum, A., and Sheffet, O., "Stability yields a ptas for kmedian and k-means clustering," in Proceedings of the IEEE Annual Symposium on Foundations of Computer Science, 2010.
[9] Awasthi, P., Blum, A., and Sheffet, O., "Center-based clustering under perturbation stability," Information Processing Letters, 2012.
[10] Awasthi, P. and Sheffet, O., "Improved spectral-norm bounds for clustering," in Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques, 2012.
[11] Bache, K. and Lichman, M., "UCI machine learning repository," 2013.
[12] Bahmani, B., Moseley, B., Vattani, A., Kumar, R., and Vassilvitskif, S., "Scalable k-means++," in Proceedings of the International Conference on Very Large Data Bases, 2012.
[13] Bai, Z.-J., Chan, R. H., and Luk, F. T., "Principal component analysis for distributed data sets with updating," in Advanced Parallel Processing Technologies, 2005.
[14] Balcan, M. F., Blum, A., and Gupta, A., "Approximate clustering without the approximation," in Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2009.
[15] Balcan, M. F. and Braverman, M., "Approximate nash equilibria under stability conditions," CoRR, vol. abs/1008.1827, 2010.
[16] Balcan, M. F. and Gupta, P., "Robust hierarchical clustering," in Proceedings of the Annual Conference on Learning Theory, 2010.
[17] Balcan, M.-F., Berlind, C., Ehrlich, S., and Liang, Y., "Efficient semisupervised and active learning of disjunctions," in Proceedings of the International Conference on Machine Learning, 2013.
[18] Balcan, M.-F., Blum, A., Fine, S., and Mansour, Y., "Distributed learning, communication complexity and privacy," in Proceedings of the Conference on Learning Thoery, 2012.
[19] Balcan, M. F. and Liang, Y., "Clustering under perturbation resilience," in Proceedings of the International Colloquium on Automata, Languages, and Programming, 2012.
[20] Balcan, M. F. and Liang, Y., "Modeling and detecting community hierarchies," in Proceedings of the International Workshop on Similarity-Based Pattern Analysis and Recognition, 2013.
[21] Balcan, M., Borgs, C., Braverman, M., Chayes, J., and Teng, S., "Finding endogenously formed communities," in Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2013.
[22] Bandelt, H. and Dress, A., "Weak hierarchies associated with similarity measures - an additive clustering technique," Bulletin of Mathematical Biology, 1989.
[23] Bartal, Y., Charikar, M., and Raz, D., "Approximating min-sum $k$ clustering in metric spaces," in Proceedings of the Annual ACM Symposium on Theory of Computing, 2001.
[24] Bellet, A., Liang, Y., Garakani, A. B., Balcan, M.-F., and Sha, F., "Distributed frank-wolfe algorithm: A unified framework for communicationefficient sparse learning," CoRR, vol. abs/1404.2644, 2014.
[25] Ben-David, S., "A framework for statistical clustering with a constant time approximation algorithms for k-median clustering," Proceedings of Annual Conference on Learning Theory, 2004.
[26] Bilu, Y. and Linial, N., "Are stable instances easy?," in Proceedings of the Symposium on Innovations in Computer Science, 2010.
[27] Bilu, Y., Daniely, A., Linial, N., and Saks, M., "On the practically interesting instances of maxcut," in Proceedings of the International Symposium on Theoretical Aspects of Computer Science, 2013.
[28] Bounova, G. and de Weck, O., "Overview of metrics and their correlation patterns for multiple-metric topology analysis on heterogeneous graph ensembles," Physical Review E, 2012.
[29] Charikar, M., Guha, S., Tardos, É., and Shmoys, D. B., "A constantfactor approximation algorithm for the k-median problem," Journal of Computer and System Sciences, 2002.
[30] Chen, K., "On k-median clustering in high dimensions," in Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2006.
[31] Cheng, D., Kannan, R., Vempala, S., and Wang, G., "A divide-andmerge methodology for clustering," ACM Transactions on Database Systems, vol. 31, no. 4, 2006.
[32] Clauset, A., Moore, C., and Newman, M., "Hierarchical structure and the prediction of missing links in networks," Nature, 2008.
[33] Clauset, A., Newman, M. E. J., and Moore, C., "Finding community structure in very large networks," Physical Review E, 2004.
[34] Condon, A. and Karp, R. M., "Algorithms for graph partitioning on the planted partition model," Random Structures and Algorithms, vol. 18, no. 2, 2001.
[35] Considine, J., Li, F., Kollios, G., and Byers, J., "Approximate aggregation techniques for sensor databases," in Proceedings of the International Conference on Data Engineering, 2004.
[36] Corbett, J. C., Dean, J., Epstein, M., Fikes, A., Frost, C., Furman, J., Ghemawat, S., Gubarev, A., Heiser, C., Hochschild, P., and othERS, "Spanner: Googles globally-distributed database," in Proceedings of the USENIX Symposium on Operating Systems Design and Implementation, 2012.
[37] Dasgupta, S., "Learning mixtures of gaussians," in Proceedings of the Annual Symposium on Foundations of Computer Science, 1999.
[38] Datta, S., Giannella, C., Kargupta, H., and others, "K-means clustering over peer-to-peer networks," in Proceedings of the International Workshop on High Performance and Distributed Mining, 2005.
[39] Daumé III, H., Phillips, J. M., Saha, A., and Venkatasubramanian, S., "Efficient protocols for distributed classification and optimization," in Algorithmic Learning Theory, pp. 154-168, Springer, 2012.
[40] de la Vega, W. F., Karpinski, M., Kenyon, C., and Rabani, Y., "Approximation schemes for clustering problems," in Proceedings of the Annual ACM Symposium on Theory of Computing, 2003.
[41] Ding, C. and He, X., "K-means clustering via principal component analysis," in Proceedings of the International Conference on Machine learning, 2004.
[42] Du, N., Liang, Y., Balcan, M.-F., and Song, L., "Budgeted influence maximization for multiple products," $C o R R$, vol. abs/1312.2164, 2013.
[43] Du, N., Liang, Y., Balcan, M.-F., and Song, L., "Influence function learning in information diffusion networks," 2014.
[44] Everitt, B. S., Landau, S., Leese, M., and Stahl, D., Hierarchical Clustering. 2011.
[45] Feldman, D. and Langberg, M., "A unified framework for approximating and clustering data," in Proceedings of the Annual ACM Symposium on Theory of Computing, 2011.
[46] Feldman, D., Schmidt, M., and Sohler, C., "Turning big data into tiny data: Constant-size coresets for k-means, pca and projective clustering," in Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2013.
[47] Feldman, D., Sugaya, A., and Rus, D., "An effective coreset compression algorithm for large scale sensor networks," in Proceedings of the International Conference on Information Processing in Sensor Networks, 2012.
[48] Forman, G. and Zhang, B., "Distributed data clustering can be efficient and exact," ACM SIGKDD Explorations Newsletter, 2000.
[49] Fortunato, S., "Community detection in graphs," Physics Reports, 2010.
[50] Ghashami, M. and Phillips, J., "Relative errors for deterministic low-rank matrix approximations," in Proceedings of the Twenty-Fifth Annual ACMSIAM Symposium on Discrete Algorithms, 2014.
[51] Girvan, M. and Newman, M. E. J., "Community structure in social and biological networks," Proceedings of the National Academy of Sciences, 2002.
[52] Girvan, M. and Newman, M. E. J., "Community structure in social and biological networks," Proceedings of the National Academy of Sciences, 2002.
[53] Greenhill, S. and Venkatesh, S., "Distributed query processing for mobile surveillance," in Proceedings of the International Conference on Multimedia, 2007.
[54] Greenwald, M. and Khanna, S., "Power-conserving computation of orderstatistics over sensor networks," in Proceedings of the ACM SIGMOD-SIGACTSIGART Symposium on Principles of Database Systems, 2004.
[55] Guha, S. and Khuller, S., "Greedy strikes back: Improved facility location algorithms," Journal of Algorithms, 1999.
[56] Har-Peled, S. and Kushal, A., "Smaller coresets for k-median and k-means clustering," Discrete 6 Computational Geometry, 2007.
[57] Har-Peled, S. and Mazumdar, S., "On coresets for k-means and k-median clustering," in Proceedings of the Annual ACM Symposium on Theory of Computing, 2004.
[58] He, J., Hopcroft, J., Liang, H., Suwajanakorn, S., and Wang, L., "Detecting the structure of social networks using $(\alpha, \beta)$-communities," in Proceedings of the International Conference on Algorithms and Models for the Web Graph, 2011.
[59] Jain, K., Mahdian, M., and Saberi, A., "A new greedy approach for facility location problems," in Proceedings of the Annual ACM Symposium on Theory of Computing, 2002.
[60] Januzaj, E., Kriegel, H., and Pfeifle, M., "Towards effective and efficient distributed clustering," in Workshop on Clustering Large Data Sets in the IEEE International Conference on Data Mining, 2003.
[61] Kannan, R., Vempala, S., and Vetta, A., "On clusterings: Good, bad and spectral," Journal of the ACM, vol. 51, no. 3, 2004.
[62] Kannan, R. and Vempala, S., "Nimble algorithms for cloud computing," CoRR, vol. abs/1304.3162, 2013.
[63] Kanungo, T., Mount, D. M., Netanyahu, N. S., Piatko, C. D., Silverman, R., and Wu, A. Y., "A local search approximation algorithm for k-means clustering," in Proceedings of the Annual Symposium on Computational Geometry, 2002.
[64] Kargupta, H., Huang, W., Sivakumar, K., and Johnson, E., "Distributed clustering using collective principal component analysis," Knowledge and Information Systems, 2001.
[65] Kim, M. and Leskovec, J., "Latent multi-group membership graph model," in Proceedings of the International Conference on Machine Learning, 2012.
[66] King, B., "Step-wise clustering procedures," Journal of the American Statistical Association, 1967.
[67] Kleinberg, J., "An impossibility theorem for clustering," in Advances in Neural Information Processing Systems, 2002.
[68] Krebs, V., "http://www.orgnet.com/," unpublished.
[69] Kumar, A. and Kannan, R., "Clustering with spectral norm and the k-means algorithm," in Proceedings of the IEEE Annual Symposium on Foundations of Computer Science, 2010.
[70] Kumar, A. and Kannan, R., "Clustering with spectral norm and the k-means algorithm," in Proceedings of the Annual IEEE Symposium on Foundations of Computer Science, 2010.
[71] Lagomarsino, M. C., Jona, P., Bassetti, B., and Isambert, H., "Hierarchy and feedback in the evolution of the Escherichia coli transcription network," Proceedings of the National Academy of Sciences, 2007.
[72] Lancichinetti, A. and Fortunato, S., "Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities," Physical Review E, 2009.
[73] Lancichinetti, A., Fortunato, S., and Kertész, J., "Detecting the overlapping and hierarchical community structure in complex networks," New Journal of Physics, 2009.
[74] Langberg, M. and Schulman, L., "Universal $\varepsilon$-approximators for integrals," in Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2010.
[75] Le Borgne, Y.-A., Raybaud, S., and Bontempi, G., "Distributed principal component analysis for wireless sensor networks," Sensors, 2008.
[76] Li, S. and Svensson, O., "Approximating k-median via pseudoapproximation," in Proceedings of the ACM Symposium on the Theory of Computing, 2013.
[77] Li, Y., Long, P. M., and Srinivasan, A., "Improved bounds on the sample complexity of learning," in Proceedings of the eleventh annual ACM-SIAM Symposium on Discrete Algorithms, 2000.
[78] Liang, Y., Balcan, M.-F., and Ehrlich, S., "Distributed k-means and kmedian clustering on general topologies," in Advances in Neural Information Processing Systems, 2013.
[79] Liang, Y., Balcan, M.-F., and Kanchanapally, V., "Distributed pca and k-means clustering," in The Big Learning Workshop, Advances in Neural Information Processing Systems, 2013.
[80] Lipton, R. J., Markakis, E., and Mehta, A., "On stability properties of economic solution concepts." Manuscript, 2006.
[81] Lloyd, S., "Least squares quantization in pcm," IEEE Transactions on Information Theory, vol. 28, no. 2, 1982.
[82] Lusseau, D., Schneider, K., Boisseau, O. J., Haase, P., Slooten, E., and Dawson, S. M. Behavioral Ecology and Sociobiology, 2003.
[83] Macua, S. V., Belanovic, P., and Zazo, S., "Consensus-based distributed principal component analysis in wireless sensor networks," in Proceedings of the IEEE International Workshop on Signal Processing Advances in Wireless Communications, 2010.
[84] Makarychev, K., Makarychev, Y., and Vijayaraghavan, A., "Bilulinial stable instances of max cut," CoRR, vol. abs/1305.1681, 2013.
[85] Matula, D. W. and Shahrokhi, F., "Sparsest cuts and bottlenecks in graphs," Discrete Applied Mathematics, vol. 27, no. 1, 1990.
[86] Mihalák, M., Schöngens, M., Šrámek, R., and Widmayer, P., "On the complexity of the metric tsp under stability considerations," in Proceedings of the International Conference on Current Trends in Theory and Practice of Computer Science, 2011.
[87] Mitra, S., Agrawal, M., Yadav, A., Carlsson, N., Eager, D., and Mahanti, A., "Characterizing web-based video sharing workloads," ACM Transactions on the Web, 2011.
[88] Newman, M. E. J., "Detecting community structure in networks," The European Physical Journal B - Condensed Matter and Complex Systems, 2004.
[89] Newman, M. E. J., "Modularity and community structure in networks," Proceedings of the National Academy of Sciences, 2006.
[90] Newman, M. E. and Girvan, M., "Mixing patterns and community structure in networks," in Statistical mechanics of complex networks, 2003.
[91] Olston, C., Jiang, J., and Widom, J., "Adaptive filters for continuous queries over distributed data streams," in Proceedings of the ACM SIGMOD International Conference on Management of Data, 2003.
[92] Ostrovsky, R., Rabani, Y., Schulman, L. J., and Swamy, C., "The effectiveness of lloyd-type methods for the k-means problem," in Proceedings of the Annual IEEE Symposium on Foundations of Computer Science, 2006.
[93] Qu, Y., Ostrouchov, G., Samatova, N., and Geist, A., "Principal component analysis for dimension reduction in massive distributed data sets," in Proceedings of IEEE International Conference on Data Mining, 2002.
[94] Ravasz, E., Somera, A. L., Mongru, D. A., Oltvai, Z. N., and Barabási, A. L., "Hierarchical Organization of Modularity in Metabolic Networks," Science, 2002.
[95] Reyzin, L., "Data stability in clustering: A closer look," in Proceedings of the International Conference on Algorithmic Learning Theory, 2012.
[96] Schölkopf, B. and Smola, A. J., Learning with Kernels. 2002.
[97] Schweinberger, M. and Snijders, T. A. B., "Settings in social networks: A measurement model," Sociological Methodology, 2003.
[98] Sneath, P. H. A. and Sokal, R. R., Numerical taxonomy. The principles and practice of numerical classification. 1973.
[99] Spielman, D. A. and Teng, S.-H., "Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems," in Proceedings of the Annual ACM Symposium on Theory of Computing, 2004.
[100] Steinbach, M., Karypis, G., Kumar, V., and others, "A comparison of document clustering techniques," in KDD workshop on text mining, vol. 400, 2000.
[101] Tasoulis, D. and Vrahatis, M., "Unsupervised distributed clustering," in Proceedings of the International Conference on Parallel and Distributed Computing and Networks, 2004.
[102] Vapnik, V. N., Statistical learning theory. 1998.
[103] Zachary, W. W., "An information flow model for conflict and fission in small groups," Journal of Anthropological Research, 1977.
[104] Zhang, Q., Liu, J., and Wang, W., "Approximate clustering on distributed data streams," in Proceedings of the IEEE International Conference on Data Engineering, 2008.

## VITA

Yingyu Liang was born in Pingnan, Guangxi Province of China. After completing his schoolwork at Guigang High School in Guangxi in 2004, Yingyu entered Tsinghua University in Beijing. He received a Bachelor of Science in July 2008 and a Master of Science in July 2010 with a major in computer science and technology from Tsinghua University. In August 2010, he entered the School of Computer Science in Georgia Institute of Technology in Atlanta, Georgia.


[^0]:    ${ }^{1}$ Bilu and Linial [26] refer to such instances as perturbation stable instances.

[^1]:    ${ }^{2}$ Of course, the score of the optimal solution will change; what the definition requires is that the partitioning induced by the optimum remains the same.

[^2]:    ${ }^{3}$ For clarity, in this paper efficient means polynomial in both $n$ (the number of points) and $k$ (the number of clusters).

[^3]:    ${ }^{1}$ For $k$-median and $k$-means in general metric spaces, the bound on the size of the coreset can be obtained by replacing $d$ with the logarithm of the total number of points. The analysis for general metric spaces is largely the same as that for $d$ dimensional Euclidean space, so we will focus on Euclidean space and point out the difference when needed.

[^4]:    ${ }^{2}$ The value communicated is the sum of the costs of approximations to the local optimal clustering. This is guaranteed to be no more than a constant factor times larger than the optimal cost.

[^5]:    ${ }^{3}$ For both $k$-median and $k$-means in general metric spaces, $\operatorname{dim}(F, P)=O(k \log |P|)$, so the bound for general metric spaces (including Euclidean space we focus on) can be obtained by replacing $d$ with $\log |P|$.

[^6]:    ${ }^{4}$ Their algorithm used coreset construction as a subroutine. The construction algorithm they used builds coreset of size $\tilde{O}\left(\frac{n k h}{\epsilon^{d}} \log |P|\right)$. Throughout this paper, when we compare to [104] we assume they use the coreset construction technique of [45] to reduce their coreset size and communication cost.

[^7]:    ${ }^{5}$ Our theoretical analysis shows that our algorithm has better bounds on the communication cost. Since the bounds are from worst-case analysis, it is meaningful to verify that our algorithm

[^8]:    ${ }^{1}$ We distinguish the nodes in the hierarchy our algorithm builds from the points in the graph.

[^9]:    ${ }^{2}$ Detailed descriptions and links for download can be found on the following website: http://www-personal.umich.edu/ mejn/netdata/.

[^10]:    ${ }^{3}$ The source code we use and details about the parameters can be found on https://sites.google.com/site/andrealancichinetti/software.

