

The Price of Hierarchical Clustering

Anna Arutyunova ✉

Universität Bonn, Germany

Heiko Röglin ✉

Universität Bonn, Germany

Abstract

Hierarchical Clustering is a popular tool for understanding the hereditary properties of a data set. Such a clustering is actually a sequence of clusterings that starts with the trivial clustering in which every data point forms its own cluster and then successively merges two existing clusters until all points are in the same cluster. A hierarchical clustering achieves an approximation factor of α if the costs of each k -clustering in the hierarchy are at most α times the costs of an optimal k -clustering. We study as cost functions the maximum (discrete) radius of any cluster (k -center problem) and the maximum diameter of any cluster (k -diameter problem).

In general, the optimal clusterings do not form a hierarchy and hence an approximation factor of 1 cannot be achieved. We call the smallest approximation factor that can be achieved for any instance the *price of hierarchy*. For the k -diameter problem we improve the upper bound on the price of hierarchy to $3 + 2\sqrt{2} \approx 5.83$. Moreover we significantly improve the lower bounds for k -center and k -diameter, proving a price of hierarchy of exactly 4 and $3 + 2\sqrt{2}$, respectively.

2012 ACM Subject Classification Theory of computation → Facility location and clustering

Keywords and phrases Hierarchical Clustering, approximation Algorithms, k -center Problem

Digital Object Identifier 10.4230/LIPIcs.ESA.2022.10

Related Version *Full Version:* <https://arxiv.org/abs/2205.01417> [4]

Funding This work has been supported by the German Research Association (DFG) grant RO 5439/1-1.

1 Introduction

Clustering is an ubiquitous task in data analysis and machine learning. In a typical clustering problem, the goal is to partition a set of objects into different clusters such that only similar objects belong to the same cluster. There are numerous ways how clustering can be modeled formally and many different models have been studied in the literature in the last decades. In many theoretical models, one assumes that the data comes from a metric space and that the desired number of clusters is given. Then the goal is to optimize some objective function like k -center, k -median, or k -means. In most cases the resulting optimization problems are NP-hard and hence approximation algorithms have been studied extensively.

One aspect of real-world clustering problems that is not captured by these models is that it is often already a non-trivial task to determine for a given data set the right or most reasonable number of clusters. One particularly appealing way to take this into account is hierarchical clustering. A hierarchical clustering of a data set is actually a sequence of clusterings, one for each possible number of clusters. It starts with the trivial clustering in which every data point forms its own cluster and then successively merges two existing clusters until all points are in the same cluster. This way for every possible number of clusters, a clustering is obtained. These clusterings help to understand the hereditary properties of the data and they provide information at different levels of granularity.



© Anna Arutyunova and Heiko Röglin;
licensed under Creative Commons License CC-BY 4.0
30th Annual European Symposium on Algorithms (ESA 2022).

Editors: Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman; Article No. 10; pp. 10:1–10:14

Leibniz International Proceedings in Informatics



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

While hierarchical clustering is successfully used in many applications, it is not as well understood from a theoretical point of view as the models in which the number of clusters is given as part of the input. One reason for this is that it is not obvious how the quality of a hierarchical clustering should be measured. A possibility that has been explored in the literature is to define the quality of a hierarchical clustering based on its worst level. To be precise, let (\mathcal{X}, d) be a metric space and $\mathcal{P} \subset \mathcal{X}$ a set of n points. Furthermore let $\mathcal{H} = (\mathcal{H}_n, \dots, \mathcal{H}_1)$ be a hierarchical clustering of \mathcal{P} , where \mathcal{H}_k denotes a k -clustering, i.e., a clustering with at most k non-empty clusters. Then \mathcal{H}_{k-1} arises from \mathcal{H}_k by merging some of the existing clusters. We assume that some objective function like k -center, k -median, or k -means is selected and denote by $\text{cost}(\mathcal{H}_k)$ the objective value of \mathcal{H}_k with respect to the selected objective function. Furthermore, let \mathcal{O}_k denote an optimal k -clustering and let $\text{cost}(\mathcal{O}_k)$ denote its objective value. Then we say that \mathcal{H} achieves an approximation factor of $\alpha \geq 1$ if $\text{cost}(\mathcal{H}_k) \leq \alpha \cdot \text{cost}(\mathcal{O}_k)$ for every k , assuming that cost is an objective that is to be minimized. In this work we consider the radius objective, which is well-known from the k -center problem. Here the cost is defined as the maximum radius of a cluster. Furthermore we consider the diameter objective, where the cost is defined as the maximum distance between any two points lying in the same cluster.

An α -approximation for small α yields a strong guarantee for the hierarchical clustering on every level. However, in general there do not exist optimal clusterings $\mathcal{O}_n, \dots, \mathcal{O}_1$ that form a hierarchy. So even with unlimited computational resources, a 1-approximation usually cannot be achieved. In the literature different algorithms for computing hierarchical clusterings with respect to different objective functions have been developed and analyzed. Dasgupta and Long [13] and Charikar et al. [8] initiated this line of research and presented both independently from each other an algorithm that computes efficiently an 8-approximate hierarchical clustering with respect to the radius and diameter objective. That is, for every level k , the maximal radius or diameter of any cluster in the k -clustering computed by their algorithms is at most 8 times the maximal radius or diameter in an optimal k -clustering. Inspired by [13], Plaxton [21] proposed a constant-factor approximation for the k -median and k -means objective. Later a general framework that also leads constant approximation guarantees for many objective functions including in particular k -median and k -means has been proposed by Lin et al. [19].

Despite these articles and other related work, which we discuss below in detail, many questions in the area of hierarchical clustering are not yet resolved. We find it particularly intriguing to find out which approximation factors can be achieved for different objectives. This question comes in two flavors depending on the computational resources available. Of course it is interesting to study which approximation factors can and cannot be achieved in polynomial time, assuming $P \neq NP$. Since in general there do not exist hierarchical clusterings that are optimal on each level, it is also interesting to study which approximation factors can and cannot be achieved in general without the restriction to polynomial-time algorithms.

For an objective function like radius or diameter we define its *price of hierarchy* as the smallest α such that for any instance there exists an α -approximate hierarchical clustering. Hence, the price of hierarchy is a measure for how much quality one has to sacrifice for the hierarchical structure of the clusterings.

Our main results are tight bounds for the price of hierarchy for the radius, discrete radius and diameter objective. Here the difference between radius and discrete radius lies in the choice of centers. For the radius objective we allow to choose the center of a cluster $C \subset \mathcal{P}$ from the whole metric space \mathcal{X} , while for the discrete radius objective the center must be contained in C itself. We will see that this has an impact on the price of hierarchy.

For all three objectives the algorithms in [13, 8] compute an 8-approximate hierarchical clustering in polynomial time. Until recently this was also the best known upper bound for the price of hierarchy in the literature for hierarchical radius and diameter. For discrete radius, Großwendt [17] shows an upper bound for the price of hierarchy of 4. The best known lower bounds are 2, proven by Das and Kenyon-Mathieu [11] for diameter and by Großwendt [17] for (discrete) radius. We improve the framework in [19] for radius and diameter and show an upper bound on the price of hierarchy of $3 + 2\sqrt{2} \approx 5.83$. The upper bound of $3 + 2\sqrt{2}$ for the radius was also recently proved by Bock [5] in independent work. However our main contribution lies in the design of clustering instances to prove a lower bound of 4 for discrete radius and $3 + 2\sqrt{2}$ for radius and diameter.

Related work. Gonzales [14] presents a simple and elegant incremental algorithm for k -center. The algorithm exhibits the following nice property: given a set \mathcal{P} which has to be clustered, it returns an ordering of the points, such that the first k points constitute the centers of the k -center solution, and this solution is a 2-approximation for every $1 \leq k \leq |\mathcal{P}|$. However the resulting clusterings are usually not hierarchically compatible. Dasgupta and Long [13] use the ordering computed by Gonzales' algorithm to compute a hierarchical clustering. The authors present an 8-approximation for the objective functions (discrete) radius and diameter. In an independent work Charikar et al. [8] also present an 8-approximation for the three objectives which outputs the same clustering as the algorithm in [13] under some reasonable conditions [11]. In a recent work, Mondal [20] gives a 6-approximation for hierarchical (discrete) radius. In the full version of this paper [4] we present an instance where this algorithm computes only a 7-approximation contradicting the claimed guarantee.

Plaxton [21] shows that a similar approach as in [13] yields a hierarchical clustering with constant approximation guarantee for the k -median and k -means objectives. Later a general framework for a variety of incremental and hierarchical problems was introduced by Lin et al. [19]. Their framework can be applied to compute hierarchical clusterings for any cost function which satisfies a certain nesting property, especially those of k -median and k -means. This yields a 20.71α -approximation for k -median and a 576β -approximation for k -means. Here $\alpha = 2.675$ and $\beta = 6.357$ are the currently best approximation guarantees for k -median [6] and k -means [2]. The algorithms presented in [8, 13, 19, 21] run in polynomial time. Unless $P=NP$ there is no polynomial-time α -approximation for $\alpha < 2$ for hierarchical (discrete) radius and diameter. For (discrete) radius this is an immediate consequence of the reduction from dominating set presented by [18]. A similar reduction from clique cover yields the statement for hierarchical diameter.

However even without time constraints it is not clear what approximation guarantee can be achieved for hierarchical clustering. It is easy to find examples, where the approximation guarantee of any hierarchical clustering for all three objectives is greater than one. Das and Kenyon-Mathieu [11] and Großwendt [17] present instances for diameter and (discrete) radius, where no hierarchical clustering has an approximation guarantee smaller than 2. On the other hand Großwendt [17] proves an upper bound of 4 on the approximation guarantee of hierarchical discrete radius by using the framework of Lin et al. [19]. In recent independent work Bock [5] improved the bound for hierarchical radius to $3 + 2\sqrt{2}$. While his approach is inspired by Dasgupta and Long [13], the resulting algorithm is similar to the algorithm we present in this paper as an improvement of [19].

Aside from the theoretical results, there also exist greedy heuristics, which are more commonly used in applications. One very simple bottom up, also called agglomerative, algorithm is the following: starting from the clustering where every point is separate, it

merges in every step the two clusters whose merge results in the smallest increase of the cost function. For (discrete) radius and diameter this algorithm is known as complete linkage and for the k -means cost this is Ward's method [23]. Ackermann et al. [1] analyze the approximation guarantee of complete linkage in the Euclidean space. They show an approximation guarantee of $O(\log(k))$ for all three objectives assuming the dimension of the Euclidean space to be constant. This was later improved by Großwendt and Röglin [15] to $O(1)$. In arbitrary metric spaces complete linkage does not perform well. There Arutyunova et al. [3] prove a lower bound of $\Omega(k)$ for all three objectives. For Ward's method Großwendt et al. [16] show an approximation guarantee of 2 under the strong assumption that the optimal clusters are well separated.

Recently other cost functions for hierarchical clustering were proposed, which do not compare to the optimal clustering on every level. Dasgupta [12] defines a new cost function for similarity measures and presents an $O(\alpha \log(n))$ -approximation for the respective problem. This was later improved to $O(\alpha)$ independently by Charikar and Chatziafratis [7] and Cohen-Addad et al. [9]. Here α is the approximation guarantee of sparsest cut. However Cohen-Addad et al. [9] prove that every hierarchical clustering is an $O(1)$ -approximation to the corresponding cost function for dissimilarity measures when the dissimilarity measure is a metric. A cost function more suitable for Euclidean spaces was developed by Wang and Moseley [22]. They prove that a randomly generated hierarchical clustering performs poorly for this cost function and show that bisecting k -means computes an $O(1)$ -approximation.

Our results. We define the *price of hierarchy* ρ_{cost} with respect to an objective function cost as the smallest number such that for every clustering instance there exists a hierarchical clustering which is a ρ_{cost} -approximation with respect to cost . Observe that the results [8, 11, 13, 17] imply that the price of hierarchy for radius and diameter is between 2 and 8 and for discrete radius between 2 and 4. We close these gaps and prove that the price of hierarchy for radius and diameter is exactly $3 + 2\sqrt{2}$ and for discrete radius exactly 4. Notice that this does not imply the existence of polynomial-time algorithms with approximation guarantee ρ_{cost} . Especially our algorithm which computes a $3 + 2\sqrt{2}$ -approximation for radius and diameter does not run in polynomial time. This is also the case for the $3 + 2\sqrt{2}$ -approximation for radius presented by Bock [5] in independent work. Our upper bound of $3 + 2\sqrt{2}$ can be achieved by a small improvement in the framework of Lin et al. [19]. However our most technically demanding contribution is the design of a clustering instance for every $\epsilon > 0$ such that every hierarchical clustering has approximation guarantee at least $3 + 2\sqrt{2} - \epsilon$ for radius and diameter and $4 - \epsilon$ for discrete radius. It requires a careful analysis of all possible hierarchical clusterings, which is highly non-trivial for complex clustering instances.

2 Preliminaries

A clustering instance $(\mathcal{X}, \mathcal{P}, d)$ consists of a metric space (\mathcal{X}, d) and a finite subset $\mathcal{P} \subset \mathcal{X}$. For a set (or cluster) $C \subset \mathcal{P}$ we denote by

$$\text{diam}(C) = \max_{p, q \in C} d(p, q)$$

the *diameter* of C . By $\text{rad}(C, c) = \max_{p \in C} d(c, p)$ we denote the radius of C with respect to a center $c \in \mathcal{X}$. This is the largest distance between c and a point in C . The *radius* of C is defined as the smallest radius of C with respect to a center $c \in \mathcal{X}$, i.e.,

$$\text{rad}(C) = \min_{c \in \mathcal{X}} \text{rad}(C, c)$$

while the *discrete radius* of C is defined as the smallest radius of C with respect to a center $c \in C$, i.e.,

$$\text{drad}(C) = \min_{c \in C} \text{rad}(C, c).$$

A k -clustering of \mathcal{P} is a partition of \mathcal{P} into at most k non-empty subsets. We consider three closely related clustering problems.

The k -diameter problem asks to minimize $\text{diam}(\mathcal{C}_k) = \max_{C \in \mathcal{C}_k} \text{diam}(C)$, i.e., the maximum diameter of a k -clustering \mathcal{C}_k . In the k -center problem we want to minimize the maximum radius $\text{rad}(\mathcal{C}_k) = \max_{C \in \mathcal{C}_k} \text{rad}(C)$, and in the discrete k -center problem we want to minimize the maximum discrete radius $\text{drad}(\mathcal{C}_k) = \max_{C \in \mathcal{C}_k} \text{drad}(C)$.

► **Definition 1.** Given an instance $(\mathcal{X}, \mathcal{P}, d)$, let $n = |\mathcal{P}|$. We call two clusterings \mathcal{C} and \mathcal{C}' of \mathcal{P} with $|\mathcal{C}| \geq |\mathcal{C}'|$ hierarchically compatible if for all $C \in \mathcal{C}$ there exists $C' \in \mathcal{C}'$ with $C \subset C'$. A hierarchical clustering of \mathcal{P} is a sequence of clusterings $\mathcal{H} = (\mathcal{H}_n, \dots, \mathcal{H}_1)$, such that

1. \mathcal{H}_i is an i -clustering of \mathcal{P}
2. for $1 < i \leq n$ the two clusterings \mathcal{H}_{i-1} and \mathcal{H}_i are hierarchically compatible.

For $\text{cost} \in \{\text{diam}, \text{rad}, \text{drad}\}$ let \mathcal{O}_i denote the optimal i -clustering with respect to cost . We say that \mathcal{H} is an α -approximation with respect to cost if for all $i = 1, \dots, n$ we have

$$\text{cost}(\mathcal{H}_i) \leq \alpha \cdot \text{cost}(\mathcal{O}_i).$$

Since optimal clusterings are generally not hierarchically compatible, there is usually no hierarchical clustering with approximation guarantee $\alpha = 1$. We have to accept that the restriction on hierarchically compatible clusterings comes with an unavoidable increase in the cost compared to an optimal solution.

► **Definition 2.** For $\text{cost} \in \{\text{diam}, \text{rad}, \text{drad}\}$ the price of hierarchy $\rho_{\text{cost}} \geq 1$ is defined as follows.

1. For every instance $(\mathcal{X}, \mathcal{P}, d)$, there exists a hierarchical clustering \mathcal{H} of \mathcal{P} that is a ρ_{cost} -approximation with respect to cost .
2. For any $\alpha < \rho_{\text{cost}}$ there exists an instance $(\mathcal{X}, \mathcal{P}, d)$, such that there is no hierarchical clustering of \mathcal{P} that is an α -approximation with respect to cost .

Thus ρ_{cost} is the smallest possible number such that for every clustering instance there is a hierarchical clustering with approximation guarantee ρ_{cost} .

3 An Upper Bound on the Price of Hierarchy

It is already known that the framework of Lin et al. [19] yields an upper bound of 4 on the price of hierarchy for the discrete radius [17]. This framework also yields upper bounds for the price of hierarchy for radius and diameter, which are not tight, however. We present an improved version that yields the following better upper bound on the price of hierarchy for radius and diameter.

► **Theorem 3.** For $\text{cost} \in \{\text{diam}, \text{rad}\}$ we have $\rho_{\text{cost}} \leq 3 + 2\sqrt{2} \approx 5.828$.

For the details we refer to the full version [4].

4 A Lower Bound on the Price of Hierarchy

The most challenging contributions of this article are matching lower bounds on the price of hierarchy for diameter, radius, and discrete radius.

► **Theorem 4.** *For $\text{cost} \in \{\text{diam}, \text{rad}\}$ we have $\rho_{\text{cost}} \geq 3 + 2\sqrt{2}$ and for $\text{cost} = \text{drad}$ we have $\rho_{\text{cost}} \geq 4$.*

There is already existing work in this area by Das and Kenyon-Mathieu [11] for the diameter and Großwendt [17] for the radius. Both show a lower bound of 2 for the respective objective. To improve upon these results we have to construct much more complex instances which differ significantly from those in [11, 17].

For every $\epsilon > 0$ we will construct a clustering instance $(\mathcal{X}, \mathcal{P}, d)$ such that for any hierarchical clustering $\mathcal{H} = (\mathcal{H}_{|\mathcal{P}|}, \dots, \mathcal{H}_1)$ of \mathcal{P} there is $1 \leq i \leq |\mathcal{P}|$ such that $\text{cost}(\mathcal{H}_i) \geq \alpha \cdot \text{cost}(\mathcal{O}_i)$, where \mathcal{O}_i is an optimal i -clustering of \mathcal{P} with respect to cost and $\alpha = (3 + 2\sqrt{2} - \epsilon)$ for $\text{cost} \in \{\text{diam}, \text{rad}\}$ and $\alpha = 4 - \epsilon$ for $\text{cost} = \text{drad}$.

The proof is divided in three parts. First we introduce the clustering instance $(\mathcal{X}, \mathcal{P}, d)$ and determine its optimal clusterings. In the second part we develop the notion of a *bad* cluster. We prove that any hierarchical clustering contains such bad clusters and develop a lower bound on their cost. In the third part we compare the lower bound to the cost of optimal clusterings and prove Theorem 4.

4.1 Definition of the Clustering Instance

For $n \in \mathbb{N}$ we denote by $[n]$ the set of numbers from 1 to n .

Let $k \in \mathbb{N}$ and $\Gamma = k + 1$. For $0 \leq \ell \leq k$ we define point sets \mathcal{Q}_ℓ and \mathcal{P}_ℓ recursively as follows:

1. For $\ell = 0$ let $\mathcal{P}_0 = \mathcal{Q}_0 = [1]$ and denote by N_0 the cardinality of \mathcal{P}_0 .
 2. For $\ell > 0$ let $\mathcal{Q}_\ell = [\Gamma \cdot N_{\ell-1}]^{N_{\ell-1}}$ and $\mathcal{P}_\ell = \prod_{i=0}^{\ell} \mathcal{Q}_i$. Furthermore set $N_\ell = |\mathcal{P}_\ell|$.
- Moreover let $\phi_\ell: \mathcal{P}_\ell \rightarrow [N_\ell]$ be a bijection for $0 \leq \ell \leq k$.

We refer to a point $X \in \mathcal{P}_k$ as a matrix with $k + 1$ rows and $N_{\ell-1}$ entries in the ℓ -th row. Thus we write

$$X = (x_{01} \mid \dots \mid x_{\ell 1}, \dots, x_{\ell N_{\ell-1}} \mid \dots \mid x_{k1}, \dots, x_{k N_{k-1}}).$$

Let $X_\ell = (x_{\ell 1}, \dots, x_{\ell N_{\ell-1}}) \in \mathcal{Q}_\ell$ for $0 \leq \ell \leq k$. For a shorter representation we can replace the ℓ -th row directly by X_ℓ and for $0 \leq i \leq j \leq k$ we can replace the i -th up to j -th row by $X_{[i:j]} = (X_i \mid \dots \mid X_j)$.

Let $X \in \mathcal{P}_k$ and $1 \leq \ell \leq k$. Notice that $X_{[0:\ell-1]} \in \mathcal{P}_{\ell-1}$ and let $m = \phi_{\ell-1}(X_{[0:\ell-1]})$, we define

$$A_\ell^X = \{(X_{[0:\ell-1]} \mid x_{\ell 1}, \dots, x_{\ell m-1}, \star, x_{\ell m+1}, \dots, x_{\ell N_{\ell-1}} \mid X_{[\ell+1:k]}) \mid \star \in [\Gamma \cdot N_{\ell-1}]\}.$$

Thus all coordinates of points in A_ℓ^X are fixed and agree with those of X except one which is variable. Here $X_{[0:\ell-1]}$ serves as prefix which indicates through $\phi_{\ell-1}$ which coordinate of X_ℓ can be changed.

We define $\mathcal{A}_\ell = \{A_\ell^X \mid X \in \mathcal{P}_k\}$ as the set containing all subsets of this form. It is clear that \mathcal{A}_ℓ is a partition of \mathcal{P}_k and that it contains only sets of size $\Gamma \cdot N_{\ell-1}$. Furthermore we set $\mathcal{A}_0 = \{\{X\} \mid X \in \mathcal{P}_k\}$.

► **Example 5.** If we perform the first three steps of the construction, we get $\mathcal{Q}_0 = [1]$, $\mathcal{Q}_1 = [\Gamma]$, $\mathcal{Q}_2 = [\Gamma^2]^\Gamma$ and

$$\mathcal{P}_1 = \{(1 \mid x_{11}) \mid x_{11} \in [\Gamma]\},$$

$$\mathcal{P}_2 = \{(1 \mid x_{11} \mid x_{21}, \dots, x_{2\Gamma}) \mid x_{11} \in [\Gamma], x_{2i} \in [\Gamma^2] \text{ for } 1 \leq i \leq \Gamma\}.$$

Since ϕ_0 is a map between two sets of cardinality one, this map is always unique. Now suppose that we picked ϕ_1 such that $\phi_1((1 \mid x_{11})) = x_{11}$ for all $(1 \mid x_{11}) \in \mathcal{P}_1$. Then the partition \mathcal{A}_1 consists of the sets

$$\{(1 \mid \star \mid x_{21}, \dots, x_{2\Gamma}) \mid \star \in [\Gamma]\}$$

with $x_{2i} \in [\Gamma^2]$ for all $1 \leq i \leq \Gamma$. The partition \mathcal{A}_2 consists of the sets

$$\{(1 \mid x_{11} \mid x_{21}, \dots, x_{2x_{11}-1}, \star, x_{2x_{11}+1}, \dots, x_{2\Gamma}) \mid \star \in [\Gamma^2]\}$$

with $x_{11} \in [\Gamma]$ and $x_{2i} \in [\Gamma^2]$ for all $1 \leq i \leq \Gamma$ with $i \neq x_{11}$. ┘

Now let $G = (V, E, w)$ denote the weighted hyper-graph with $V = \mathcal{P}_k$ and $E = \bigcup_{i=1}^k \mathcal{A}_i$. The weight of a hyper-edge $e \in E$ is set to ℓ iff $e \in \mathcal{A}_\ell$. For $0 \leq \ell \leq k$, the sub-graph $G_\ell = (V_\ell, E_\ell, w_\ell)$ is given by $V_\ell = \mathcal{P}_k$, $E_\ell = \bigcup_{i=0}^\ell \mathcal{A}_i$ and $w_\ell = w|_{E_\ell}$.

We extend G to a hyper-graph $H = (V', E', w')$ as follows. Let $V' = V \cup \bigcup_{i=0}^k \{v_A \mid A \in \mathcal{A}_i\}$ and $E' = E \cup \bigcup_{i=0}^k \{\{v, v_A\} \mid A \in \mathcal{A}_i, v \in A\}$. Thus H contains one vertex for every $A \in \bigcup_{i=0}^k \mathcal{A}_i$ and this vertex is connected by edges to every vertex $v \in A$. For $e \in E$ we set $w'(e) = w(e)$ and for $e = \{v, v_A\}$ for some $A \in \mathcal{A}_\ell$ and $v \in A$ we set $w'(e) = \ell/2$.

The clustering instance $(\mathcal{X}, \mathcal{P}, d)$ is given by $\mathcal{X} = V'$, $\mathcal{P} = V$, and d as the shortest path metric on H . Observe that the extension of G to H is only necessary for the lower bound for the radius but not for the diameter and the discrete radius. This is because the additional points $V' \setminus V$ do not belong to \mathcal{P} and are hence irrelevant for the clustering instance for the diameter and discrete radius. In the lower bound for the radius they will be used as centers, however.

► **Example 6.** For $k = 2$ we obtain $\Gamma = 3$ and $\mathcal{P} = \mathcal{P}_2$. Suppose that we again picked ϕ_1 such that $\phi_1((1 \mid x_{11})) = x_{11}$ for all $(1 \mid x_{11}) \in \mathcal{P}_1$. By the above definition the shortest path between the two points $X = (1 \mid 1 \mid 1, 1, 1)$, $Y = (1 \mid 1 \mid 2, 2, 1) \in \mathcal{P}_2$ in G is of the form

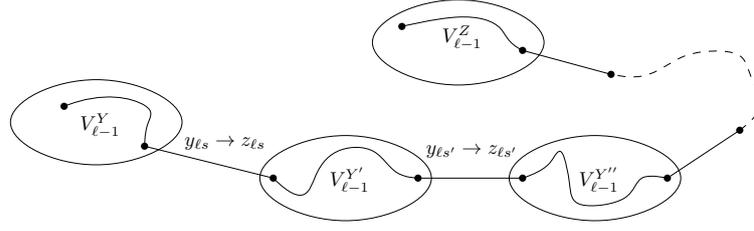
$$X = (1 \mid 1 \mid 1, 1, 1), (1 \mid 1 \mid 2, 1, 1), (1 \mid 2 \mid 2, 1, 1), (1 \mid 2 \mid 2, 2, 1), (1 \mid 1 \mid 2, 2, 1) = Y.$$

Thus the distance is given by $d(X, Y) = 2 + 1 + 2 + 1 = 6$. ┘

► **Lemma 7.** *Let $p, q \in V$, then $d(p, q)$ is the length of a shortest path between p and q in G .*

Proof. By definition $d(p, q)$ is the length of a shortest path between p and q in H . Suppose the shortest path contains a vertex v_A for some $A \in \bigcup_{i=0}^k \mathcal{A}_i$ with $v \in A$ as predecessor and $w \in A$ as ancestor. Since v and w are connected in H by the hyper-edge A we can delete v_A from the path and the length of the path does not change. The resulting path is also a path in G , so $d(p, q)$ is also the length of a shortest path between p and q in G . ◀

Next we state some structural properties of the graph G and the clustering instance $(\mathcal{X}, \mathcal{P}, d)$. To establish a lower bound on the approximation factor of a hierarchical clustering we first focus on the optimal clusterings of the instance $(\mathcal{X}, \mathcal{P}, d)$. One can already guess that \mathcal{A}_ℓ is an optimal clustering with $\frac{N_k}{\Gamma N_{\ell-1}}$ clusters with respect to $\text{cost} \in \{\text{diam}, \text{rad}, \text{drad}\}$ and we will prove this in this section. First we need the following statement about the connected components of G_ℓ .



■ **Figure 1** Here we see the construction of the path. It corresponds to changing the coordinates of Y successively until they match Z . We use an edge in \mathcal{A}_ℓ to change $y_{\ell s}$ to $z_{\ell s}$, next we change $y_{\ell s'}$ to $z_{\ell s'}$ and proceed like this until we obtain Z . The respective edges are then connected to a path from $V_{\ell-1}^X$ to $V_{\ell-1}^Z$.

► **Lemma 8.** *The vertex set of every connected component in G_ℓ has cardinality N_ℓ and is of the form $V_\ell^X = \{(X' | X) | X' \in \mathcal{P}_\ell\}$ for a given $X = (X_{\ell+1} | \dots | X_k) \in \prod_{i=\ell+1}^k \mathcal{Q}_i$.*

Proof. Notice that $|V_\ell^X| = N_\ell$ and that $\{V_\ell^X | X \in \prod_{i=\ell+1}^k \mathcal{Q}_i\}$ is a partition of V . Furthermore since $E_\ell = \bigcup_{i=0}^{\ell} \mathcal{A}_i$ any edge $e \in E_\ell$ is either completely contained in or disjoint to V_ℓ^X .

It is left to show that V_ℓ^X is connected. We prove this via induction over ℓ . For $\ell = 0$ this is clear because $|V_0^X| = 1$. For $\ell > 0$ let $Y = (Y_\ell | X), Z = (Z_\ell | X) \in \prod_{i=\ell}^k \mathcal{Q}_i$. By the induction hypothesis we know that the sets $V_{\ell-1}^Y, V_{\ell-1}^Z$ are connected. To prove that V_ℓ^X is connected it is sufficient to show that there is a path from a point in $V_{\ell-1}^Y$ to a point in $V_{\ell-1}^Z$. We show this claim by induction over the number m of coordinates in which Y and Z differ. For $m = 0$ there is nothing to show. If $m > 0$ pick $1 \leq s \leq N_{\ell-1}$ such that $y_{\ell s} \neq z_{\ell s}$ and let $P = \phi_{\ell-1}^{-1}(s) \in \prod_{i=0}^{\ell-1} \mathcal{Q}_i$. Consider the point $(P | Y_\ell | X)$ which is contained in $V_{\ell-1}^Y$. This point is also contained in the set

$$\{(P | y_{\ell 1}, \dots, y_{\ell s-1}, \star, y_{\ell s+1}, \dots, y_{\ell N_{\ell-1}} | X) | \star \in [\Gamma \cdot N_{\ell-1}]\} \in E_\ell.$$

Thus there is an edge in G_ℓ connecting a point in $V_{\ell-1}^Y$ to a point in $V_{\ell-1}^{Y'}$ with $Y' = (y_{\ell 1}, \dots, y_{\ell s-1}, z_{\ell s}, y_{\ell s+1}, \dots, y_{\ell N_{\ell-1}} | X)$. Now Y' and Z differ in $m - 1$ coordinates, thus there is a path between two points in $V_{\ell-1}^{Y'}$ and $V_{\ell-1}^Z$ by induction hypothesis. If we combine this with the induction hypothesis that $V_{\ell-1}^{Y'}$ is connected this yields the claim (see Figure 1 for an illustration). ◀

► **Lemma 9.** *Any clustering of $(\mathcal{X}, \mathcal{P}, d)$ with less than $\frac{N_k}{N_{\ell-1}}$ clusters costs at least ℓ if $\text{cost} \in \{\text{diam}, \text{drad}\}$ and $\ell/2$ if $\text{cost} = \text{rad}$.*

Proof. The shortest path in G between any two points which lie in different connected components of $G_{\ell-1}$ must contain an edge of weight $\geq \ell$. Thus any set of points $M \subset V$ which is disconnected in $G_{\ell-1}$ has diameter $\geq \ell$. Remember that the discrete radius of M is given by $\text{drad}(M) = \min_{c \in M} \max_{p \in M} d(p, c)$. For every possible choice of $c \in M$ there exists a point $p \in M$ which is not in the same connected component of $G_{\ell-1}$ as c , thus $d(c, p) \geq \ell$ and therefore $\text{drad}(M) \geq \ell$ and $\text{rad}(M) \geq \text{diam}(M)/2 \geq \ell/2$.

We conclude that if $\text{cost} \in \{\text{diam}, \text{drad}\}$ any cluster of cost smaller than ℓ is contained in one of the sets $V_{\ell-1}^X$ for some $X \in \prod_{i=\ell}^k \mathcal{Q}_i$ by Lemma 8 and any clustering with less than $|\prod_{i=\ell}^k \mathcal{Q}_i|$ clusters costs at least ℓ . By the same argument if $\text{cost} = \text{rad}$ any cluster of cost smaller than $\ell/2$ is contained in one of the sets $V_{\ell-1}^X$ for some $X \in \prod_{i=\ell}^k \mathcal{Q}_i$ by Lemma 8 and any clustering with less than $|\prod_{i=\ell}^k \mathcal{Q}_i|$ clusters costs at least $\ell/2$. Since

$$\left| \prod_{i=\ell}^k \mathcal{Q}_i \right| = \frac{\left| \prod_{i=0}^k \mathcal{Q}_i \right|}{\left| \prod_{i=0}^{\ell-1} \mathcal{Q}_i \right|} = \frac{N_k}{N_{\ell-1}}$$

this proves the lemma. ◀

► **Corollary 10.** *For $1 \leq \ell \leq k$ and $\text{cost} \in \{\text{diam}, \text{rad}, \text{drad}\}$ the clustering \mathcal{A}_ℓ is an optimal $\frac{N_k}{\Gamma N_{\ell-1}}$ -clustering for the instance $(\mathcal{X}, \mathcal{P}, d)$. Furthermore $\text{diam}(\mathcal{A}_\ell) = \text{drad}(\mathcal{A}_\ell) = \ell$ and $\text{rad}(\mathcal{A}_\ell) = \ell/2$.*

4.2 Characterization of Hierarchical Clusterings

Let from now on $\mathcal{H} = (\mathcal{H}_{N_k}, \dots, \mathcal{H}_1)$ denote a hierarchical clustering of $(\mathcal{X}, \mathcal{P}, d)$. We introduce the notion of *bad clusters* in $\mathcal{H}_{\frac{N_k}{\Gamma N_{\ell-1}}}$ which are clusters whose cost increases repeatedly, as we will see later. In this section we prove the existence of such clusters in \mathcal{H} and we give a lower bound on their cost.

► **Definition 11.** *We call all clusters $C \in \mathcal{H}_{N_k}$ bad at time 0 and denote by $\text{Ker}_0(C) = C$ the kernel of C at time 0 and set $\text{Bad}(0) = \mathcal{H}_{N_k}$.*

For $1 \leq \ell \leq k$ we say that a cluster $C \in \mathcal{H}_{\frac{N_k}{\Gamma N_{\ell-1}}}$ is anchored at $\ell \leq \ell' \leq k$ if the set

$$\bigcup_{D \in \text{Bad}(\ell-1): D \subset C} \text{Ker}_{\ell-1}(D) \text{ is}$$

1. *connected in $G_{\ell'}$,*
2. *disconnected in $G_{\ell'-1}$.*

We call C bad at time ℓ if C is anchored at some $\ell' \geq \ell$. We denote by $\text{Bad}(\ell) \subset \mathcal{H}_{\frac{N_k}{\Gamma N_{\ell-1}}}$ the set of all bad clusters at time ℓ . If C is bad we define the kernel of C as the union of all kernels of bad clusters at time $\ell - 1$ contained in C , i.e.,

$$\text{Ker}_\ell(C) = \bigcup_{D \in \text{Bad}(\ell-1): D \subset C} \text{Ker}_{\ell-1}(D).$$

All clusters in $\mathcal{H}_{\frac{N_k}{\Gamma N_{\ell-1}}} \setminus \text{Bad}(\ell)$ are called good.

The example in Figure 2 shows that a bad cluster at time ℓ can contain clusters which are good at time $\ell - 1$. However we are only interested in points that are contained exclusively in bad clusters at any time $t < \ell$. The set $\text{Ker}_\ell(C)$ contains exactly such points.

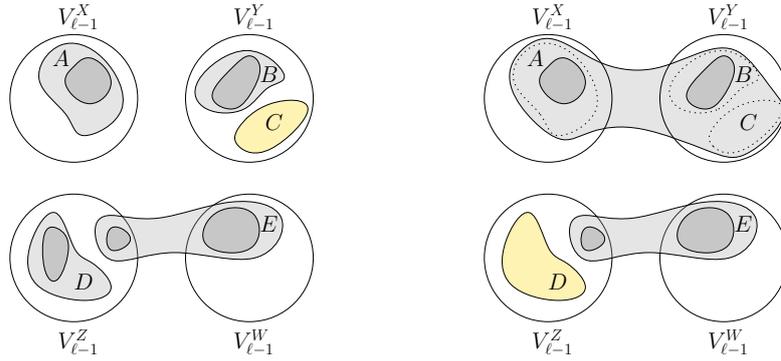
We will use two crucial properties to prove the final lower bound on the approximation factor of any hierarchical clustering \mathcal{H} of $(\mathcal{X}, \mathcal{P}, d)$. We first observe that bad clusters exist in \mathcal{H} for every time-step $1 \leq \ell \leq k$ and second that these clusters have a large cost compared to the optimal clustering.

► **Lemma 12.** *For all $0 \leq \ell \leq k$ we have $\sum_{C \in \text{Bad}(\ell)} |\text{Ker}_\ell(C)| \geq \frac{\Gamma - \ell}{\Gamma} N_k$.*

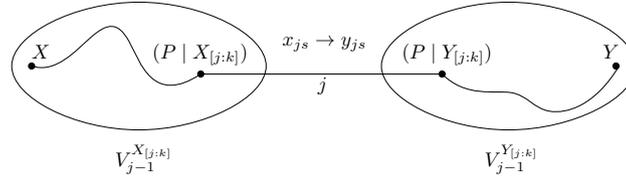
An immediate consequence of Lemma 12 is the existence of bad clusters at time ℓ for any $0 \leq \ell \leq k$. To prove that their (discrete) radius and diameter is indeed large we need a lower bound on the distance between two points $X, Y \in \mathcal{P}$ that lie in different connected components of G_{j-1} for some $1 \leq j \leq k$.

Suppose that the points X and Y only differ in one coordinate, i.e., there is a $1 \leq s \leq N_{j-1}$ such that $x_{js} \neq y_{js}$, while X and Y agree in all other coordinates. There is only one edge in G_j connecting $V_{j-1}^{X_{[j:k]}}$ with $V_{j-1}^{Y_{[j:k]}}$. Let $P = \phi_{j-1}^{-1}(s)$, then this edge connects the points

10:10 The Price of Hierarchical Clustering



■ **Figure 2** An illustration of the evolution of good and bad clusters: In the example, we see five clusters at time $\ell - 1$. The clusters A, B, D, E are assumed to be bad, with their kernels depicted in dark gray, while C is assumed to be a good cluster. At time ℓ , clusters A, B and C are merged. The resulting cluster is bad because the kernels of A and B lie in different connected components of $G_{\ell-1}$. Clusters D and E are still present at time ℓ , but now D is a good cluster since its kernel is completely contained in $V_{\ell-1}^Z$, while E is still bad, since its kernel is disconnected in $G_{\ell-1}$.



■ **Figure 3** A shortest path between X and Y . It consists of two shortest paths inside the connected components of G_{j-1} and the unique edge of weight j between these components.

$(P | X_{[j:k]})$ and $(P | Y_{[j:k]})$. If we connect X to $(P | X_{[j:k]})$ and $(P | Y_{[j:k]})$ to Y via a shortest path, this results in a path from X to Y , see Figure 3. We show that this path is indeed a shortest path between X and Y and generalize this to arbitrary X and Y which are disconnected in G_{j-1} .

► **Lemma 13.** *Let $X, Y \in \mathcal{P}$ be two points and suppose there is $1 \leq j \leq k$ and $1 \leq s \leq N_{j-1}$ such that $x_{js} \neq y_{js}$. Let $P = \phi_{j-1}^{-1}(s) \in \prod_{i=0}^{j-1} \mathcal{Q}_i$. Then*

$$d(X, Y) \geq d(X, (P | X_{[j:k]})) + j + d(Y, (P | Y_{[j:k]})).$$

We now define the so called *anchor set* $\text{Anc}_\ell(C)$ of a bad cluster C at time ℓ . If C is anchored at ℓ' then $\text{Anc}_\ell(C)$ is the union of ℓ' and the anchor set of some bad cluster $D \subset C$ at time $\ell - 1$. If we choose D appropriately the sum of anchors in $\text{Anc}_\ell(C)$ is a lower bound on the discrete radius of C , as we show later. It is clear that ℓ' itself is a lower bound on the discrete radius since $\text{Ker}_\ell(C)$ is disconnected in $G_{\ell'-1}$ by definition. If we additionally assume that the discrete radius of D is large, e.g., lower bounded by the sum of anchors in $\text{Anc}_{\ell-1}(D)$, then it is reasonable to assume that the discrete radius of C is lower bounded by some function in ℓ' and the sum of anchors in $\text{Anc}_{\ell-1}(D)$. First we give a formal definition of $\text{Anc}_\ell(C)$ and how to choose D .

► **Definition 14.** *Let $1 \leq \ell \leq k$ and C be a bad cluster at time ℓ which is anchored at $\ell' \geq \ell$. If $\ell = 1$ we define the anchor set of C as $\text{Anc}_1(C) = \{\ell'\}$ and set $\text{prev}(C) = \{X\}$ for some $X \in C$.*

For $\ell > 1$ we distinguish two cases.

Case 1: C contains a bad cluster D which is bad at time $\ell - 1$ and anchored at ℓ' . We then set $\text{Anc}_\ell(C) = \text{Anc}_{\ell-1}(D)$ and $\text{prev}(C) = D$.

Case 2: C does not contain such a cluster. Then let $D \subset C$ be a bad cluster at time $\ell - 1$ minimizing

$$\sum_{a \in \text{Anc}_{\ell-1}(D)} a$$

among all clusters $D' \in \text{Bad}(\ell - 1)$ with $D' \subset C$. We set $\text{Anc}_\ell(C) = \text{Anc}_{\ell-1}(D) \cup \{\ell'\}$ and $\text{prev}(C) = D$.

Observe that in Case 2 of the previous definition, the bad cluster D must be anchored at some $\ell_D < \ell'$.

With the help of Lemma 13 we are able to show how the discrete radius and diameter of a bad cluster, depends on the sum of anchors.

► **Lemma 15.** *Let $1 \leq \ell \leq k$ and C be a bad cluster at time ℓ anchored at ℓ' . Then for any point $Z \in \mathcal{P}$ there is $X \in \text{Ker}_\ell(C)$ such that*

$$d(Z, X) \geq \sum_{a \in \text{Anc}_\ell(C)} a.$$

► **Lemma 16.** *Let $1 \leq \ell \leq k$ and C be a bad cluster at time ℓ anchored at ℓ' . Then there are two points $X, Y \in \text{Ker}_\ell(C)$ such that*

$$d(X, Y) \geq \ell' + 2 \sum_{a \in \text{Anc}_\ell(C) \setminus \{\ell'\}} a.$$

4.3 Comparison to Optimal Clusterings

Our initial motivation was to construct an instance where any hierarchical clustering has a high approximation ratio. If we consider an arbitrary time $1 \leq \ell \leq k$ then the hierarchical clustering \mathcal{H} on $(\mathcal{X}, \mathcal{P}, d)$ may be even optimal at time ℓ . Thus the bounds which we develop in Lemma 15 and Lemma 16 on the discrete radius and diameter of bad clusters are useless without linking the cost of a bad cluster at time ℓ to the cost of bad clusters at other time steps. Therefore we construct a sequence of clusters $C_1 \subset C_2 \dots \subset C_k$ where C_i is a bad cluster at time i such that $\text{Anc}_1(C_1) \subset \text{Anc}_2(C_2) \subset \dots \subset \text{Anc}_k(C_k)$. We then show with the help of Lemma 15 and Lemma 16 that at least one of these clusters has a high discrete radius and diameter compared to the optimal cost.

► **Lemma 17.** *Let C_k be a bad cluster at time k . For $1 \leq i \leq k - 1$ we define $C_i = \text{prev}(C_{i+1})$. For all $1 \leq i \leq k - 1$ cluster C_i is bad at time i and one of the following two cases occurs:*

1. $\text{Anc}_i(C_i) = \text{Anc}_{i+1}(C_{i+1})$,
2. $\text{Anc}_{i+1}(C_{i+1}) \setminus \{\ell\} = \text{Anc}_i(C_i)$, where $\ell = \max \text{Anc}_{i+1}(C_{i+1})$.

Proof. For $i = k$ cluster C_k is bad at time k by assumption. If C_{i+1} is a bad cluster at time $i + 1$ then $C_i = \text{prev}(C_{i+1})$ is a bad cluster at time i , by definition of prev .

Let C_i be anchored at $\ell' \geq i$ and C_{i+1} be anchored at $\ell \geq i + 1$. Since C_i is a bad cluster at time i with $C_i \subset C_{i+1}$ we have by definition of $\text{Ker}_{i+1}(C_{i+1})$ that $\text{Ker}_i(C_i) \subset \text{Ker}_{i+1}(C_{i+1})$ and thus $\ell' \leq \ell$. If $\ell' = \ell$ we obtain by Definition 14, that $\text{Anc}_i(C_i) = \text{Anc}_{i+1}(C_{i+1})$, so the lemma holds in this case.

If $\ell' < \ell$ we know by Definition 14 that $\text{Anc}_i(C_i) = \text{Anc}_{i+1}(C_{i+1}) \setminus \{\ell\}$. So the lemma also holds in this case. ◀

10:12 The Price of Hierarchical Clustering

► **Corollary 18.** *Let C_k be a bad cluster at time k . For $1 \leq i \leq k-1$ we define $C_i = \text{prev}(C_{i+1})$. Let $\text{Anc}_k(C_k) = \{\ell_1, \dots, \ell_s\}$ such that $\ell_{t-1} < \ell_t$ for all $2 \leq t \leq s$ and let $\ell_0 = 0$. Then for any $1 \leq t \leq s$ and for any i with $\ell_{t-1} < i \leq \ell_t$, we have $\{\ell_1, \dots, \ell_t\} \subset \text{Anc}_i(C_i)$.*

Proof. We prove this via induction over i , starting from $i = k$ in decreasing order. There is nothing to show for $i = k$. For $i < k$ we distinguish two cases. If $\text{Anc}(C_i) = \text{Anc}_{i+1}(C_{i+1})$, the lemma follows from the induction hypothesis.

Otherwise remember that $\text{Anc}_i(C_i) \subset \text{Anc}_k(C_k)$ and $\ell_{t-1} < i$. Since $\max \text{Anc}_i(C_i) \geq i$ we obtain that $\max \text{Anc}_i(C_i) \in \{\ell_t, \dots, \ell_s\}$ and therefore $\ell_t \leq \max \text{Anc}_i(C_i)$.

By Lemma 17 we know that $\text{Anc}_i(C_i) = \text{Anc}_{i+1}(C_{i+1}) \setminus \{\ell\}$, where $\ell = \max \text{Anc}_{i+1}(C_{i+1})$. Thus $\ell_t \leq \max \text{Anc}_i(C_i) < \max \text{Anc}_{i+1}(C_{i+1}) = \ell$ and by induction hypothesis we obtain

$$\{\ell_1, \dots, \ell_t\} \subset \text{Anc}_{i+1}(C_{i+1}) \setminus \{\ell\} = \text{Anc}_i(C_i). \quad \blacktriangleleft$$

Before we are able to prove the theorem we need some final lemma.

► **Lemma 19.** *For every $\epsilon > 0$ there exists $k \in \mathbb{N}$ such that for every $s \in \mathbb{N}$ any sequence of $s+1$ numbers $(\ell_0, \dots, \ell_s) \in \mathbb{R}_{\geq 0}^{s+1}$ with $\ell_0 = 0$ and $\ell_s = k$ satisfies the following.*

1. *There exists $1 \leq t \leq s$ such that for $\alpha_1 = 4 - \epsilon$ and $\Delta_1 = 1$ we have*

$$\frac{\ell_t + \Delta_1 \sum_{i=0}^{t-1} \ell_i}{\ell_{t-1} + 1} > \alpha_1.$$

2. *There exists $1 \leq t \leq s$ such that for $\alpha_2 = 3 + 2\sqrt{2} - \epsilon$ and $\Delta_2 = 2$ we have*

$$\frac{\ell_t + \Delta_2 \sum_{i=0}^{t-1} \ell_i}{\ell_{t-1} + 1} > \alpha_2.$$

► **Theorem 4.** *For $\text{cost} \in \{\text{diam}, \text{rad}\}$ we have $\rho_{\text{cost}} \geq 3 + 2\sqrt{2}$ and for $\text{cost} = \text{drad}$ we have $\rho_{\text{cost}} \geq 4$.*

Proof. Let $\epsilon > 0$ and k be the respective number from Lemma 19. We claim that the approximation factor of any hierarchical clustering $\mathcal{H} = (\mathcal{H}_{N_k}, \dots, \mathcal{H}_1)$ on the instance $(\mathcal{X}, \mathcal{P}, d)$ is larger than $3 + 2\sqrt{2} - \epsilon$ if $\text{cost} \in \{\text{diam}, \text{rad}\}$ and larger than $4 - \epsilon$ if $\text{cost} = \text{drad}$. First we use Lemma 12 to observe that there is a cluster $C_k \in \mathcal{H}_{\frac{N_k}{\Gamma N_{k-1}}}$ that is bad at time k . For $1 \leq i \leq k-1$ we define $C_i = \text{prev}(C_{i+1})$. Let $\text{Anc}_k(C_k) = \{\ell_1, \dots, \ell_s\}$ with $\ell_{t-1} < \ell_t$ for $2 \leq t \leq s$ and let $\ell_0 = 0$. We know by Corollary 18, that for any $1 \leq t \leq s$ and for $i = \ell_{t-1} + 1$ we have $\{\ell_1, \dots, \ell_t\} \subset \text{Anc}_i(C_i)$. Let $\ell' = \max \text{Anc}_i(C_i)$, we obtain by Lemma 16 and Lemma 15 that

$$\begin{aligned} \text{diam}(C_i) &\geq \ell' + 2 \sum_{a \in \text{Anc}_i(C_i) \setminus \{\ell'\}} a \geq \ell_t + 2 \sum_{u=1}^{t-1} \ell_u, \\ \text{rad}(C_i) &\geq \frac{\text{diam}(C_i)}{2} \geq \frac{\ell_t + 2 \sum_{u=1}^{t-1} \ell_u}{2}, \\ \text{drad}(C_i) &\geq \sum_{a \in \text{Anc}_i(C_i)} a \geq \sum_{u=1}^t \ell_u. \end{aligned}$$

Remember that by Corollary 10 \mathcal{A}_i is an optimal $\frac{N_k}{\Gamma N_{i-1}}$ -clustering with $\text{cost}(\mathcal{A}_i) = i$ if $\text{cost} \in \{\text{diam}, \text{drad}\}$ and $\text{cost}(\mathcal{A}_i) = i/2$ if $\text{cost} = \text{rad}$. We obtain

$$\frac{\text{rad}(C_i)}{\text{rad}(\mathcal{A}_i)} = \frac{2\text{rad}(C_i)}{2\text{rad}(\mathcal{A}_i)} \geq \frac{\text{diam}(C_i)}{\text{diam}(\mathcal{A}_i)} \geq \frac{\ell_t + 2\sum_{u=1}^{t-1} \ell_u}{\ell_{t-1} + 1}$$

$$\frac{\text{drad}(C_i)}{\text{drad}(\mathcal{A}_i)} \geq \frac{\sum_{u=1}^t \ell_u}{\ell_{t-1} + 1}$$

which are lower bounds on the approximation factor of \mathcal{H} .

We apply Lemma 19 on (ℓ_0, \dots, ℓ_s) to observe that there is $1 \leq t' \leq s$ such that

$$\frac{\ell_{t'} + 2\sum_{u=1}^{t'-1} \ell_u}{\ell_{t'-1} + 1} > 3 + 2\sqrt{2} - \epsilon$$

and an $1 \leq t'' \leq s$ such that

$$\frac{\sum_{u=1}^{t''} \ell_u}{\ell_{t''-1} + 1} > 4 - \epsilon.$$

This proves the theorem. ◀

5 Conclusions and Open Problems

We have proved tight bounds for the price of hierarchy with respect to the diameter and (discrete) radius. It would be interesting to also obtain a better understanding of the price of hierarchy for other important objective functions like k -median and k -means. The best known upper bound is 16 for k -median [10] and 32 for k -means [17] but no non-trivial lower bounds are known. Closing this gap also for these objectives is a challenging problem for further research.

Another natural question is which approximation factors can be achieved by polynomial-time algorithms. The algorithm we used in this article to prove the upper bounds is not a polynomial-time algorithm because it assumes that for each level k an optimal k -clustering is given. The approximation factors worsen if only approximately optimal clusterings are used instead. It is known that 8-approximate hierarchical clusterings can be computed efficiently with respect to the diameter and (discrete) radius [13]. It is not clear whether or not it is NP-hard to obtain better hierarchical clusterings. The only NP-hardness results come from the problems with given k . Since computing a $(2 - \epsilon)$ -approximation for k -clustering with respect to the diameter and (discrete) radius is NP-hard, this is also true for the hierarchical versions. However, this is obsolete due to our lower bound, which shows that in general there does not even exist a $(2 - \epsilon)$ -approximate hierarchical clustering.

References

- 1 Marcel R. Ackermann, Johannes Blömer, Daniel Kuntze, and Christian Sohler. Analysis of agglomerative clustering. *Algorithmica*, 69(1):184–215, 2014. doi:10.1007/s00453-012-9717-4.
- 2 Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward. Better guarantees for k -means and Euclidean k -median by primal-dual algorithms. In *58th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 61–72, 2017. doi:10.1109/FOCS.2017.15.
- 3 Anna Arutyunova, Anna Großwendt, Heiko Röglin, Melanie Schmidt, and Julian Wargalla. Upper and lower bounds for complete linkage in general metric spaces. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 18:1–18:22, 2021. doi:10.4230/LIPIcs.APPROX/RANDOM.2021.18.

- 4 Anna Arutyunova and Heiko Röglin. The price of hierarchical clustering. *CoRR*, abs/2205.01417, 2022. doi:10.48550/arXiv.2205.01417.
- 5 Felix Bock. Hierarchy cost of hierarchical clusterings. *Journal of Combinatorial Optimization*, 2022. doi:10.1007/s10878-022-00851-4.
- 6 Jaroslaw Byrka, Thomas W. Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k -median and positive correlation in budgeted optimization. *ACM Trans. Algorithms*, 13(2):23:1–23:31, 2017. doi:10.1145/2981561.
- 7 Moses Charikar and Vaggos Chatziafratis. Approximate hierarchical clustering via sparsest cut and spreading metrics. In *Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 841–854, 2017. doi:10.1137/1.9781611974782.53.
- 8 Moses Charikar, Chandra Chekuri, Tomás Feder, and Rajeev Motwani. Incremental clustering and dynamic information retrieval. *SIAM J. Comput.*, 33(6):1417–1440, 2004. doi:10.1137/S0097539702418498.
- 9 Vincent Cohen-Addad, Varun Kanade, Frederik Mallmann-Trenn, and Claire Mathieu. Hierarchical clustering: Objective functions and algorithms. In *Proc. of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 378–397, 2018. doi:10.1137/1.9781611975031.26.
- 10 Wenqiang Dai. A 16-competitive algorithm for hierarchical median problem. *SCIENCE CHINA Information Sciences*, 57(3):1–7, 2014. doi:10.1007/s11432-014-5065-0.
- 11 Aparna Das and Claire Kenyon-Mathieu. On hierarchical diameter-clustering and the supplier problem. *Theory Comput. Syst.*, 45(3):497–511, 2009. doi:10.1007/s00224-009-9186-6.
- 12 Sanjoy Dasgupta. A cost function for similarity-based hierarchical clustering. In *Proc. of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 118–127, 2016. doi:10.1145/2897518.2897527.
- 13 Sanjoy Dasgupta and Philip M. Long. Performance guarantees for hierarchical clustering. *Journal of Computer and System Sciences*, 70(4):555–569, 2005. doi:10.1016/j.jcss.2004.10.006.
- 14 Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985. doi:10.1016/0304-3975(85)90224-5.
- 15 Anna Großwendt and Heiko Röglin. Improved analysis of complete-linkage clustering. *Algorithmica*, 78(4):1131–1150, 2017. doi:10.1007/s00453-017-0284-6.
- 16 Anna Großwendt, Heiko Röglin, and Melanie Schmidt. Analysis of ward’s method. In *Proc. of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2939–2957, 2019. doi:10.1137/1.9781611975482.182.
- 17 Anna-Klara Großwendt. *Theoretical Analysis of Hierarchical Clustering and the Shadow Vertex Algorithm*. PhD thesis, University of Bonn, 2020. URL: <http://hdl.handle.net/20.500.11811/8348>.
- 18 Dorit S. Hochbaum and David B. Shmoys. A unified approach to approximation algorithms for bottleneck problems. *J. ACM*, 33(3):533–550, 1986. doi:10.1145/5925.5933.
- 19 Guolong Lin, Chandrashekar Nagarajan, Rajmohan Rajaraman, and David P. Williamson. A general approach for incremental approximation and hierarchical clustering. *SIAM Journal on Computing*, 39(8):3633–3669, 2010. doi:10.1137/070698257.
- 20 Sakib A. Mondal. An improved approximation algorithm for hierarchical clustering. *Pattern Recognit. Lett.*, 104:23–28, 2018. doi:10.1016/j.patrec.2018.01.015.
- 21 C. Greg Plaxton. Approximation algorithms for hierarchical location problems. *Journal of Computer and System Sciences*, 72(3):425–443, 2006. doi:10.1016/j.jcss.2005.09.004.
- 22 Yuyan Wang and Benjamin Moseley. An objective for hierarchical clustering in euclidean space and its connection to bisecting k -means. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):6307–6314, 2020. doi:10.1609/aaai.v34i04.6099.
- 23 Joe H. Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, 1963. doi:10.1080/01621459.1963.10500845.