



It's Hard to HAC Average Linkage!

MohammadHossein Bateni  

Google Research, New York, NY, USA

Laxman Dhulipala  

University of Maryland, College Park, MD, USA

Kishen N. Gowda  

University of Maryland, College Park, MD, USA

D. Ellis Hershkowitz  

Brown University, Providence, RI, USA

Rajesh Jayaram  

Google Research, New York, NY, USA

Jakub Łącki  

Google Research, New York, NY, USA

Abstract

Average linkage Hierarchical Agglomerative Clustering (HAC) is an extensively studied and applied method for hierarchical clustering. Recent applications to massive datasets have driven significant interest in near-linear-time and efficient parallel algorithms for average linkage HAC.

We provide hardness results that rule out such algorithms. On the sequential side, we establish a runtime lower bound of $n^{3/2-\epsilon}$ on n node graphs for sequential combinatorial algorithms under standard fine-grained complexity assumptions. This essentially matches the best-known running time for average linkage HAC. On the parallel side, we prove that average linkage HAC likely cannot be parallelized even on simple graphs by showing that it is CC-hard on trees of diameter 4. On the possibility side, we demonstrate that average linkage HAC can be efficiently parallelized (i.e., it is in NC) on paths and can be solved in near-linear time when the height of the output cluster hierarchy is small.

2012 ACM Subject Classification Theory of computation → Parallel algorithms; Theory of computation → Streaming, sublinear and near linear time algorithms; Theory of computation → Graph algorithms analysis

Keywords and phrases Clustering, Hierarchical Graph Clustering, HAC, Fine-Grained Complexity, Parallel Algorithms, CC

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.18

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: <https://arxiv.org/abs/2404.14730> [6]

Funding LD and KNG were supported by NSF grant CNS-2317194.

Acknowledgements We thank the anonymous reviewers for their useful comments.

1 Introduction

Hierarchical clustering is a fundamental method for data analysis which organizes data points into a hierarchical structure so that similar points appear closer in the hierarchy. Unlike other common clustering methods, such as k -means, hierarchical clustering does not require the the number of clusters to be fixed ahead of time. This allows it to capture structures that are inherently hierarchical – such as phylogenies [18] and brain structure [11]. One of the most widely used and studied methods for hierarchical clustering is Hierarchical Agglomerative



© MohammadHossein Bateni, Laxman Dhulipala, Kishen N. Gowda, D. Ellis Hershkowitz, Rajesh Jayaram, and Jakub Łącki; licensed under Creative Commons License CC-BY 4.0


51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).

Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;

Article No. 18; pp. 18:1–18:18



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany




18:2 It's Hard to HAC Average Linkage!

Clustering (HAC) [22, 24, 37]. HAC produces a hierarchy by first placing each point in its own cluster and then iteratively merging the two *most similar* clusters until all points are aggregated into a single cluster. The similarity of two clusters is given by a *linkage function*. HAC is included in many popular scientific computing libraries such as scikit-learn [35], SciPy [40], ALGLIB [36], Julia, R, MATLAB, Mathematica and many more [33, 34]. This cluster hierarchy is often equivalently understood as a binary tree – a.k.a. dendrogram – whose internal nodes correspond to cluster merges.

The proliferation of massive datasets with billions of points has driven the need for more efficient HAC algorithms that can overcome the inherent $\Theta(n^2)$ complexity required to read all pairwise distances [16, 17, 29]. Finer-grained running time bounds for HAC were recently obtained by assuming that only $m = o(n^2)$ pairs of points have nonzero similarity, and analyzing the running time as a function of both n and m . This is a natural assumption in practice, as in large datasets of billions of datapoints, typically a small fraction of pairs exhibit nonnegligible similarity. In this case, the input to HAC is an edge-weighted graph, where each vertex represents an input point and each edge weight specifies the similarity between its endpoints. This approach is convenient for large-scale applications since (1) very large clustering instances can be compactly represented as sparse weighted graphs and (2) the running time of HAC can be decoupled from the running time of nearest-neighbor search.

A particularly common linkage function for HAC is *average linkage*, which both optimizes reasonable global objectives [31] and exhibits good empirical performance [5, 12, 20, 23, 27, 29, 30, 32, 44]. Here, the similarity of two clusters is the average edge weight between them (non-present edges are treated as having weight 0). In other words, average linkage HAC repeatedly merges the two clusters with the highest average edge weight between them (see Figure 1 for an example).



■ **Figure 1** An example of average linkage HAC run on an input graph G . Edges labeled with weights. 1a gives G . 1b gives the cluster hierarchy output by HAC. 1c gives the corresponding dendrogram with internal nodes labeled with the weight of their corresponding merge.

A natural algorithmic question then is how quickly can we solve average linkage HAC on n node and m edge graphs? Recent work has provided a partial answer to this basic question in sequential and parallel models of computation. In particular, [15] showed that average linkage HAC can be solved in $\tilde{O}(n\sqrt{m})$ time, thus providing a sub-quadratic time algorithm for sufficiently sparse graphs. A follow-up paper studied average linkage HAC in the parallel setting and showed that the problem is P-complete and so likely does not admit NC algorithms [16]. However, the P-completeness result of [16] holds for worst case graphs whereas typical applications of HAC are on highly structured graphs – namely those which are meant to capture relevant properties of an underlying metric – and so there is still hope for parallelizing average linkage HAC on more structured instances.

In fact, such structured instances of average linkage HAC are known to admit much faster algorithms in the sequential setting: the sequential algorithm of [15] implies that if the input graph is planar (or, more generally, minor-free) average linkage HAC can be solved in time $\tilde{O}(m)$. More generally, if each graph obtained by contracting all clusters at each step of average linkage HAC has $O(1)$ arboricity¹, then it is possible to solve average linkage HAC in time $\tilde{O}(m)$; it follows that average linkage HAC can be solved in sequential time $\tilde{O}(m)$ on trees or planar graphs. In light of these improved sequential results for highly structured graphs, it becomes natural to hope for efficient parallel algorithms on structured graphs such as low arboricity graphs or, even, just trees.

1.1 Our Contributions

In this work, we continue the line of work which studied the computational complexity of different variants of HAC [1, 15, 16, 19, 39] and perform a careful investigation into the complexity of average linkage HAC. In particular, we study HAC on n node and m edge graphs and investigate whether near-linear time algorithms, or more efficient parallel algorithms are possible, namely:

1. **Near-Linear Time Algorithms:** Can we improve over the best known $\tilde{O}(n\sqrt{m})$ upper bound for average linkage HAC and obtain near-linear time sequential algorithms?
2. **NC Algorithms:** are there $\text{polylog}(n)$ depth parallel algorithms for average linkage HAC with $\text{poly}(n)$ work for highly structured instances, e.g., trees, or minor-closed graphs?

We give both new lower bounds which (conditionally) rule out near-linear time and NC algorithms, and provide conditions under which these impossibility results can be bypassed.

First, we demonstrate that near-linear time algorithms are impossible under standard fine-grained complexity assumptions.

► **Theorem 1.** *If average linkage HAC can be solved by a combinatorial algorithm in $O(n^{3/2-\epsilon})$ time for any $\epsilon > 0$, then the Combinatorial Boolean Matrix Multiplication (Combinatorial BMM) Conjecture is false.*

Our reduction also implies a second (weaker) conditional lower bound that also holds for non-combinatorial algorithms (e.g., algebraic algorithms) based on the running time of matrix multiplication. In particular, for two $n \times n$ binary matrices, it is well known that matrix multiplication can be solved in time $O(n^\omega)$ where $2 \leq \omega < 2.3716$ [43]. In this setting, we obtain the following result:

► **Theorem 2.** *If average linkage HAC can be solved by an algorithm in $O(n^{\omega/2-\epsilon})$ time for some $\epsilon > 0$, then boolean matrix multiplication can be solved in $O(n^{\omega-\epsilon'})$ time for some $\epsilon' > 0$.*

Notably, Theorem 1 shows that the prior running time of $\tilde{O}(n\sqrt{m})$ of [15] is *optimal* up to logarithmic factors under standard fine-grained complexity assumptions, at least for graphs consisting of $O(n)$ many edges. We obtain this conditional lower bound by showing that a carefully constructed instance of HAC can be used to solve the triangle detection problem, which is sub-cubically equivalent to Boolean Matrix Multiplication [42]. We obtain a bound of (essentially) $\Omega(n^{3/2})$ since our reduction incurs a quadratic time and space blowup when transforming an input triangle detection instance to an instance of average linkage HAC.

¹ A graph has arboricity at most α if all of its edges can be covered by at most α trees.

We next turn to the parallel setting. Here, we show that HAC – *even on trees* – is unlikely to admit efficient parallel algorithms. More formally, we show that average linkage HAC on low diameter trees is as hard as any problem in the complexity class Comparator Circuit (CC) [13, 38]. It is believed that CC is incomparable with NC and that CC-hardness is evidence that a problem is not parallelizable [13, 28].

► **Theorem 3.** *Average linkage HAC is CC-hard, even on trees of diameter 4.*

We note that it is known that $CC \subseteq P$ and so the P-hardness of [16] already suggests the impossibility of efficient parallel algorithms on *general graphs*. However, our result suggests the impossibility of efficient parallel algorithms *even on very simple graphs* (trees of diameter 4). We obtain this result by reducing from the lexicographically first maximal matching (LFM Matching) problem and an intermediate problem which we call Adaptive Minimum, which captures some of what makes HAC intrinsically difficult to parallelize.

On the positive side, we demonstrate that average linkage HAC on path graphs is in NC, under the mild assumption that the aspect ratio is polynomial. While the class of path graphs is restrictive, even on paths average linkage is highly non-trivial and naively running HAC requires resolving chains of $\Omega(n)$ sequential dependencies. For example, consider a path of vertices (v_1, v_2, \dots, v_n) where the edge $\{v_i, v_{i+1}\}$ has weight $1 + i \cdot \epsilon$ for some small $\epsilon > 0$ and initially each vertex is in its own cluster. Initially, v_n 's most similar neighbor is v_{n-1} and so v_n would like to merge with v_{n-1} but v_{n-1} 's most similar neighbor is v_{n-2} and so on. Thus, whether or not v_n gets to merge with v_{n-1} depends on the merge behavior of $\Theta(n)$ other clusters and so it is not at all clear that NC algorithms should be possible for this setting. Nonetheless, we show the following.

► **Theorem 4.** *Average linkage HAC on paths is in NC. In particular, there is an algorithm for average linkage HAC that runs in $O(\log^2 n \log \log n)$ depth with $O(n \log n \log \log n)$ work.*

The above algorithm leverages the fact that in average linkage HAC the maximum edge similarity monotonically decreases. In particular, it works in $O(\log n)$ phases where each phase consists of merges of equal similarity up to constants. The goal then becomes to efficiently perform merges until every edge is no longer within a constant of the starting maximum similarity of the phase. The starting point of the algorithm is to observe that $\Omega(n)$ sequential dependencies of clusters of equal size can be resolved efficiently in parallel in a phase by noting that in this phase only the odd-indexed edges merge in the chain. Thus, each edge can decide if it is odd-indexed in parallel by, e.g., using `prefix-sum`, which is well known to be solvable in linear work in NC.

For chains with clusters of general weights, we decompose dependency chains into short ($O(\log n)$ -length) subchains where resolving dependencies within the subchain must be done sequentially but in the current phase each subchain's merge behavior only depends on whether or not its closest neighboring subchains merges into it or not. Thus, each subchain can compute its merge behavior for these two cases and then, similar to the equal weights setting, we propagate merge behavior across subchains efficiently in parallel.

To complement our sequential lower bound with a positive result, we demonstrate that it is possible to achieve near-linear running time, provided the dendrogram has low height. Thus, if the output dendrogram is a relatively balanced tree, then near-linear time algorithms are possible.

► **Theorem 5.** *There is an implementation of the nearest-neighbor chain algorithm for average linkage HAC that runs in $O(m \cdot h \log n)$ time where h is the height of the output dendrogram.*


The above result is in fact obtained by a relatively simple (but to the best of our knowledge, new) analysis of *existing* classic HAC algorithms. In particular, we show that the nearest-neighbor chain [7, 21] and heap-based algorithms [27] for HAC, which were developed over 40 years ago achieve this bound. Due to space constraints, we prove Theorem 5 in the full version [6].

2 Preliminaries

The input to the HAC algorithm is an undirected weighted graph $G = (V, E, w)$, where $w : V \times V \rightarrow \mathbb{R}_+ \cup \{0\}$ is a function assigning nonnegative weights to the edges. For convenience we assume $w(x, y) = 0$ when $xy \notin E$. The vanilla version of average linkage HAC is given as Algorithm 1. It starts by putting each vertex in a cluster of size 1 and then repeats the following step. While there is a pair of clusters of positive similarity, find two most similar clusters and merge them together, that is, replace them by their union. The similarity between two clusters is the total edge weight between them divided by the product of the cluster sizes. We refer to this version as the *static graph* version, since the graph is not changed throughout the run of the algorithm.

Throughout the paper we usually work with a different (equivalent) way of presenting the same algorithm which is given as Algorithm 2. In this version we maintain a graph G whose vertices are clusters. The *size* of the vertex is the size of the cluster it represents. The *normalized* weight of an edge xy in G is $w(x, y)$ divided by the product of the sizes of x and y .

Whenever two clusters merge, their corresponding vertices are merged into one, i.e., the edge between them is contracted and the size of the new vertex is the sum of the sizes of the vertices that merged. In the following we sometimes say that a vertex x merges into vertex y . In this case we simply assume that the name of the resulting vertex is y and the size of y is increased by the size of x . See Figure 2.



■ **Figure 2** An example of average linkage HAC run on an input graph G where we imagine we contract merged clusters. Intermediate vertices labeled with the vertices of G their corresponding cluster contains. Edges labeled with their weight and next merged edge is dashed.

The output of HAC is a *dendrogram* – a rooted binary tree representing the cluster merges performed by the algorithm. Every node of the dendrogram is a cluster built by the algorithm. There are exactly $|V|$ leaves corresponding to the single-element clusters that are formed in the beginning of the algorithm. Whenever two clusters C_1 and C_2 are merged, we add to the dendrogram a new node $C_1 \cup C_2$ whose children are C_1 and C_2 . See Figure 1c for the dendrogram of Figure 2.

We use the classic multithreaded model [4, 9, 10] (formally, the MP-RAM [9]) to analyze the parallel algorithms. We assume a set of threads that share the memory. Each thread acts like a sequential RAM plus a `fork` instruction that forks two new child threads. When a

■ **Algorithm 1** Average linkage HAC – static graph version.

Input: $G = (V, E, w)$

- 1 **Function** $\text{Similarity}(C_1, C_2, w)$:
- 2 \lfloor **return** $\sum_{x \in C_1, y \in C_2} w(x, y) / (|C_1| \cdot |C_2|)$
- 3 **Function** $\text{HAC}(G)$:
- 4 $\mathcal{C} \leftarrow$ clustering where each vertex of G is in a separate cluster
- 5 **while** $\exists_{C_1, C_2 \in \mathcal{C}} s.t. C_1 \neq C_2$ and $\text{Similarity}(C_1, C_2, w) > 0$ **do**
- 6 $(C_1, C_2) = \arg \max_{(C_1, C_2) \in \mathcal{C} \times \mathcal{C}} \text{Similarity}(C_1, C_2, w)$
- 7 $\mathcal{C} := (\mathcal{C} \setminus \{C_1, C_2\}) \cup \{C_1 \cup C_2\}$.

■ **Algorithm 2** Average linkage HAC – graph contraction version.

Input: $G = (V, E, w)$

- 1 **Function** $\text{Similarity}(x, y, w, S)$:
- 2 \lfloor **return** $w(x, y) / (S(x) \cdot S(y))$
- 3 **Function** $\text{HAC}(G)$:
- 4 $S :=$ a function mapping each element of V to 1
- 5 **while** $\exists_{xy \in E} s.t. \text{Similarity}(x, y, w, S) > 0$ **do**
- 6 $xy = \arg \max_{xy \in E} \text{Similarity}(x, y, w, S)$
- 7 Contract x with y in G creating a vertex z . The parallel edges that are created are merged into a single edge whose weight is the sum of the merged edge weights. Any resulting self-loops are removed.
- 8 Set $S(z) := S(x) + S(y)$

thread performs a fork, the two child threads can both start by running their next instructions, and the original thread is suspended until both children terminate. A computation starts with a single root thread and finishes when that root thread finishes. A parallel for-loop can be viewed as executing forks for a logarithmic number of levels. A computation can thus be viewed as a DAG (directed acyclic graph). We say the *work* is the total number of operations in this DAG and *span* (*depth*) is equal to the longest path in the DAG. We note that computations in this model can be cross-simulated in standard variants of the PRAM model in the same work (asymptotically), and losing at most a single logarithmic factor in the depth [9].

3 An $\Omega(n^{3/2-\epsilon})$ Conditional Lower Bound for Average Linkage HAC

In this section, we show an $\Omega(n^{3/2-\epsilon})$ conditional lower bound on the time required to solve average linkage HAC on general weighted graphs. Specifically, we show this lower bound assuming the Combinatorial Boolean Matrix Multiplication (BMM) conjecture, a central conjecture in fine-grained complexity about the time required to multiply two $n \times n$ boolean matrices [2, 42].

► **Conjecture 6** (Combinatorial BMM). *Combinatorial algorithms cannot solve Boolean Matrix Multiplication in time $O(n^{3-\epsilon})$ for $\epsilon > 0$.*

We refer to [2] for an in-depth discussion of the somewhat informal notion of “combinatorial” algorithms and more on Conjecture 6 and its history.

In this work we will make use of an equivalent characterization of the BMM conjecture due to [42]. Specifically, [42] shows that the BMM problem is sub-cubically equivalent to the Triangle Detection problem: the problem of deciding whether or not an input graph G contains a triangle (i.e., cycle of 3 vertices). The following summarizes this result.

► **Theorem 7** (Theorem 1.3 of [42]). *Combinatorial algorithms cannot solve Triangle Detection in time $O(n^{3-\epsilon})$ for $\epsilon > 0$ unless Conjecture 6 is false.*

Thus, we give a reduction from Triangle Detection to average linkage HAC. Our reduction will quadratically increase the number of vertices of the input Triangle Detection instance, and therefore give an $\Omega(n^{3/2-\epsilon})$ lower bound for average linkage HAC. In the rest of this section, we show the following quadratic-blowup reduction from Triangle Detection to average linkage HAC.

► **Theorem 8.** *Given a Triangle Detection instance on graph G with t vertices and m edges, there is a reduction that runs in $O(t^2)$ time and constructs an instance of average linkage HAC on graph G' with $t + t^2$ vertices and $t^2 + m$ edges. Furthermore, given the sequence of merges performed by average linkage HAC on G' , we can solve Triangle Detection on G in time $O(t^2)$.*

As a corollary of this reduction and Theorem 7, we obtain the following conditional lower-bound on the running time of HAC.

► **Theorem 1.** *If average linkage HAC can be solved by a combinatorial algorithm in $O(n^{3/2-\epsilon})$ time for any $\epsilon > 0$, then the Combinatorial Boolean Matrix Multiplication (Combinatorial BMM) Conjecture is false.*

As a second corollary, we obtain a conditional lower-bound in terms of the optimal running time of matrix multiplication for two $n \times n$ binary matrices. Matrix multiplication can be solved in time $O(n^\omega)$ where $2 \leq \omega < 2.3716$ [43]. An extensive line of research on matrix multiplication over the past thirty years has only improved ω from 2.376 to 2.3716, with the current state-of-the-art being due to a very recent result of Williams et al. [43] (for a subset of the historical advances in this area see, e.g., [3, 14, 26, 41]). The fastest known algorithm for triangle detection works by simply reducing the problem to matrix multiplication and therefore runs in $O(n^\omega)$ time. Surprisingly, despite triangle detection only returning a single bit (whether a triangle exists or not in G), the problem can be used to give a sub-cubic reduction for boolean matrix multiplication (where the output is n^2 bits). In particular, an algorithm for triangle detection running in time $O(n^{3-\delta})$ for some $\delta > 0$ yields an algorithm for matrix multiplication in time $O(n^{3-\delta/3})$ [42]. Using this fact, we can derive a conditional lower bound based on the value of ω .

► **Theorem 2.** *If average linkage HAC can be solved by an algorithm in $O(n^{\omega/2-\epsilon})$ time for some $\epsilon > 0$, then boolean matrix multiplication can be solved in $O(n^{\omega-\epsilon'})$ time for some $\epsilon' > 0$.*

An interesting open question is whether there are faster non-combinatorial algorithms that can leverage fast matrix multiplication or Strassen-like techniques and improve over the $\Omega(n^{3/2-\epsilon})$ barrier for combinatorial algorithms for average linkage HAC.

3.1 Reduction

We now prove Theorem 8 by giving a quadratic-time reduction from triangle detection to average linkage HAC. The reduction is loosely inspired by a recent lower-bound result for multidimensional range queries [25]. The input to the reduction is an unweighted graph G

18:8 It's Hard to HAC Average Linkage!

on t vertices with m edges; the problem is to detect whether G has a triangle. To do this, we will construct a HAC instance on an edge-weighted graph G' with $t + t^2$ vertices and $t^2 + m$ edges. We will show that the specific way in which an exact HAC algorithm merges the edges in this instance reveals whether or not G has a triangle.

Constructing G'

Let $N_G(v)$ denote the neighbors of a vertex $v \in G$ (note that $v \notin N_G(v)$). We define G' as follows. We start by adding all vertices and edges from G , that is the t vertices v_1, \dots, v_t from G , including all of their incident edges $N_G(v_i)$. We call these the *core* vertices. The initial weight of the edges between any two core vertices is set to 1.

In addition to the core vertices, we add an additional t^2 leaf vertices that we connect to the core vertices with specific edge weights. We add the t^2 leaf vertices over a sequence of t rounds where the i -th round connects one new leaf vertex to every core vertex. The weights to the newly added leaves depend on the neighbors of the node v_i in the original graph G , and are set as follows:

- (1) A core vertex v_j is connected to its new leaf with an edge of weight $(1/i) - \epsilon$ if $v_j \in N_G(v_i)$.
 - (2) A core vertex v_j is connected to its new leaf with an edge of weight $(1/i) + \epsilon$ if $v_j \notin N_G(v_i)$.
- See Figure 3 for an illustration of our reduction.

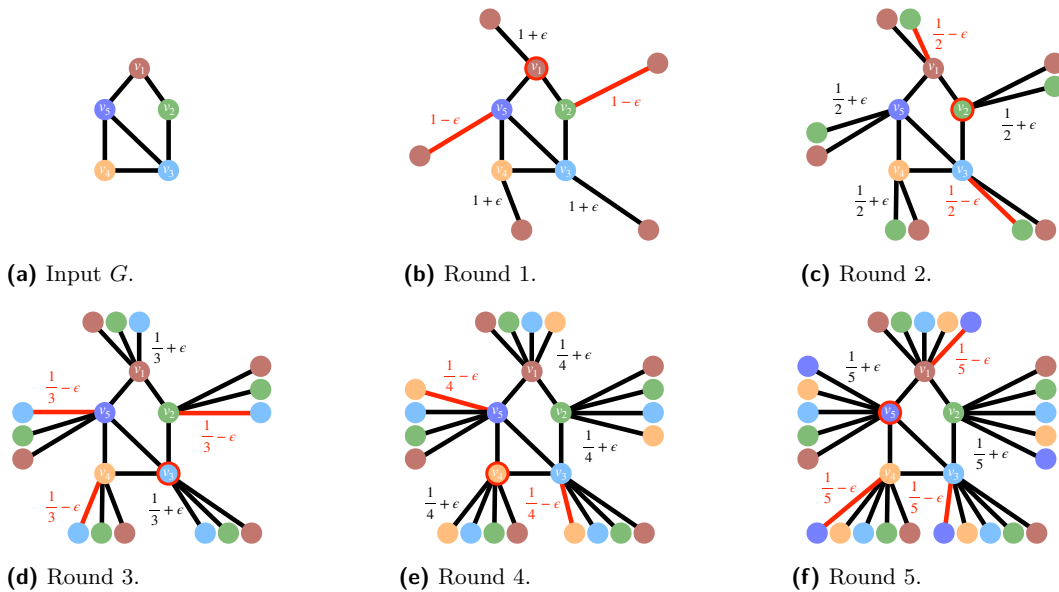


Figure 3 Our triangle detection reduction where we compute G' from G by adding $t = 5$ nodes over t rounds. 3f gives G' . Each node labeled according to its round and corresponding vertex in G . Edges labelled with their weight in the round they are added (edges of G have weight 1). For the i th round we highlight in red v_i and the edges added with weight $1/i - \epsilon$.

Running HAC on G'

Having defined G' , let us consider the merges that the exact HAC algorithm will make on this instance. In this section, for succinctness we use *weight* to refer to the normalized (i.e., average linkage) weight. HAC will begin by merging the maximum weight edges. The maximum weight initially depends on the structure of $N_G(v_1)$. First, all core vertices that

are not in $N_G(v_1)$ will merge with their round 1 leaves (since these edges have weight $1 + \epsilon$) increasing their cluster size to 2. This leaves all core vertices in $N_G(v_1)$. Any edge in G' between two such core vertices will have weight 1, and will be merged next. Crucially, any merge in this round with a weight of 1 indicates a triangle incident on v_1 since the two core vertex endpoints of the edge must be contained in $N_G(v_1)$, hence connected by edges to v_1 . If no edges of weight 1 merge, the remaining leaves that we added to core vertices in $N_G(v_1)$ merge into their neighboring core vertex.

Assuming we did not merge any weight 1 edges in the first round, at this point the cluster size of each core vertex is 2, and so the weight of any edge originally in G (between two core vertices) will be $1/4$. The edge weights to leaves in round 2 will be $((1/2) \pm \epsilon)/2 = 1/4 \pm (\epsilon/2)$, which is larger than $1/4$ for edges of Type 2. Therefore, the same argument for how edges merge in round 1 can be inductively applied to the next round. The edge weights in round i for edges between core vertices will be $1/i^2$, and by the same argument as before, the Type 2 (Type 1) edges will be larger (smaller) by ϵ/i . As a result of how an exact average linkage HAC will merge the edges of G' , we obtain the following lemma:

► **Lemma 9.** *Consider the sequence of merges performed by the HAC algorithm on G' . If the merge sequence consists of t^2 merges, which first merge all leaf vertices, and only then makes merges between core vertices, then G does not contain any triangles. If the merge sequence merges any edge between two core vertices in the first t^2 merges, then G contains a triangle.*

Completing the Reduction

We will now complete the proof of Theorem 8. Suppose we are given an instance of TRIANGLE DETECTION on n vertices. Conjecture 6 implies that this instance cannot be solved by combinatorial algorithms in $O(n^{3-\epsilon})$ time for any $\epsilon > 0$.

Let the time complexity of HAC on a graph G with n vertices and m edges be $T_{\text{HAC}}(n, m)$. Suppose HAC can be solved combinatorially in $O(n^{3/2-\epsilon})$ time. Given a TRIANGLE DETECTION instance on n vertices we create a graph G' with $O(n^2)$ vertices and $O(n^2)$ edges, and run HAC on G' . The running time of the reduction is $O(n^2)$, and the running time of HAC on G' is $O((n^{2 \cdot (3/2-\epsilon)}) = O(n^{3-2\epsilon})$, which will falsify Conjecture 6 by Theorem 7. Thus, conditional on Conjecture 6, there is no algorithm for HAC running in time $T_{\text{HAC}}(n, m) = O(n^{3/2-\epsilon})$ for any constant $\epsilon > 0$, completing the proof of Theorems 8 and 1. The same argument, under the assumption that triangle detection cannot be solved in $O(n^{\omega-\epsilon})$ time for any constant $\epsilon > 0$ implies Theorem 2.

4 Average Linkage HAC is Hard to Parallelize Even on Trees

In this section we prove that average linkage HAC is likely hard to parallelize by showing it is CC-hard even on low depth trees. We begin with some preliminaries. The formal definition of CC-hardness we will use is as follows.

► **Definition 10 (CC-Hard).** *A problem is CC-hard if all problems of CC are logspace-reducible to it.*

For our purposes we will not need to define the class CC. Rather, we only need the above definition of CC-hardness and a single CC-hard problem, LFM Matching. Recall that a matching of a graph $G = (V, E)$ is a subset of edges $M \subseteq E$ if each vertex is incident to at most one edge of M . A matching is said to be maximal if each $e = \{u, v\} \notin M$ satisfies the property that either u or v is incident to an edge of M . The greedy algorithm for maximal matching initializes M as \emptyset and then simply iterates over the edges of E in some order and adds the current edge e to M if the result of doing so is a matching.

18:10 It's Hard to HAC Average Linkage!

► **Problem 1 (LFM Matching).** An instance of lexicographically first maximal matching (LFM Matching) consists of a bipartite graph $G = (V = L \sqcup R, E)$ with vertices ordered as $L = (l_0, l_1, \dots, l_{n-1})$ and $R = (r_0, r_1, \dots, r_{n-1})$. The lexicographically first maximal matching is the matching obtained by running the greedy algorithm for maximal matching on edges ordered first by their endpoint in L and then by their endpoint in R . That is, in this ordering $e = \{l_i, r_j\}$ precedes $e' = \{l_{i'}, r_{j'}\}$ iff (1) $i < i'$ or (2) $i = i'$ and $j < j'$. Our goal is to decide if a designated input edge is in the LFM Matching.

The following summarizes known hardness of LFM Matching.

► **Theorem 11** ([28, 38]). LFM Matching is CC-hard.


Next, we introduce the search variant of the HAC problem whose CC-hardness we will prove.

► **Problem 2 (Average Linkage HAC).** An instance of Average Linkage HAC consists an undirected graph $G = (V, E)$, along with edge weights $w : E \rightarrow \mathbb{R}_{\geq 0}$. Consider the sequence $(C_1, C'_1), (C_2, C'_2), \dots$ of cluster merges produced by the procedure HAC(G) from Algorithm 1. Given any pair of vertices $u, v \in G$, the goal of the Average Linkage HAC problem is to output the index i such that u, v first merge together at step i , namely $u \in C_i$ and $v \in C'_i$ (or $u \in C'_i$ and $v \in C_i$).

To prove the hardness of Average Linkage HAC, we will first prove the hardness of an intermediate problem, called Adaptive Minimum. The construction of the Adaptive Minimum problem will be more amenable to our reductions, and therefore simplify the following exposition. See Figure 4 for an illustration of Adaptive Minimum.

► **Problem 3 (Adaptive Minimum).** An instance of Adaptive Minimum consists of a (0-based indexed) $n \times n$ matrix A where each row contains a permutation of $\{0, \dots, n-1\}$ and some index $x \in [0, n)$. The goal is to simulate the following algorithm. Start with $I = \{0, \dots, n-1\}$ and execute the following steps for $i = 0, \dots, x$:

1. Let $k_i = \arg \min_{j \in I} A[i, j]$.
 2. Set $I := I \setminus \{k_i\}$.
- Our goal is to compute k_x .



■ **Figure 4** Adaptive Minimum on matrix A . The row considered in each step is shown in blue. k_i for the i -th row written to the right of A in green with witnessing entry of A also in green. Indices removed from I in relevant rows crossed out in red.

Observe that both problems 2 and 3 are defined as having an algorithm output an index $i \in \{1, 2, \dots\}$. Thus, these can be considered search problems instead of decision problems. We choose to work with the search versions of these problems for simplicity of our reductions, however, our reduction naturally extends to the decision variants (e.g., where the algorithm is given $u, v \in V$ and an index i and asked if u, v merge on step i).

We first prove the CC-hardness of this intermediate problem. See Figure 5 for an illustration of our reduction from LFM Matching to Adaptive Minimum.

► **Lemma 12.** *Adaptive Minimum is CC-hard.*

Proof. By Theorem 11 and the definition of CC-hardness (Definition 10), it suffices to argue that LFM Matching (Problem 1) is logspace reducible to Adaptive Minimum (Problem 3). We begin by describing our reduction and then observe that it only requires logarithmic space. The basic idea of the reduction is to associate with each vertex on the left side of our instance of LFM Matching a row of the matrix of Adaptive Minimum and each vertex on the right side of LFM Matching a column of the matrix of Adaptive Minimum.

More formally, consider an instance of LFM Matching on graph $G = (V = L \sqcup R, E)$ where our goal is to decide if a given edge e is in the LFM matching. We consider the following instance of Adaptive Minimum to solve this on a $2n \times 2n$ size matrix A . We will refer to an index i as a *dummy index* if $i \geq n - 1$.

Consider a vertex $l_i \in L$ connected to vertices $R_i \subseteq R$ in G , for some $i \leq n - 1$. We will construct the i th row of A to correspond to a permutation π_i that first gives the indices of all neighbors of l_i in R sorted according to the ordering of R then gives all dummy indices then gives the indices of non-neighbors of l_i in R . Specifically, the first $|R_i|$ indices of π_i will be the indices of R_i (sorted by their order in R), the next n indices will be dummy indices $n, n + 1, \dots, 2n - 1$ and the remaining $n - |R_i|$ indices will be the indices of vertices in $R \setminus R_i$ (sorted, say, by their order in R). For $i \geq n$ we can construct our permutation arbitrarily. Lastly, let x (the index for which we would like to compute k_x in our instance of Adaptive Minimum) be the index of the endpoint of e in L . Once Adaptive Minimum computes k_x , we verify whether or not it corresponds to the endpoint of e in R to determine the final output of the LFM Matching instance. Again, see Figure 5.

We now argue correctness of the reduction. A straightforward proof by induction on i demonstrates that at the beginning of the i th round of the Adaptive Minimum algorithm we have that I consists of at least $n - i$ dummy indices and $j < n$ is not in I only if r_j is in the LFM Matching and is matched to some $l_{i'}$ for $i' < i$. It follows that $e = (l_i, r_j)$ is in the LFM Matching iff $k_i = j$, showing correctness of our reduction.

It remains to show that the above reduction can be done with logspace. In order to do so, we must argue that $A[i, j]$ can be computed with logspace for every i and j . Doing so is trivial if i is a dummy index, so consider $i < n$.

- If $j \leq |R_i|$ then $A[i, j]$ is just the index of the j th vertex of R_i (i.e., neighbor of l_i) in the ordering given by R .
- If $j \in [|R_i|, |R_i| + n]$ then $A[i, j]$ just is $j - |R_i| + n$.
- If $j > |R_i| + n$ then $A[i, j]$ is the index of the $(j - |R_i| - n)$ th vertex in $R \setminus R_i$ (i.e., non-neighbors of l_i) when vertices of $R \setminus R_i$ are sorted according to the ordering on R .

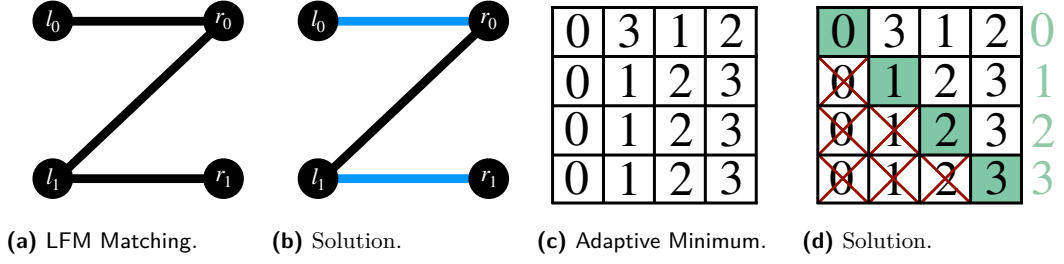
All three of the above quantities can easily be computed in logspace. ◀

Concluding, we use the CC-hardness of Adaptive Minimum to prove the CC-hardness of Average Linkage HAC.

► **Theorem 3.** *Average linkage HAC is CC-hard, even on trees of diameter 4.*

Proof. Our reduction shows how to reduce an instance of Adaptive Minimum (Problem 3) of size n to an instance of average linkage HAC on a tree. We build a rooted tree, in which each root-to-leaf path has length 2 (i.e., the tree has depth 2). We call the neighbors of the root *internal* nodes. Observe that each node is either the root, an internal node or a leaf. The fact that the tree is rooted is only for the convenience of the description. In the construction, we will begin by assigning each node an initial size (see the definition of *size* in Section 2) which is possibly larger than 1 (but at most $\text{poly}(n)$). We will later show how to remove these variable sizes, and reduce to the case where all nodes have initial size 1 (as in the original definition of HAC).

18:12 It's Hard to HAC Average Linkage!



■ **Figure 5** Reduction from LFM Matching to Adaptive Minimum. 5a gives the LFM Matching instance and 5b its solution. 5c gives the Adaptive Minimum instance from the reduction and 5d its solution.

The basic idea of our construction is as follows. Our HAC instance will consist of a rooted tree where each child of the root corresponds to a column of A in our Adaptive Minimum instance. HAC merges will then happen in phases where each phase corresponds to a row of A . In a given phase, exactly one internal node will merge with the root which will correspond to this internal node's column being minimum for the corresponding row of A . In order to guarantee this, each child of the root will have its own carefully selected children such that merging with these children guarantees the desired behavior in every phase.

More formally, the tree is constructed as follows. The root r of the tree has initial size n^8 . It has n children, each of initial size n^4 – we denote them by v_0, \dots, v_{n-1} . The root is connected to its children using edges of weight 1, i.e., $w(r, v_i) = 1$ for all $i = 0, 1, \dots, n-1$ (thus, the normalized weight of the edges $\{rv_i\}_{i=0}^{n-1}$ are each $\frac{1}{n^{12}}$ at the start). Each internal node has $n(n+1)$ children (leaf nodes) grouped into n groups of $n+1$ leaves each. We write $C_{i,j} = \{v_{i,j,0}, v_{i,j,1}, \dots, v_{i,j,n}\}$ to denote the j -th group of children of the i -th internal vertex. All the leaves $v_{i,j,k}$ have initial size 1. Thus, the vertices in the full graph in our construction consists of the root r , internal nodes $\{v_0, v_1, \dots, v_{n-1}\}$ and leaves $\cup_{i=0}^{n-1} \cup_{j=0}^{n-1} \cup_{k=0}^n \{v_{i,j,k}\}$

Let $r_i = n^8 + i \cdot n^4$ (for $0 \leq i < n$). For each pair of an internal node and each of its groups there are only two distinct edge weights for edges between the internal node and the leaves in the group. Specifically, for an internal node v_j and group $C_{j,i}$ of its children we have $A[i, j] + 1$ edges of weight $\frac{1}{r_i - 1}$ and $n - A[i, j]$ edges of weight $\frac{1}{r_i + i \cdot n^3}$. Specifically, we set

$$w(v_j, v_{j,i,k}) = \begin{cases} \frac{1}{r_i - 1} & \text{if } 0 \leq k \leq A[i, j] \\ \frac{1}{r_i + i \cdot n^3} & \text{if } A[i, j] < k \leq n \end{cases}$$

We call the two weights *high-weight* and *low-weight* edges, respectively. Note that their normalized weights are $\frac{1}{n^4(r_i - 1)}$ and $\frac{1}{n^4(r_i + i \cdot n^3)}$ respectively. Observe that the setting of *high-weight* and *low-weight* edges is independent of the internal node v_j , although the number of high versus low-weight edges depends on $A[i, j]$. Moreover, note that even low-weight edges of any group $C_{j,i}$ have higher weights than high-weight edges of group $C_{j,i+1}$:

$$\begin{aligned} \frac{1}{n^4(r_i + i \cdot n^3)} &> \frac{1}{n^4(r_{i+1} + (i+1) \cdot n^3)} \iff \\ \frac{1}{n^8 + i \cdot n^4 + i \cdot n^3} &> \frac{1}{n^8 + (i+1)n^4 - 1} \iff \\ i \cdot n^3 &< n^4 - 1. \end{aligned}$$

which follows from the fact that $i \leq n - 1$. We now demonstrate that the average linkage HAC on this instance works in n phases numbered from 0 to $n - 1$, where in each phase i ,

- $n - 1$ internal nodes contract all of their incident group i edges, and
- one internal node contracts all of its high-weight edges to group i , after which it merges with the root.

Because of the internal node merging with the root, the root has incident leaves, but they are connected with edges of (normalized) edge weights $\leq \frac{1}{n^{15}}$ and so they will be irrelevant until all phases have been completed. We show that if we denote by k_i the index of the internal node which merges with the root in phase i , the sequence k_0, \dots, k_{n-1} is a correct solution to the Adaptive Minimum problem.

In order to analyze the algorithm, we prove the following claim. For convenience, let us define $w_i = n^4 + i \cdot (n + 1)$.

▷ **Claim 13.** In the beginning of phase i , the graph is as follows:

1. The size of the root node is $r_i + i^2 \cdot \Delta_i$ for some $\Delta_i \in [0, n + 1]$.
2. Exactly i internal nodes have been merged with the root, and the corresponding values k_0, \dots, k_{i-1} have been computed correctly.
3. The size of each of the $n - i$ remaining internal nodes is w_i .
4. For all remaining internal nodes, *all* leaves in groups $0, \dots, i - 1$ have been merged into their parents, and no leaves in groups $i, \dots, n - 1$ have been merged.
5. The root may have incident leaves (resulting from internal nodes contracting into it) connected to the root with edges of normalized weights $\leq \frac{1}{n^{15}}$.

Proof. We prove the above claim using induction on i . The base case of $i = 0$ follows directly from how the tree is constructed.

We now simulate a single phase. The edges between the root and the internal nodes have (normalized) weights $\frac{1}{w_i(r_i + i^2 \cdot \Delta_i)} > \frac{1}{n^{15}}$. Hence, the additional leaf nodes incident to the root (see Item 5 of the Claim) are irrelevant. Thus, the highest weight edge in the graph is surely incident to one of the internal nodes. Observe that the relative order of edge weights between an internal node v and its children does not change as the leaves are merged into v . Therefore, given that groups $0, \dots, i - 1$ do not exist anymore, among edges between the internal nodes and leaves, the edges of group i have the highest weights. In the beginning of a phase the high-weight edges in that group have normalized weights $\frac{1}{w_i(r_i - 1)}$ and the low-weight edges have weight $\frac{1}{w_i(r_i + i \cdot n^2)}$.

Hence, we have that if we sort the edges by their normalized weights, the top 3 classes of edges are, starting from the highest weight:

1. High-weight edges between internal nodes and leaves of group i .
2. Edges between the root and the internal nodes.
3. Low-weight edges between internal nodes and leaves of group i .

We will show that the phase consists of the following sub-phases.

1. First, there is some number of subphases, where each of $n - i$ internal nodes contract one incident high-weight edge.
2. Then, there is exactly one subphase, where $n - i - 1$ nodes contract an incident high-weight edge and one internal node merges with the root.
3. Then, the remaining $n - i - 1$ internal nodes merge with all of their group i leaves (we do not analyze the order in this subphase, as it is irrelevant).

18:14 It's Hard to HAC Average Linkage!

Assume that each internal node has at least one high-weight edge in group i . Then, the algorithm will execute a Type 1 subphase: the first $n - i$ steps of the algorithm would merge exactly one high-weight edge incident to each internal node. Note that when an edge incident to an internal node v merges, the weight of v increases, and so the incident edge weights decrease. This guarantees that in the considered $n - i$ steps exactly one merge per internal node happens.

Type 1 subphases of $n - i$ steps continue as long as each internal node has at least one high-weight group i edge in the beginning of the subphase. Each Type 1 subphase also causes the weight of each internal node to increase by 1. Clearly, since nodes are being merged into internal nodes, the ordering of edge weights incident to any internal node does not change.

At some point, in the beginning of a subphase there is an internal node that does not have any incident high-weight edge in group i . Assume that this happened after p Type 1 subphases have completed. Thus, the size of each internal node is $w_i + p$. Since by the construction each internal node had a different number of high-weight edges in group i , there is exactly one node v with no high-weight incident edges and that node merges with the root. This is when Type 2 subphase happens. First, $n - i - 1$ internal nodes contract with a high-weight incident group i edge. At this point the edge weights are as follows. The weight of an edge between v and the root is $\frac{1}{(w_i+p)(r_i+i^2\cdot\Delta_i)}$ and the weight of a high-weight edge in group i is $\frac{1}{(w_i+p+1)(r_i-1)}$. We have that the former is larger since

$$\begin{aligned} (w_i + p)(r_i + i^2 \cdot \Delta_i) &< (w_i + p + 1)(r_i - 1) \iff \\ (w_i + p) \cdot i^2 \cdot \Delta_i &< r_i - 1 \iff \\ (n^4 + i \cdot (n + 1) + p) \cdot i^2 \cdot \Delta_i &< n^8 + i \cdot n^4 - 1 \iff \\ (n^4 + n \cdot (n + 1) + n) \cdot n \cdot n^2 &< n^8 + i \cdot n^3 - 1. \end{aligned}$$

Thus, the internal node with no incident high-weight edges in group i merges with the root. Observe that this is exactly the internal node which had the lowest number of high-weight edges in group i among all remaining internal nodes. This immediately implies that k_i is computed correctly, proving Item 2.

We now show that in the remaining part of the phase the $n - i - 1$ remaining internal nodes contract their incident group i edges. First, observe that the new size of the root node is

$$r_i + w_i + p + i^2 \cdot \Delta_i = (n^8 + i \cdot n^4 + n^4) + (i \cdot (n + 1) + i^2 \cdot \Delta_i + p) = r_{i+1} + (i + 1)^2 \Delta_{i+1}$$

for some $\Delta_{i+1} \in [0, n + 1]$. Note that we use the fact that both Δ_i and p are upper bounded by $n + 1$, which implies $i \cdot (n + 1) + i^2 \cdot \Delta_i + p \leq (i + 1)^2(n + 1)$. This proves Item 1. Thus for an internal node of weight w , the weight of its edge to the root is

$$\frac{1}{w \cdot (r_{i+1} + (i + 1)^2 \Delta_{i+1})} \leq \frac{1}{w \cdot r_{i+1}} = \frac{1}{w \cdot (n^8 + (i + 1) \cdot n^4)}.$$

On the other hand, its low-weight edges to group i leaves have weight

$$\frac{1}{w(n^8 + i(n^4 + n^2))}.$$

As a result, in the remaining part of the current phase all internal nodes will contract all their incident group i edges. This implies Item 4. Thus, the size of each internal node within the phase increases to $w_i + (n + 1) = n^4 + i \cdot (n + 1) + n + 1 = n^4 + (i + 1)(n + 1)$, as required. This proves Item 3 and completes the proof. \triangleleft

Finally, we now claim that, given an instance of Adaptive Minimum with input index $x \in [0, n)$ and an algorithm which can compute the solutions to Problem 2, we can compute the solution k_x to Problem 3 in logspace. To see this, note that it suffices to determine the value k_x as defined above given an algorithm for Average Linkage HAC. To see this, note that for any internal node i , we can query the Average Linkage HAC algorithm to determine which time step t_i it merged with the root. This does not directly tell us which phase i merged with the root, but for a given i we can determine if it merged in phase x by comparing t_i with t_j for all $j \in \{0, 1, \dots, n-1\} \setminus \{i\}$, and checking if there are exactly $x-1$ values of t_j smaller than t_i . This clearly can be verified in log-space. Repeating for all $i \in \{0, 1, \dots, n-1\}$ allows us to correctly determine the identity of the internal node that merged with the root in phase x , and therefore the value of k_x , in logspace as required.

To complete the proof of the Lemma it remains to show how to drop the assumption on the node sizes being initially not all equal to 1. In order to obtain a node of size w it suffices to create a node of weight 1 and initially connect it to $w-1$ auxiliary nodes using very high weight edges. This will force the algorithm to merge all these auxiliary nodes and increase the size of that node to w . Since the auxiliary leaves are connected only to the root and internal nodes (the leaves in our construction have weight 1), the diameter of the tree does not increase. \blacktriangleleft

5 Average Linkage HAC on Paths in NC

In this section, we present an $\tilde{O}(n)$ work and $O(\text{polylog}(n))$ depth algorithm for solving average linkage HAC on path graphs, provided that the aspect ratio of the input instance is bounded by $\text{poly}(n)$. The *aspect ratio* is defined as $\mathcal{A} = W_{\max}/W_{\min}$, where $W_{\max} = \arg \max_{e \in E} w(e)$ and $W_{\min} = \arg \min_{e \in E} w(e)$ (note that this definition excludes all non-edges, which implicitly have a weight of 0).

In average linkage HAC, the weight (i.e., similarity) of edges monotonically decreases over time. Thus, our idea is to partition the edges into *buckets* where the edges in any bucket have the same similarity, up to constant factors. Next, we process these buckets in *phases*, from the highest similarity bucket to the lowest. In each phase, we perform a modified version of the classic nearest-neighbor chain algorithm, wherein we compute the nearest-neighbor chains for the graph induced on the edges in that bucket, and process each chain independently. We note that when we use the terminology *nearest neighbor* of a vertex in what follows, we refer to the neighbor along the *highest weight edge* incident to the vertex.

Initially, each cluster is a singleton, and we might end up with $\Omega(n)$ sequential dependencies to resolve. However, we observe that in this special case when the size of every cluster is equal, starting with the reciprocal pair, every *alternate* edge in this chain can be merged independently, and the rest of the edges will be moved to a later bucket. We can compute the edges that will be merged easily via a simple *prefix-sum* routine [8]. However, when the cluster sizes are arbitrary, this observation no longer holds. Nonetheless, we show that we can partition each chain further into $O(\log n)$ -sized subchains such that, even though the dependencies within a subchain must be resolved sequentially, the dependencies across subchains can be resolved in parallel using a similar application of *prefix-sum*.

We state the following theorems showing (1) that our parallel algorithm is highly efficient (it runs near-linear time in the number of nodes) and runs in poly-logarithmic depth and (2) that our algorithm implies that the dendrogram height of a path input with polynomial aspect ratio is always poly-logarithmic. Due to space constraints, we defer the proofs and the details of the algorithm to the full version [6].

► **Theorem 4.** *Average linkage HAC on paths is in NC. In particular, there is an algorithm for average linkage HAC that runs in $O(\log^2 n \log \log n)$ depth with $O(n \log n \log \log n)$ work.*

► **Theorem 14.** *Average linkage HAC on path graphs with $\text{poly}(n)$ aspect-ratio returns a dendrogram with height at most $O(\log^2 n)$.*

6 Conclusion

In this paper, we studied the parallel and sequential complexity of hierarchical graph clustering. We gave new classic and fine-grained reductions for Hierarchical Agglomerative Clustering (HAC) under the average linkage measure that likely rule out efficient algorithms for *exact* average linkage, parallel or otherwise. We also showed that such impossibility results can be circumvented if the output dendrogram has low height or is a path. An interesting question is whether such structure can be leveraged for other variants of interest of average linkage HAC: for example, can we obtain dynamic algorithms for HAC that are also parameterized by the height?

References

- 1 Amir Abboud, Vincent Cohen-Addad, and Hussein Houdrouge. Subquadratic high-dimensional hierarchical clustering. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Annual Conference on Neural Information Processing Systems (NeurIPS)*, volume 32. Curran Associates, Inc., 2019.
- 2 Amir Abboud, Nick Fischer, and Yarín Shechter. Faster combinatorial k-clique algorithms. *arXiv preprint*, 2024. [arXiv:2401.13502](https://arxiv.org/abs/2401.13502).
- 3 Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 522–539. SIAM, 2021.
- 4 N. S. Arora, R. D. Blumofe, and C. G. Plaxton. Thread scheduling for multiprogrammed multiprocessors. *ACM Transactions on Computer Systems*, 34(2), April 2001.
- 5 MohammadHossein Bateni, Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Raimondas Kiveris, Silvio Lattanzi, and Vahab Mirrokni. Affinity clustering: Hierarchical clustering at scale. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 6864–6874, 2017.
- 6 MohammadHossein Bateni, Laxman Dhulipala, Kishen N Gowda, D Ellis Hershkowitz, Rajesh Jayaram, and Jakub Łącki. It's hard to hac with average linkage!, 2024. [arXiv:2404.14730](https://arxiv.org/abs/2404.14730).
- 7 J-P Benzécri. Construction d'une classification ascendante hiérarchique par la recherche en chaîne des voisins réciproques. *Cahiers de l'analyse des données*, 7(2):209–218, 1982.
- 8 Guy E Blelloch. Scans as primitive parallel operations. *IEEE Transactions on computers*, 38(11):1526–1538, 1989.
- 9 Guy E. Blelloch, Jeremy T. Fineman, Yan Gu, and Yihan Sun. Optimal parallel algorithms in the binary-forking model. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2020.
- 10 Robert D. Blumofe and Charles E. Leiserson. Space-efficient scheduling of multithreaded computations. *SIAM J. on Computing*, 27(1), 1998.
- 11 Mélanie Boly, Vincent Perlbarg, Guillaume Marrelec, Manuel Schabus, Steven Laureys, Julien Doyon, Mélanie Pélégriani-Issac, Pierre Maquet, and Habib Benali. Hierarchical clustering of brain activity during human nonrapid eye movement sleep. *Proceedings of the National Academy of Sciences*, 109(15):5856–5861, 2012.
- 12 Vincent Cohen-Addad, Varun Kanade, Frederik Mallmann-Trenn, and Claire Mathieu. Hierarchical clustering: Objective functions and algorithms. *Journal of the ACM (JACM)*, 66(4), 2019.

- 13 Stephen A Cook, Yuval Filmus, and Dai Tri Man Le. The complexity of the comparator circuit value problem. *ACM Transactions on Computation Theory (TOCT)*, 6(4):1–44, 2014.
- 14 Don Coppersmith and Shmuel Winograd. On the asymptotic complexity of matrix multiplication. *SIAM Journal on Computing*, 11(3):472–492, 1982.
- 15 Laxman Dhulipala, David Eisenstat, Jakub Lacki, Vahab Mirrokni, and Jessica Shi. Hierarchical agglomerative graph clustering in nearly-linear time. In *International Conference on Machine Learning (ICML)*, pages 2676–2686, 2021.
- 16 Laxman Dhulipala, David Eisenstat, Jakub Lacki, Vahab Mirrokni, and Jessica Shi. Hierarchical agglomerative graph clustering in poly-logarithmic depth. *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 35:22925–22940, 2022.
- 17 Laxman Dhulipala, Jakub Łącki, Jason Lee, and Vahab Mirrokni. Terahac: Hierarchical agglomerative clustering of trillion-edge graphs. *Proceedings of the ACM on Management of Data*, 1(3):1–27, 2023.
- 18 Michael B Eisen, Paul T Spellman, Patrick O Brown, and David Botstein. Cluster analysis and display of genome-wide expression patterns. *Proceedings of the National Academy of Sciences*, 95(25):14863–14868, 1998.
- 19 Raymond Greenlaw and Sanpawat Kantabutra. On the parallel complexity of hierarchical clustering and cc-complete problems. *Complexity*, 14(2):18–28, 2008.
- 20 Guan-Jie Hua, Che-Lun Hung, Chun-Yuan Lin, Fu-Che Wu, Yu-Wei Chan, and Chuan Yi Tang. MGUPGMA: a fast UPGMA algorithm with multiple graphics processing units using NCCL. *Evolutionary Bioinformatics*, 13:1176934317734220, 2017.
- 21 J Juan. Programme de classification hiérarchique par l’algorithme de la recherche en chaîne des voisins réciproques. *Cahiers de l’analyse des données*, 7(2):219–225, 1982.
- 22 Benjamin King. Step-wise clustering procedures. *Journal of the American Statistical Association*, 62(317):86–101, 1967.
- 23 Ari Kobren, Nicholas Monath, Akshay Krishnamurthy, and Andrew McCallum. A hierarchical algorithm for extreme clustering. In *International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 255–264, 2017.
- 24 Godfrey N Lance and William Thomas Williams. A general theory of classificatory sorting strategies: 1. hierarchical systems. *The computer journal*, 9(4):373–380, 1967.
- 25 Joshua Lau and Angus Ritossa. Algorithms and hardness for multidimensional range updates and queries. In *Innovations in Theoretical Computer Science Conference (ITCS)*, 2021.
- 26 François Le Gall. Faster algorithms for rectangular matrix multiplication. In *Symposium on Foundations of Computer Science (FOCS)*, pages 514–523. IEEE, 2012.
- 27 Christopher D Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- 28 Ernst W Mayr and Ashok Subramanian. The complexity of circuit value and network stability. *Journal of Computer and System Sciences*, 44(2):302–323, 1992.
- 29 Nicholas Monath, Kumar Avinava Dubey, Guru Guruganesh, Manzil Zaheer, Amr Ahmed, Andrew McCallum, Gokhan Mergen, Marc Najork, Mert Terzihan, Bryon Tjanaka, et al. Scalable hierarchical agglomerative clustering. In *International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1245–1255, 2021.
- 30 Benjamin Moseley and Joshua R. Wang. Approximation bounds for hierarchical clustering: Average linkage, bisecting k-means, and local search. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 3094–3103, 2017.
- 31 Benjamin Moseley and Joshua R Wang. Approximation bounds for hierarchical clustering: Average linkage, bisecting k-means, and local search. *Journal of Machine Learning Research*, 24(1):1–36, 2023.
- 32 Daniel Müllner. Modern hierarchical, agglomerative clustering algorithms. *arXiv preprint*, 2011. [arXiv:1109.2378](https://arxiv.org/abs/1109.2378).
- 33 Fionn Murtagh and Pedro Contreras. Algorithms for hierarchical clustering: an overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 2(1):86–97, 2012.

18:18 It's Hard to HAC Average Linkage!

- 34 Fionn Murtagh and Pedro Contreras. Algorithms for hierarchical clustering: an overview, ii. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 7(6):e1219, 2017.
- 35 Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- 36 JM Shearer and Michael A Wolfe. Alglib, a simple symbol-manipulation package. *Communications of the ACM*, 28(8):820–825, 1985.
- 37 Peter Henry Andrews Sneath. The principles and practice of numerical classification. *Numerical taxonomy*, 573, 1973.
- 38 Ashok Subramanian. *A new approach to stable matching problems*. Stanford University, 1989.
- 39 Tom Tseng, Laxman Dhulipala, and Julian Shun. Parallel batch-dynamic minimum spanning forest and the efficiency of dynamic agglomerative graph clustering. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 233–245, 2022.
- 40 Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.
- 41 Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *Annual ACM Symposium on Theory of Computing (STOC)*, pages 887–898, 2012.
- 42 Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *Symposium on Foundations of Computer Science (FOCS)*, pages 645–654, 2010.
- 43 Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. New bounds for matrix multiplication: from alpha to omega. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3792–3835. SIAM, 2024.
- 44 Ying Zhao and George Karypis. Evaluation of hierarchical clustering algorithms for document datasets. In *Conference on Information and Knowledge Management (CIKM)*, pages 515–524, 2002.