

# Parallel Filtered Graphs for Hierarchical Clustering

Shangdi Yu  
MIT CSAIL

Julian Shun  
MIT CSAIL

**Abstract**—Given all pairwise weights (distances) among a set of objects, filtered graphs provide a sparse representation by only keeping an important subset of weights. Such graphs can be passed to graph clustering algorithms to generate hierarchical clusters. In particular, the directed bubble hierarchical tree (DBHT) algorithm on filtered graphs has been shown to produce good hierarchical clusters for time series data.

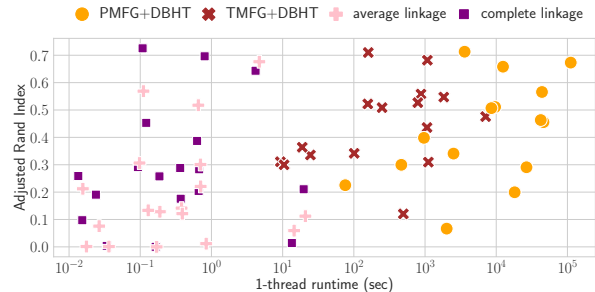
We propose a new parallel algorithm for constructing triangulated maximally filtered graphs (TMFG), which produces valid inputs for DBHT, and a scalable parallel algorithm for generating DBHTs that is optimized for TMFG inputs. In addition to parallelizing the original TMFG construction, which has limited parallelism, we also design a new algorithm that inserts multiple vertices on each round to enable more parallelism. We show that the graphs generated by our new algorithm have similar quality compared to the original TMFGs, while being much faster to generate. Our new parallel algorithms for TMFGs and DBHTs are 136–2483x faster than state-of-the-art implementations, while achieving up to 41.56x self-relative speedup on 48 cores with hyper-threading, and achieve better clustering results compared to the standard average-linkage and complete-linkage hierarchical clustering algorithms. We show that on a stock data set, our algorithms produce clusters that align well with human experts’ classification.

## I. INTRODUCTION

Clustering is an unsupervised machine learning method that has been widely used in many fields including finance, biology, and computer vision, to discover structures in a data set. Often times, one is interested in exploring clusters at varying resolutions. A *hierarchical clustering* algorithm can be used to produce a tree, also known as a *dendrogram*, that represents clusters at different scales.

Running a metric clustering algorithm on a set of  $n$  points often involves working with  $\Theta(n^2)$  pairwise distances, and is computationally prohibitive on large data sets. One approach to improving efficiency is to use a *filtered graph* that keeps only a subset of the pairwise distances, and then pass the resulting graph to a graph clustering algorithm. Filtered graphs reduce the number of distances considered while retaining the most important features, both locally and globally. Simply removing all edges with weights below a certain threshold may not perform well in practice, as the threshold may require significant tuning, and the importance of an edge is not only determined by its weight, but also its location in the graph. Several other methods for constructing filtered graphs have been proposed, including  $k$ -nearest neighbor graphs [1], minimum spanning trees [2], [3], and weighted maximal planar graphs [4], [5].

A variant of weighted maximal planar graphs, called planar maximally filtered graphs (PMFG) [4], has been shown to perform well in practice in combination with the directed



**Fig. 1:** Sequential runtime vs. clustering quality for hierarchical clustering methods on different data sets.

bubble hierarchy tree (DBHT) technique, especially on financial [6]–[8] and biological [9], [10] data sets. The PMFG is constructed by repeatedly adding an edge between the pair of points with the highest weight, while preserving planarity. The DBHT technique constructs a dendrogram based on certain triangles in the input graph, along with various shortest path calculations. It has the benefit that no prior information about the data is required, and no parameter tuning is needed. However, the state-of-the-art PMFG and DBHT algorithms are sequential and do not scale to large data sets. Massara et. al [5] proposed the triangulated maximally filtered graphs (TMFG), a variant of maximal weighted planar graph that can be generated more efficiently. Instead of repeatedly adding a single edge as in PMFG, the TMFG is constructed by repeatedly adding a single vertex to a triangle with its three edges to the triangle corners, while respecting planarity. In other words, we search for an uninserted vertex  $x$  that is “close” to three vertices that already form a triangle in the graph, and insert  $x$  by adding three edges from  $x$  to the three vertices. However, the state-of-the-art TMFG implementation is also sequential, and to the best of our knowledge, the clustering quality of using DBHTs with TMFGs has not been evaluated.

One important type of data that PMFGs, TMFGs, and DBHTs have been shown to perform well on is *time series data*. Time series data arise in a multitude of applications, including in finance, signal processing, biology, astronomy, and weather forecasting. To extract insights from the data, one is often interested in finding correlations between different time series, and clustering the data based on these correlations. The *correlation matrix* stores the correlation between every pair of time series. It is important to construct a filtered graph on the correlation matrix to enable efficient and scalable clustering. We show in Figure 1 the runtime and cluster quality (using the Adjusted Rand Index [11]) for PMFG and TMFG combined with DBHT, compared with the standard average-linkage and complete-linkage methods for hierarchical

clustering, for a collection of time series data sets. For all methods, the dendrogram is cut such that the number of resulting clusters is the same as the number of ground truth clusters. We see that, while the runtimes of PMFG and TMFG are higher than those of average-linkage and complete-linkage, they are able to generate higher quality clusters. Thus, PMFG, TMFG, and DBHT are preferable for hierarchical clustering, when one is willing to use a larger time budget to obtain better quality clusters.

To reduce the runtime of TMFG and DBHT, we propose new parallel algorithms and fast implementations for constructing TMFGs and a scalable parallel DBHT algorithm that is optimized for TMFGs. Besides parallelizing the original TMFG construction, which has limited parallelism, we also design a new algorithm that inserts multiple vertices on each round to enable more parallelism. We show that the filtered graph generated by our new algorithm has similar quality to that of the original TMFG, while being much faster. The key challenge in our algorithm is in resolving the conflicts when inserting multiple vertices into the TMFG. In TMFG construction, we resolve the conflict of vertices by designating a single triangle for each vertex based on edge weights. Our DBHT algorithm leverages the special topological structure of TMFGs. We are able to construct a rooted undirected bubble tree on the fly with a special invariant while constructing the TMFG, and we use this invariant to efficiently direct the bubble tree with a recursive process. This reduces the complexity of these two steps from quadratic (which the original DBHT algorithm requires) to linear.

We compare the speed and quality of our parallel algorithms with parallelized versions of the average-linkage and complete-linkage hierarchical clustering algorithms. As a baseline, we also compare with  $k$ -means, which is a non-hierarchical clustering algorithm and only produces clusters at a single resolution. On a collection of 16 data sets generated from time series and image data, we find that the DBHT using TMFG produces similar and sometimes better clusters than the DBHT using PMFG. It also gives better clusters than average-linkage and complete-linkage clustering, and is competitive with  $k$ -means on most data sets.

We show that our new algorithm is up to 15589 times faster than the sequential DBHT on PMFG and has Adjusted Rand Index scores up to 0.65 higher than agglomerative clustering algorithms. We show that on time series data sets of stock prices from 2013–2019 from the US stock market, DBHT on TMFG is able to produce clusters that align well with human experts’ classification, and in our experiment it produced better clusters than the original TMFG algorithm. On 48 cores with hyper-threading, our new algorithm for constructing a TMFG variant and DBHT are 136–2483x faster times faster than the state-of-the-art TMFG implementation, and achieves up to 41.56x self-relative speedup on 48 cores.

We summarize our contributions below:

- We design a new parallel algorithm for constructing a TMFG variant that produces high-quality graphs.

- We design a new parallel algorithm for constructing the DBHT, which is optimized for TMFG-like inputs.
- We perform experiments showing that our parallel algorithms achieve significant speedups over state-of-the-art algorithms, while producing clusters of similar or better quality.

Our source code and data is available at <https://github.com/yushangdi/par-filtered-graph-clustering>.

## II. BACKGROUND

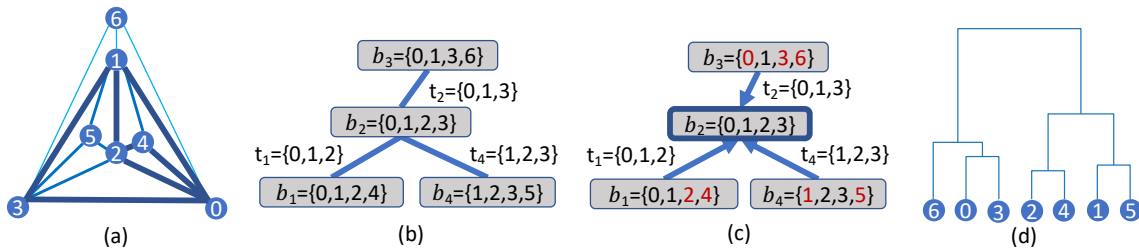
We briefly describe the sequential construction of PMFGs and TMFGs, and introduce the notation and primitives that we use.

A *planar* graph can be drawn on the plane such that no edges cross each other. Both PMFGs and TMFGs are maximal planar subgraphs of a complete undirected, edge-weighted graph (“maximal” means that no additional edge can be added to the vertex set without violating planarity). They give an approximation to the NP-hard Weighted Maximum Planar Graph problem, which requires the sum of the edge weights kept to be maximized [12]. The complete edge-weighted graph input can also be viewed as a similarity matrix  $S$ , where  $S[i, j]$  is the weight of edge  $(i, j)$  in the graph.

**Planar Maximally Filtered Graph (PMFG).** The sequential algorithm for constructing PMFGs [4] first sorts all of the edges by their edge weights. It then starts with an empty graph  $G$  and considers each edge in the sorted order, from highest to lowest weight. An edge is added to the graph if and only if it does not violate the planarity of  $G$ , which is checked by running a planarity testing algorithm. This process is computationally expensive due to the need to run a planarity testing algorithm  $\Theta(n^2)$  times.

**Triangulated Maximally Filtered Graph (TMFG).** TMFGs [5] have been proposed to improve the efficiency of PMFG construction. Instead of considering edges one by one, the original sequential TMFG algorithm adds one vertex  $x$  to the vertex set of the subgraph under construction, as well as three edges (each from  $x$  to one of the three vertices of a triangular face in a planar embedding of the subgraph) on every iteration, and eliminates the vertex from consideration in future iterations. The vertex-triangle pair that gives the highest increase in total edge weight is chosen. This leads to only  $\Theta(n)$  iterations for TMFG, compared to  $\Theta(n^2)$  for PMFG. The algorithm starts with four vertices with maximum weighted degree and all six edges connecting them. By definition, it ensures that edges added do not violate planarity, and so planarity testing is not needed. Figure 2(a) shows an example of a TMFG. If there were additional vertices to insert, they could go into any of the faces, which are all triangles.

**Directed Bubble Hierarchy Tree (DBHT).** After obtaining the TMFG or PMFG, we can then generate a dendrogram from it for hierarchical clustering using the DBHT [13] algorithm, which has been shown to produce high-quality dendrograms for financial time series [6]–[8] and biological [9], [10] data. The original, sequential DBHT algorithm first constructs a *bubble tree* [14], which is a tree where nodes correspond to planar subgraphs, and edges between two nodes corre-



**Fig. 2:** (a) An example TMFG. The edges have weights 0.8, 0.4, or 0.2. The darker and thicker edges have larger weights. For example,  $w(0, 1) = 0.8$ ,  $w(2, 3) = 0.4$ , and  $w(0, 6) = 0.2$ . (b) The undirected bubble tree. Each bubble node is a gray box and nodes are connected by blue edges. Each bubble is marked with vertices in the bubble and a name  $b_i$ . Each edge is marked with the triangle it corresponds to and a name  $t_i$ . (c) The directed bubble tree. The bolded node is the only converging bubble in this example. The red vertices are assigned to the bubble in which they are marked. (d) The dendrogram for DBHT.

spond to triangles in the original graph that separate the two corresponding planar subgraphs. The DBHT algorithm then directs the tree edges by computing the strength of connection from the separating face to its interior and exterior. Next, the algorithm runs a number of shortest distance computations to assign vertices to bubbles, and finally uses complete-linkage clustering to generate a hierarchy. We will describe more details of the DBHT algorithm along with how we parallelize and optimize it in Section V.

**Notation and Terminology.** Given a planar graph  $G$  and its planar embedding, a *face* of  $G$  is a maximal region bounded by edges, except for the *outer face*, which is unbounded. For example, in Figure 2(a),  $\{0, 3, 6\}$  is the outer face in this embedding. A bounded face is an *inner face*. We *insert* vertex  $v$  into a triangular face  $t$  with three corners  $v_x$ ,  $v_y$ , and  $v_z$  by adding three weighted edges from  $v$  to each of  $v_x$ ,  $v_y$ , and  $v_z$ . The *gain* of inserting  $v$  to  $t$  is the sum of these three edge weights. Let  $V$  be a set of vertices. We say that a vertex  $v \in V$  is the *best vertex in  $V$*  for face  $t$  if inserting  $v$  to  $t$  gives the maximum gain among all vertices in  $V$ . The weight of an edge  $(i, j)$  in graph  $G$  is denoted as  $w(i, j)$ .

A face is *separating* if removing it disconnects the graph. A *bubble* is a maximal planar graph whose 3-cliques are non-separating cycles. If we create a *bubble node* for each bubble and connect bubble nodes that share separating triangles, then the bubble nodes and connection edges form a tree [13], [14], which we call a *bubble tree*. We will denote a clique with  $\mathcal{C}$  and a bubble node in the bubble tree as  $b$ .

The output of our algorithm is a tree called a *dendrogram*, where the height of each node corresponds to the dissimilarity between the two clusters represented by its two children. The *height* of a dendrogram node is a value computed by our algorithm, and a node's height is at most the height of its parent. Cutting the dendrogram at different heights gives subtrees that correspond to clusters at different scales.

**Model of Computation.** To analyze our parallel algorithms, we use the *work-span model* [15], [16], a standard model for analyzing shared-memory algorithms. The work  $W$  of an algorithm is the total number of operations, and the span  $S$  is the length of the longest dependency path of the algorithm. The work of a sequential algorithm is the same as its time. We can bound the running time of the algorithm on  $P$  processors by  $W/P + O(S)$  using a randomized work-stealing

**TABLE I:** Parallel Primitives.

Operator	Description	Work	Span
Filter	Filter out desired elements	$O(n)$	$O(\log n)$
Sort	Sort elements	$O(n)$	$O(\log n)$ w.h.p.
Maximum	Find the maximum	$O(n)$	$O(1)$ w.h.p.
WRITEMIN	concurrently write the smallest value	$O(1)$	$O(1)$
WRITEMAX	concurrently write the largest value	$O(1)$	$O(1)$
WRITEADD	atomically add	$O(1)$	$O(1)$

scheduler [17].

**Parallel Primitives [15], [18].** Table I gives a summary of the parallel primitives that we use. A parallel *filter* takes an array  $A$  and a predicate function  $f$ , and returns a new array containing  $a \in A$  for which  $f(a)$  is true, in the same order that they appear in  $A$ . Filter takes  $O(n)$  work and  $O(\log n)$  span. A parallel *sort* takes an array  $A$  and a binary function  $f(a, b)$  that returns true if  $a <_k b$  where  $<_k$  is a total ordering, and returns a new array containing  $a \in A$  in non-decreasing order. Sorting takes  $O(n \log n)$  work and  $O(\log n)$  span. A parallel *integer sort* can be done in  $O(n)$  work and  $O(\log n)$  span with high probability (w.h.p.)<sup>1</sup> for integers in the range  $[1, \dots, O(n \log^c n)]$  [19]. A parallel *maximum* takes an array  $A$  and returns the largest element in  $A$ . A maximum can be computed in  $O(n)$  work and  $O(1)$  span with high probability. WRITEMIN is a priority concurrent write that takes as input two arguments, where the first argument is the location to write to and the second argument is the value to write; on concurrent writes, the smallest value is written [20]. WRITEMAX is similar but writes the largest value. WRITEADD is a priority concurrent write that takes as input two arguments, where the first argument is the location to write to and the second argument is the value to atomically add to the value at the first location. We assume that these priority concurrent writes take constant work and span.

### III. RELATED WORK

There is a rich literature in designing both sequential and parallel hierarchical clustering algorithms. The most popular hierarchical clustering algorithm variants are hierarchical agglomerative clustering algorithms (HAC) [21]–[25]. Different variants of HAC use different linkage functions to compute the distance between clusters. Popular linkage functions include complete,

<sup>1</sup>We say that a bound holds *with high probability (w.h.p.)* if it holds with probability at least  $1 - n^{-q}$  for  $q \geq 1$ , where  $n$  is the input size.

average, single, centroid, and median linkage. Several papers have been written on the topic of parallel HAC, such as parallel HAC under the single-linkage metric [22], [23], [26]–[28] as well as other metrics [22]. Song et al. [13] showed that DBHT with PMFG is superior to complete-linkage and average-linkage HAC. We also show in Section VII that our method has higher quality than complete-linkage and average-linkage HAC. Musmeci et al. [6] showed that DBHT with PMFG produces better clusters on stock data sets than single linkage, average linkage, complete linkage, and  $k$ -medoids.

There has also been work on other hierarchical clustering methods, such as partitioning hierarchical clustering algorithms, algorithms that combine agglomerative and partitioning methods [29], [30], and density-based methods [23], [31]. Compared to the density-based methods [23], [31], our algorithm has two advantages, the first is that many density-based methods are most suitable for low-dimensional data, while our algorithm does not have this constraint. Moreover, these clustering methods require some apriori information to set the parameters to obtain high-quality clusters. In comparison, DBHT does not require any additional parameters. Dash et al. [29] proposed a parallel hierarchical clustering algorithm that divides data into overlapping cells and clusters the cells. However, the parallel version is only described for low-dimensional data sets. Cohen-Addad et al. [32] designed a parallel algorithm for the hierarchical  $k$ -median problem, but it is limited to the Euclidean metric. There are also algorithms designed for specific types of data, such as categorical data [33] and genomics data [30]. While most parallel algorithms have been designed for shared memory, there have also been algorithms designed for distributed memory [34].

The PMFG is an approximate solution to the weighted maximal planar graph (WMPG) problem, which is NP-hard [12]. There have also been other approximate solutions designed for the WMPG. Massara et al. [5] propose the TMFG, which is based on local topological moves. Katakai et al. [35] design a heuristic that is based on a transformation to the connected spanning subgraph problem and the dual graph. Osman et al. [36] use a greedy random adaptive search procedure to solve the WMPG problem. Other heuristics have been developed for the WMPG problem [37]–[39].

Besides filtering graphs based on topological constraints, there are other graph filtering methods that are used for clustering. For example, Ma et al. [40] use a low-pass filter to extract useful data representations for subspace clustering, and Tremblay et al. [41] use graph filtering for faster spectral clustering. Another graph construction algorithm used for clustering is  $b$ -matching [42]–[44]. The  $b$ -matching algorithm filters graph edges such that the in-degree and out-degree of each node is at most  $b$ , and can be solved in polynomial time.

#### IV. PARALLEL TMFG

This section introduces our algorithm for parallelizing TMFG construction, whose pseudocode is shown in Algorithm 1. We first give a high-level overview. Our parallel algorithm simultaneously adds multiple vertices to the graph. Adding

---

#### Algorithm 1: Parallel TMFG

---

```

Input:  $n \times n$  similarity matrix  $S$ , prefix size PREFIX  $\geq 1$ 
// The four vertices that have highest total
sum across its row in  $S$ 
1  $\mathcal{C} = \{v_1, v_2, v_3, v_4\}$ 
2  $\mathcal{E} = \{(v_1, v_2)(v_1, v_3), (v_1, v_4), (v_2, v_3), (v_2, v_4), (v_3, v_4)\}$ 
// The four triangular faces in the initial
graph
3  $\mathcal{F} = \{\{v_1, v_2, v_3\}, \{v_1, v_2, v_4\}, \{v_1, v_3, v_4\}, \{v_2, v_3, v_4\}\}$ 
4  $V = \{v_5, \dots, v_n\}$ 
// The best vertex in  $V$  for each triangle
5 GAINS = [argmax $_{u \in V} (\sum_{y \in t} S[v, u])$  for  $t \in \mathcal{F}$ ]
6 Initialize bubble tree  $T$  with  $\mathcal{C}$ 
7 OUTERFACE =  $\{v_1, v_2, v_3\}$ 
8 while  $|V| > 0$  do
9    $L$  = the PREFIX vertex-face pairs with the largest gains in GAINS using
parallel sorting.
10  For vertices paired with multiple faces in  $L$ , only keep the pair with
maximum gain, using a parallel filter and parallel sorting.
11   $V = V \setminus \{\text{vertices} \in L\}$ 
12  parallel for  $(v, t = \{v_x, v_y, v_z\}) \in L$  do
13     $\mathcal{E} = \mathcal{E} \cup \{(v, v_x), (v, v_y), (v, v_z)\}$ 
14     $\mathcal{F} = \mathcal{F} \cup \{\{v, v_x, v_y\}, \{v, v_y, v_z\}, \{v, v_x, v_z\}\} \setminus t$ 
15    parallel for  $t' \in \{t' : \text{GAINS}[t'] =$ 
16       $v\} \cup \{\{v, v_x, v_y\}, \{v, v_y, v_z\}, \{v, v_x, v_z\}\}$  do
17      | GAINS[ $t'$ ] = argmax $_{u \in V} (\sum_{v \in t'} S[v, u])$ 
18      | UpdateBubbleTree( $v, t', T$ )
18 return  $\mathcal{E}$ 

```

---

multiple vertices may give results that differ from the sequential TMFG algorithm because the sequential algorithm updates the graph after every insertion, giving future vertices more faces to choose from, whereas our parallel algorithm will only update the graph after all of the vertices in a round have been added. Therefore, the more vertices we add in parallel, the more we deviate from the sequential algorithm. This gives a tradeoff in the quality of the graph and the amount of available parallelism. We therefore consider adding only a prefix of the vertices that give the highest gain in total edge weight. The prefix size can be tuned for the application and available parallelism.

In the rest of the section, we describe the details of our parallel TMFG algorithm (Algorithm 1). The input is an  $n \times n$  symmetric matrix  $S$  that represents the complete undirected graph, and a parameter PREFIX. The highlighted lines are used for building DBHT, and will be explained in Section V. On Lines 1–4, we first find the four vertices  $\{v_1, v_2, v_3, v_4\}$  that have the highest total sum across their rows in  $S$ , and add all edges among them to the resulting graph’s edge list  $\mathcal{E}$ . The four faces formed by the four vertices are added to the set of faces  $\mathcal{F}$ . The rest of the vertices are in set  $V$ , and will be added to the graph later. On Line 5, for each face in  $\mathcal{F}$ , we find its best vertex (the vertex that maximizes the gain if inserted into this face).

On Lines 8–18, we insert the remaining vertices into the graph in batches of up to size PREFIX. On Lines 9–11, we obtain the batch of vertices to insert and their corresponding faces to insert into. We first obtain the PREFIX vertex-face pairs with the largest gains in the GAINS array, and store the pairs in  $L$ . This can be done using a parallel sort on the GAINS array. On Line 10, we ensure that a given vertex is only added to a single face to avoid race conditions. Conflicts for a given vertex are resolved by only allowing the vertex to add itself to the face that gives the highest gain. We can use parallel sorting to group vertices based on which face (among all faces) gives

**Algorithm 2: UpdateBubbleTree**


---

```

1 Function UpdateBubbleTree( $v, t = \{v_x, v_y, v_z\}, T$ ):
2    $b^* = \text{new bubble } \{v, v_x, v_y, v_z\}$ 
3    $b = \text{the bubble that } t \text{ is in}$ 
4   if  $t = \text{OUTERFACE}$  then
5      $b.\text{PARENT} = b^*$ 
6     Add  $b$  to  $b^*.\text{CHILDREN}$ 
7      $\text{OUTERFACE} = \{v, v_x, v_y\}$ 
8   else
9      $b^*.\text{PARENT} = b$ 
10    Add  $b^*$  to  $b.\text{CHILDREN}$ 

```

---

the highest gain, and have each face pick the vertex (among all vertices that choose this face) that gives the highest gain. The available parallelism increases as the number of faces in the graph grows. If PREFIX = 1, Lines 9–10 can be simplified to a single parallel maximum computation. On Line 11, we remove the vertices in  $L$  from  $V$  using a parallel filter.

On Lines 12–16, we loop over the vertex-face pairs  $(v, t)$  in  $L$ , add  $v$  to  $t$ , and update  $\mathcal{E}$  and  $\mathcal{F}$ . To update  $\mathcal{F}$ , we add the three new faces created and remove  $t$ . We also update GAINS by computing the new best vertex among  $V$  for the three new faces and faces that previously had  $v$  as their best vertex. Unlike the original algorithm [5], which loops over all of the faces to find the faces that previously had  $v$  as their best vertex, we keep an array for each vertex that stores such faces to improve the efficiency of this step.

If we choose PREFIX = 1 in Algorithm 1, we obtain the same result as the sequential TMFG algorithm, but we still have some parallelism from the parallel maximum call on Line 9 on Lines 15–16. This case gives the same parallel algorithm that is described (but not implemented) by Massara et al. [5]. However, we show in Section VII that using PREFIX = 1 has limited parallelism because on each round only a single vertex can be inserted.

## V. PARALLEL DBHT FOR TMFG

We describe our parallel algorithm for building the DBHT, which leverages the special structure of the bubble tree for TMFGs. The original DBHT construction of Song et al. [13] takes a planar graph, and two weights on each edge—a similarity measure and a dissimilarity measure. The similarity measure is from the similarity matrix  $S$ , and the dissimilarity measure needs to be additionally supplied. The construction has the following steps: building an undirected bubble tree [14] from a planar graph; assigning directions to the bubble tree edges; assigning graph vertices to bubbles; and using complete-linkage clustering to generate a hierarchy. Below, we describe how we accelerate and parallelize each step.

### A. Bubble Tree for TMFG

As described in Section II, a *bubble tree* [14] is a tree where nodes correspond to planar subgraphs, and edges between nodes correspond to triangles in the original graph that separate the two corresponding planar subgraphs. Specifically, a *bubble* is a maximal planar graph whose triangles are non-separating [14]. The original DBHT construction is inefficient, as it involves first finding all of the triangles in the graph, and then testing for every triangle whether removing it would disconnect the graph.

The key observation is that the bubble tree can be more efficiently constructed during TMFG construction, as a vertex that is added in the TMFG algorithm will correspond to exactly one bubble node and one edge in the bubble tree. In TMFG, every time we insert a new vertex, we create a 4-clique, which is a bubble because all faces of a 4-clique are non-separating triangles. This will also produce a separating triangle, which corresponds to a new edge in the bubble tree. This separating triangle is shared by the 4-cliques corresponding to the two bubble nodes incident to it. For example, consider the graph without vertex 5 in Figure 2(a). If we now insert vertex 5 into triangle  $t_4 = \{1, 2, 3\}$ , we created a new bubble  $b_4 = \{1, 2, 3, 5\}$  and  $t_4$  now becomes a separating triangle, which is shared by  $b_2 = \{0, 1, 2, 3\}$  and  $b_4$ . We incorporate the bubble tree construction into our parallel TMFG algorithm to save a significant amount of work.

Our undirected bubble tree for TMFG has the invariant that each bubble node has a parent and at most three children (because for each 4-clique, there is one outer face and three inner faces), except for the root, which does not have a parent. Moreover, all descendants of an edge are on the interior side of the separating triangle corresponding to the edge. This invariant is important for us to accelerate the direction computation, which we describe later.

We now describe the details of bubble tree construction, which are the highlighted lines in Algorithm 1. On Line 6, we initialize our bubble tree with a single bubble node, which contains the 4-clique that we start our TMFG with. On Line 7, we initialize our outer face to be  $\{v_1, v_2, v_3\}$ . This outer face can be chosen arbitrarily among the four faces of  $\mathcal{C}$  because the bubble tree’s topological structure is independent of this choice [14]. On Line 17, we add a new node and a new edge to the bubble tree  $T$  for each vertex  $v$  inserted into face  $t$ . This can be done in parallel because each vertex  $v$  is inserted into a single, unique face  $t$ , so there will not be any conflicts.

Algorithm 2 shows the subroutine for building the bubble tree. On Line 2, we create a new bubble node  $b^*$ . On Line 3, we find the bubble  $b$  in  $T$  that  $t$  is in. This is the bubble that is created when  $t$  is created (in some previous call to UpdateBubbleTree).  $b^*$  is the bubble that the faces  $\{v, v_x, v_y\}$ ,  $\{v, v_y, v_z\}$ , and  $\{v, v_x, v_z\}$  are in. If  $t$  is the current outer face, this means we are inserting  $v$  into the outer face, and  $b$  is the current root of  $T$ . So on Lines 4–7, we let  $b$ ’s parent be  $b^*$ , and add  $b$  to  $b^*$ ’s children. This maintains the invariant above because the vertices in the bubbles before inserting  $v$  are in the interior of the outer face, and after the procedure these bubbles become descendants of the edge corresponding to the outer face  $t$ . The OUTERFACE is then updated to be a face  $\{v, v_x, v_y\}$  in the 4-clique corresponding to  $b^*$ . If  $t$  is not the current outer face, then we do not need to change the outer face, because the vertex is inserted into some inner face of  $b$ . In this case, on Lines 8–10, we let  $b^*$ ’s parent be  $b$ , and add  $b^*$  to  $b$ ’s children. This maintains the invariant above because  $v$  must be in the interior of  $t$  and its ancestors, and the bubble  $b^*$  containing  $v$  is a descendant of the edge corresponding to  $t$ .

**Example 1.** We now give an example of the bubble tree



construction by running our algorithm on the example TMFG in Figure 2. We will first describe inserting a single vertex, and then describe inserting multiple vertices in parallel. Suppose we start with the 4-clique  $\mathcal{C} = \{0, 1, 2, 4\}$ . We have four faces  $\{\{0, 1, 2\}, \{0, 1, 4\}, \{0, 2, 4\}, \{1, 2, 4\}\}$ , where  $t_1 = \{0, 1, 2\}$  is the outer face. We initialize the bubble tree  $T$  with node  $b_1 = \{0, 1, 2, 4\}$ , which corresponds to  $\mathcal{C}$ . We insert vertices 3, 5, and 6, in order, into faces  $\{0, 1, 2\}$ ,  $\{1, 2, 3\}$ , and  $\{0, 1, 3\}$ , respectively. We first insert 3 into  $t = t_1 = \{0, 1, 2\}$ . For this insertion, the new bubble  $b^*$  on Line 2 is  $b_2 = \{0, 1, 2, 3\}$ , and the  $b$  on Line 3 is  $b_1$ . Since  $t$  is the current outer face, we let  $b_2$  be  $b_1$ 's parent. The new outer face is  $t_2 = \{0, 1, 3\}$ . Now suppose we insert 5 and 6 into  $\{1, 2, 3\}$  and  $\{0, 1, 3\}$ , respectively, in parallel. This is similar to inserting 3, and we add  $b_3 = \{0, 1, 3, 6\}$  as  $b_2$ 's parent and  $b_4 = \{1, 2, 3, 5\}$  as  $b_2$ 's child in parallel.

### B. Directing Bubble Tree Edges

We now describe how to direct the bubble tree edges after obtaining the undirected bubble tree from Algorithms 1 and 2, and how we significantly improve the efficiency of this step over the original sequential DBHT algorithm. Each edge of the bubble tree corresponds to a separating triangle. The direction is decided by computing the sum over the weights of the edges in the TMFG connecting the triangle with its interior versus its exterior [13]. Figure 2(c) shows an example of directing the edges.

We denote the sum over the weights of the edges to the interior as  $\text{INVAL}$  and to the exterior as  $\text{OUTVAL}$ . In the original algorithm, the two values are computed by running a breadth-first-search (BFS) on  $G \setminus t$  for each separating triangle  $t$  to find its exterior and interior, and then computing the sum of edge weights. This takes  $\Theta(n^2)$  work because PMFGs and TMFGs have  $\Theta(n)$  edges [13] and each BFS takes  $\Theta(n)$  work.

In our bubble tree, the interior of a separating triangle contains all of the vertices in the descendants of the edge corresponding to this separating triangle, and the exterior contains all other vertices. Using this property, we can compute the direction of all edges in  $\Theta(n)$  work using a novel recursive algorithm, starting from the root of bubble tree. At each bubble node  $b$ , we compute the direction of the edge from itself to its parent, and recursively call the procedure on the bubble's children. This gives a significant improvement over the quadratic work of the original algorithm.

Specifically, we compute the direction of all edges by recursively calling the `COMPUTEDIRECTION` function, shown in Algorithm 3, with the root of bubble tree as the argument to the initial call. At each bubble node  $b$ , we compute the direction of the edge from itself to its parent. If  $b$  is the root, then it has no edge to its parent, and so we do not need to compute anything besides initializing the computation on its children (Lines 20–22). Otherwise, we compute the  $\text{INVAL}$  and  $\text{OUTVAL}$  for the separating triangle corresponding to the edge from the bubble to its parent. If  $\text{INVAL} > \text{OUTVAL}$ , then the edge goes from  $b$ 's parent to  $b$ , and vice versa (Lines 14–17).

On Lines 3–4, we obtain the vertices  $\{v_x, v_y, v_z\}$  in the

### Algorithm 3: ComputeDirection

---

```

1 Function ComputeDirection(bubble tree node b):
2   if  $b$  has a parent then
3      $\{v_x, v_y, v_z\}$  = vertices shared by  $b$  and its parent
4      $v = b \setminus \{v_x, v_y, v_z\}$ 
5      $r = \{v\}$ 
6      $r[v_x] = w(v_x, v), r[v_y] = w(v_y, v), r[v_z] = w(v_z, v)$ 
7     parallel_for  $b^* \in b$ 's children do
8        $r^* = \{v_x^* : val_x^*, v_y^* : val_y^*, v_z^* : val_z^*\} =$ 
9         computeDirection( $b^*$ )
10      if  $v_x^* \in r$  then WRITEADD( $r[v_x^*], val_x^*$ )
11      if  $v_y^* \in r$  then WRITEADD( $r[v_y^*], val_y^*$ )
12      if  $v_z^* \in r$  then WRITEADD( $r[v_z^*], val_z^*$ )
13       $\text{INVAL} = r[v_x] + r[v_y] + r[v_z]$ 
14       $\text{OUTVAL} = \text{deg}(v_x) + \text{deg}(v_y) + \text{deg}(v_z) - \text{INVAL} -$ 
15         $2(w(v_x, v_y) + w(v_x, v_z) + w(v_y, v_z))$ 
16      if  $\text{INVAL} > \text{OUTVAL}$  then
17        Direct the edge from  $b$ 's parent to  $b$ 
18      else
19        Direct the edge from  $b$  to  $b$ 's parent
20      return  $r$ 
21   else /*  $b$  is the root
22     parallel_for  $b^* \in b$ 's children do
23       computeDirection( $b^*$ )
24     return  $\{v\}$ 

```

---

### Algorithm 4: Parallel DBHT for TMFG

---

```

Input:  $n \times n$  dissimilarity matrix  $D$ , weighted undirected planar graph  $G$ 
generated by TMFG, bubble tree  $T$ 
1 computeDirection(root of  $T$ )
2  $v.g = (-\infty, -\infty), v.q = (-\infty, -\infty) \forall$  vertices  $v \in G$ 
3  $\text{BB} =$  set containing  $T$ 's nodes
4  $\text{CVGBB} = \{b \in \text{BB} : b.\text{out-degree} = 0\}$ 
5 parallel_for  $b \in \text{BB}$  do
6   BFS( $T, b$ )
7 // all-pairs shortest paths
8  $A = \text{allPairsShortestPaths}(G, D)$ 
9 // assign vertices to converging bubbles
10 parallel_for  $b \in \text{CVGBB}$  do
11   parallel_for  $v \in b$  do
12      $\chi = \text{computeChi}(v, b)$ 
13     WRITEMAX( $v.g, (\chi, b)$ )
14  $V_b^0 = \{v \in G : v.g = b\} \forall b$ 
15  $v.g = (\infty, \infty)$  if  $v.g$  is  $(-\infty, -\infty) \forall v$ 
16 parallel_for  $b \in \text{CVGBB}$  do
17   parallel_for  $v \rightarrow b$  and  $v.g == (\infty, \infty)$  do
18      $L = \text{computeAverageShortestPath}(v, V_b^0)$ 
19     WRITEMIN( $v.g, (L, b)$ )
20 // assign vertices to bubbles
21 parallel_for  $b \in \text{BB}$  do
22    $\chi_{\text{total}} = 0$ 
23   for  $v \in b$  do
24      $\chi_v = \text{computeChi}(v, b); \chi_{\text{total}} += \chi_v$ 
25   for  $v \in b$  do
26     WRITEMAX( $v.q, (\chi_v / \chi_{\text{total}}, b)$ )
27 // build hierarchy using complete linkage
28  $\{\mathcal{Z}_1, \dots, \mathcal{Z}_n\}$  /* initialize dendrogram nodes */
29 parallel_for  $b^c \in \text{CVGBB}$  do
30    $\text{BB}_{b^c} = \{v.q : v \in b^c\}$ 
31   parallel_for  $b \in \text{BB}_{b^c}$  do
32     Let  $\mathcal{Z}_{(b^c, b)} = \text{completeLinkage}(\{\mathcal{Z}_v : v.q = b \wedge v.g = b^c\})$ 
33   parallel_for  $b^c \in \text{CVGBB}$  do
34     Let  $\mathcal{Z}_{b^c} = \text{completeLinkage}(\{\mathcal{Z}_{(b^c, *)} : b^{c'} = b^c\})$ 
35    $\mathcal{Z} = \text{completeLinkage}(\{\mathcal{Z}_{b^c} : b^c \in \text{CVGBB}\})$ 
36   dendrogram = computeHeight( $\mathcal{Z}$ )
37 return dendrogram

```

---

separating triangle represented by the edge from  $b$  to its parent, and let  $v$  be the remaining vertex in the bubble. On Lines 5–6, we initialize the sum of edges to the interior from each corner of the triangle with the edge weights from each corner to  $v$  since  $v$  is in the interior. Then, on Lines 7–11 we recursively, and in parallel, compute the sum of edge weights from the corners of  $b$ 's children to the children's interior. Note that this computation is nested parallel, meaning that parallel tasks are recursively created. Since the children's interior is also  $b$ 's interior, we can sum the weights of  $b$ 's corners obtained from its children to compute the  $\text{INVAL}$  of  $b$ . On Line 13, we compute  $\text{OUTVAL}$  by subtracting the  $\text{INVAL}$  and the edge

weights of the triangle from the weighted degrees of the corners of the triangle. On Line 18, we return the weights at the corners to  $b$ 's parent.

Now, we describe the derivation of the formula on Line 13 of Algorithm 3. Let  $D = \text{deg}(v_x) + \text{deg}(v_y) + \text{deg}(v_z)$ ,  $I_k$  be the total weighted degree of  $v_k$  to the interior, and  $O_k$  be the total weighted degree of  $v_k$  to the exterior. We can rewrite  $D$  as follows:

$$\begin{aligned} D &= I_x + O_x + w(v_x, v_y) + w(v_x, v_z) \\ &\quad + I_y + O_y + w(v_y, v_x) + w(v_y, v_z) \\ &\quad + I_z + O_z + w(v_z, v_x) + w(v_z, v_y) \\ &= (I_x + I_y + I_z) + (O_x + O_y + O_z) \\ &\quad + 2(w(v_x, v_y) + w(v_x, v_z) + w(v_y, v_z)) \\ &= \text{INVAL} + \text{OUTVAL} + 2(w(v_x, v_y) + w(v_x, v_z) + w(v_y, v_z)) \end{aligned}$$

Rearranging gives the following formula for computing OUTVAL:  $\text{OUTVAL} = D - \text{INVAL} - 2(w(v_x, v_y) + w(v_x, v_z) + w(v_y, v_z))$

**Example 2.** We continue to our example for Algorithm 3. Given the rooted bubble tree in Figure 2(b), we start our computation at the root  $b_3$ . Since  $b_3$  does not have a parent, we recurse down to its child  $b_2$ . For  $b = b_2$ , the vertices shared with its parent are  $\{v_x, v_y, v_z\} = t_2 = \{0, 1, 3\}$  and the remaining vertex is  $v = 2$ . Then, we initialize  $r[0] = w(0, 2)$ ,  $r[1] = w(1, 2)$ , and  $r[3] = w(3, 2)$  on Line 6. Next, we recurse down to  $b_2$ 's children,  $b_1$  and  $b_4$ . For  $b_1$ , the resulting  $r^*$  array should contain  $r^*[0] = w(0, 4)$ ,  $r^*[1] = w(1, 4)$ , and  $r^*[2] = w(2, 4)$  from the recursion on Line 8 because the shared vertices with its parent are  $t_1 = \{0, 1, 2\}$  and the remaining vertex is 4. On Lines 9–11, since  $v_x^* = 0 \in r$  and  $v_y^* = 1 \in r$ , we increment  $r[0]$  by  $w(0, 4)$  and  $r[1]$  by  $w(1, 4)$ . Since  $v_z^* = 2 \notin r$ , we do not process it. Now  $r[0] = w(0, 4) + w(0, 2)$  and  $r[1] = w(1, 4) + w(1, 2)$ . Similarly for  $b_4$ , we increment  $r[1]$  by  $w(1, 5)$  and  $r[3]$  by  $w(3, 5)$ . Now  $r[0] = w(0, 4) + w(0, 2)$ ,  $r[1] = w(1, 4) + w(1, 2) + w(1, 5)$  and  $r[3] = w(3, 2) + w(3, 5)$ .

When we get to Line 12 for  $b_2$ , the sum of  $r[0]$ ,  $r[1]$ , and  $r[3]$  is INVAL because it contains the edge weights from the three corners of  $t_2$  to its interior. OUTVAL is computed by summing the weights of all edges from  $t_2$  and then subtracting INVAL and the weight of  $t_2$ 's edges (once in each direction). In our example, we find that  $\text{INVAL} > \text{OUTVAL}$ , and so the edge is directed from  $b_3$  to  $b_2$  (Figure 2(c)). The other edges are directed similarly by comparing INVAL and OUTVAL.

### C. Assigning Vertices

A **converging bubble** is a bubble with only incoming edges, and no outgoing edges. Intuitively, converging bubbles are the “end points of a directional path that follows the strongest connections” [13], so they are considered the center of local clusters. The first level of clustering assigns each vertex to a unique converging bubble. If a vertex is in at least one converging bubble, then it is assigned to the converging bubble with the strongest attachment  $\chi(v, b) = \frac{\sum_{u \in b} w(u, v)}{3(|b| - 2)}$ , where  $3(|b| - 2)$  is the number of edges in the bubble  $b$ . For

TMFG, all bubbles have 6 edges, and so we can simplify it to  $\chi(v, b) = \sum_{u \in b} w(u, v)$ . For a vertex that is not in any converging bubble, it is assigned to the converging bubble that has the minimum mean average shortest path distance:

$$\bar{L}(v, b) = \text{mean}\{l_D(u, v) \mid u \in V_b^0 \wedge v \rightarrow b\}.$$

$v \rightarrow b$  means  $v$  is in some bubble that can reach  $b$  in the directed bubble tree, and  $l_D(u, v)$  is the shortest path distance from  $u$  to  $v$  in the TMFG using the dissimilarity measure  $D$  as the edge weights. For all bubbles  $b$ , we let  $V_b^0$  be the set of vertices in converging bubbles that have already been assigned to  $b$  from computing  $\chi$ . Let vertices assigned to the same converging bubble  $b$  in the procedure above be a **group**.

After this initial partitioning of vertices, we investigate how each of these groups is internally structured by performing a second level of clustering. This time we assign each vertex to a unique bubble, but not necessarily a converging bubble. A vertex  $v$  is assigned to the bubble  $b$  that maximizes a different attachment score

$$\chi'(v, b) = \frac{\sum_{u \in b} w(u, v)}{\sum_{u', v' \in b} w(u', v')}.$$

The pseudocode for our parallel DBHT algorithm is in Algorithm 4. On Lines 1–23, we show how the two assignments are computed. We first compute the necessary auxiliary data used in the vertex assignment computation. On Line 1, we compute the directions of bubble tree edges using Algorithm 3. On Line 2, we initialize the fields  $g$  and  $q$  for each vertex to  $(-\infty, -\infty)$  to prepare for the WRITEMAX operation on them.  $g$  is the group assignment and  $q$  is the bubble assignment. On Line 3, we initialize the set BB containing the bubble nodes. On Line 4, we obtain all of the converging bubbles CVGGB by using a parallel filter based on the out-degree of the tree nodes. On Lines 5–6, we run a BFS in the directed bubble tree for each bubble in parallel, and record for vertices in the bubbles which converging bubbles they can reach. This helps us to do the  $v \rightarrow b$  computation later. On Line 7, we compute all-pairs shortest paths in the TMFG by running Dijkstra's shortest path algorithm from each bubble node in parallel.

Now that we have obtained all necessary auxiliary data, we start to compute the assignments. On Lines 8–17, we compute the vertex assignments to converging bubbles, i.e., the groups. First, we compute the group assignment for vertices in at least one converging bubble by computing all of the attachment scores  $\chi$  in parallel and using concurrent WRITEMAXes to write the group assignment (Lines 8–11). On Line 12, we compute  $V_b^0$  for all converging bubbles  $b$  by first using a parallel integer sort to sort the vertices by group assignment, and then using a parallel filter to obtain the start and end indices of the groups in the sorted vertex array. Then on Lines 13–17, we assign the rest of the unassigned vertices to converging bubbles. Finally on Lines 18–23, we compute the bubble assignment in parallel using concurrent WRITEMINS.

**Example 3.** We continue with our example in Figure 2(c). Here we only have a single converging bubble  $b_2$  because this is the only bubble with no outgoing edges. Therefore, all

vertices are assigned to this converging bubble on Lines 8–17. Next, we look at what happens on Lines 18–23. Vertices 4, 5, and 6 are only in a single bubble, and so they are assigned to the bubble they are in. Vertices 0, 1, 2, and 3 are in multiple bubbles, and so we compute the  $\chi'$  score and assign each vertex to the bubble with the maximum  $\chi'$ . For example, for vertex 3, we will compute  $\chi'(3, b_2)$ ,  $\chi'(3, b_3)$ , and  $\chi'(3, b_4)$ . In our example, we find that  $\chi'(3, b_3)$  is the largest, and so we assign vertex 3 to  $b_3$ .

#### D. Complete Linkage

Now we describe how to obtain the dendrogram from the vertex assignments. At a high level, we build the complete-linkage hierarchy at three levels: intra-bubble, inter-bubble, and inter-group [13]. For the complete-linkage algorithm, the distance between two vertices  $u$  and  $v$  is their shortest path distance  $l_D(u, v)$ , and the distance between two set of vertices  $V_1$  and  $V_2$  is  $d(V_1, V_2) = \max\{l_D(u, v) \mid u \in V_1, v \in V_2\}$ . We use the parallel complete-linkage algorithm by Yu et al. [22].

The steps for building the dendrogram are shown on Lines 24–33 of Algorithm 4. On Line 24, we initialize  $n$  dendrogram nodes, one for each vertex in the TMFG. We define a **subgroup** to be the set of vertices in a group that belong to the same bubble. On Lines 25–28, for each subgroup we run the complete-linkage algorithm on the dendrogram nodes corresponding to vertices within the subgroup. The result is a dendrogram for each subgroup. We use subgroups so that vertices within the same group, but in different bubbles will not be processed together, which is consistent with the sequential algorithm [13]. Then, on Line 30, we run the complete-linkage algorithm on dendrogram nodes in each group. The result is a dendrogram for each group. Finally, on Line 31, we run the complete-linkage algorithm on the group dendrogram nodes to obtain the final dendrogram.

**Dendrogram Heights.** In conventional complete-linkage clustering, the dendrogram height is the distance between the two merged clusters. The dendrogram’s height requires that nodes closer to the dendrogram root (clusters merged later) have a height at least the height of nodes further away from the root. However, since we are concatenating three levels of dendrograms, the shortest path distance might not satisfy the height requirement. For example, the maximum shortest path distance might be larger between two vertices in a bubble than between two bubbles. As a result, we will have to re-assign heights to the dendrogram nodes. We use the same height assignment in the implementation by Aste [45], which ensures that all nodes that contain exactly one group are at the same height.

For inter-group dendrogram nodes merged on Line 31, the height is the number of converging bubbles in its descendants. This can be computed in a top-down traversal from the root  $\mathcal{Z}$ . Intra-bubble and inter-bubble dendrogram nodes that belong to the same converging bubble  $b^c$  have heights chosen from  $[\frac{1}{n_b-1}, \frac{1}{n_b-2}, \dots, \frac{1}{2}, 1]$ , where  $n_b$  is the number of vertices assigned to converging bubble  $b^c$ . Each bubble’s dendrogram nodes have a contiguous segment of the heights in the height

list. To obtain the dendrogram height for descendants of each  $\mathcal{Z}_{(b^c)}$ , we sort these dendrogram nodes such that nodes merged on Line 28 appear before nodes merged on Line 30. Nodes merged on Line 30 are sorted by the distance when they are merged. Nodes merged on Line 28 are sorted first by bubble assignment and then by the distance when they are merged. Now all dendrogram nodes in the same bubble and converging bubble are contiguous. We then assign heights  $[\frac{1}{n_b-1}, \frac{1}{n_b-2}, \dots, \frac{1}{2}, 1]$  to the dendrogram nodes in the sorted order. This sorted order ensures that the inter-group hierarchy is below the inter-bubble hierarchy in the dendrogram.

**Example 4.** In Figure 2(c), since  $\{2, 4\}$ ,  $\{0, 3, 6\}$ , and  $\{1, 5\}$  are all assigned to the same bubble, we first have three dendrograms  $\mathcal{Z}_{(b_2, b_1)}$ ,  $\mathcal{Z}_{(b_2, b_3)}$ , and  $\mathcal{Z}_{(b_2, b_4)}$ . Then we get a single dendrogram  $\mathcal{Z}_{b_2}$  by running complete linkage on the three dendrogram roots. In this example, we only have a single converging bubble, and so we are done. If there were multiple converging bubbles, our algorithm would run complete linkage on the group dendrogram roots, which would all be at the same height 1, to produce the final dendrogram. The resulting dendrogram is shown in Figure 2(d).

## VI. ANALYSIS

In this section, we analyze the complexity of our parallel algorithms using the work-span model defined in Section II.

We use  $\rho$  to denote the number of rounds required by the parallel TMFG algorithm, i.e., the number of times we iterate the while loop on Line 8 of Algorithm 1. We use  $\pi$  to denote the size of the PREFIX.

**Parallel TMFG.** First, we analyze the complexity of the parallel TMFG algorithm given in Section IV. The first step of computing the initial four vertices takes  $O(n^2)$  work and  $O(1)$  span. Initializing the GAINS array takes  $O(n)$  work and  $O(1)$  span because initially there are  $n - 4$  nodes in  $V$  and four faces in  $\mathcal{F}$ .

Now we analyze the while loop. Line 9 takes  $O(n \log n)$  work and  $O(\log n)$  span for parallel sorting. Lines 10–11 takes  $O(\pi \log \pi)$  time and  $O(\log \pi)$  span for parallel filtering and sorting. In the for-loop, each vertex in  $L$  contributes  $O(1)$  work and span on Lines 13–14. On Lines 15–16, we maintain a sorted list of vertices for each face. Every time we create a new face, we take  $O(n \log n)$  work to sort all vertices by their gains with respect to this face, and we maintain a linked list of sorted vertices for each face. Each vertex stores pointers to its positions in these linked lists. When we insert a vertex to a face, we remove the vertex from the linked lists of all other faces. Since there are  $O(n)$  faces, the total work of sorting is  $O(n^2 \log n)$  and the total work of removing vertices from linked lists is  $O(n^2)$  because each vertex is removed exactly once for each face. The span of this algorithm is  $O(\rho \log n)$  because the span of each round is  $O(\log n)$ .

**Bubble Tree Construction.** The bubble tree is constructed by the highlighted lines in Algorithm 1 (which calls Algorithm 2). Building the bubble tree takes  $O(n)$  work because we insert  $n$  vertices, and each insertion takes  $O(1)$  work. Since we update



the bubble tree on each round in  $O(1)$  span, the total span is  $O(\rho)$ .

**Directing Edges.** The directions of bubble tree edges are computed in Algorithm 3. Computing the bubble tree direction takes  $\Theta(n)$  work because for each bubble, we perform  $O(1)$  work (and span), and there are  $n-3$  bubbles in total. The span of computing the directions is asymptotically the same as the height of the bubble tree. This is because on each round of Algorithm 1 we only add tree edges adjacent to existing bubble node, so the height can increase by at most 2 in each round, and thus the span of this step is  $O(\rho)$ .

**Assigning Vertices.** Now, we give the complexity of the vertex assignment step in parallel DBHT algorithm (Lines 4–23 in Algorithm 4). Lines 4, 8–13, and 17–23 take  $O(n)$  work because there are  $O(n)$  bubbles, and each bubble has four vertices. Lines 14–17 takes  $O(n^2)$  work because there are at most  $O(n)$  converging bubbles, and each vertex can reach  $O(n)$  converging bubbles in the worst case. Lines 8–23 have span  $O(1)$  except Line 12 which has  $O(\log n)$  span for integer sorting. Lines 5–6 take  $O(n^2)$  work because we run  $O(n)$  BFS's. The span of running a BFS is  $O(\rho \log n)$  because the bubble tree has diameter  $O(\rho)$  [46].

The all-pairs shortest paths (APSP) on Line 7 can either be computed by running single-source shortest path (SSSP) from all nodes or can be computed by a dedicated APSP algorithm. The APSP and SSSP problems, both in the sequential and parallel settings, have been well-studied [47]–[56]. There are also APSP and SSSP algorithms specifically designed for planar graphs, which have better complexity than their general counterparts [57]–[69]. While Federickson [59] gives a sequential algorithm that computes APSP in  $O(n^2)$ , we are not aware of any parallel algorithms that achieve the same work with lower depth. Some examples for APSP on planar undirected graphs include an algorithm by Pan and Reif that runs in  $O(n^2 \log n)$  work and  $O(\log^2 n)$  span [70], and an algorithm by Henzinger et al. that runs in  $O(n^{7/3} \log n \log(nD))$  work and  $O(n^{2/3} \log^{7/3} n \log(nD))$  span, where  $D$  is the sum of the absolute values of the edge weights [57]. Let the work of the APSP algorithm used be  $W_{APSP}$  and the span be  $S_{APSP}$ .

**Complete Linkage.** Finally, we consider the complete linkage step. Let  $\rho_l$  be the number of rounds used by the parallel complete linkage algorithm [22]. The span is  $O(\rho_l \log n)$  and the work is  $O(\rho_l n^2)$  in the worst case. Yu et al. [22] show that in practice, the running time of the complete linkage algorithm is close to  $O(n^2)$ .

**Summary.** In total, TMFG takes  $O(n^2 \log n)$  work and  $O(\rho \log n)$  span and DBHT takes  $O(W_{APSP} + \rho_l n^2)$  work and  $O(S_{APSP} + \rho + \rho_l \log n)$  span. Compared to Song et al. [13], our algorithm has a lower span in all steps and lower work in bubble tree construction and direction computations. Song et al. take  $O(n^2)$  work for bubble tree construction and direction computations, while our algorithm takes  $O(n)$  work. We show in Section VII that all of our steps are faster. The fastest part of their DBHT code was the APSP computation, and all other parts were up to orders of magnitude slower; we have significantly reduced the running times of all parts of the

DBHT code, and now our bottleneck is APSP.

## VII. EXPERIMENTS

**Testing Environment.** We perform experiments on a c5.24xlarge machine on Amazon EC2, with 2 Intel Xeon Platinum 8275CL (3.00GHz) CPUs for a total of 48 hyper-threaded cores, and 192 GB of RAM. By default, we use all cores with hyper-threading. For the C++ code tested, we use the g++ compiler (version 7.5) with the `-O3` flag, and use ParlayLib [71] for parallelism in our code.

- **PMFG-DBHT** is the existing sequential PMFG and DBHT implementation in MATLAB, which we obtained online [45]. The bottleneck of PMFG-DBHT is in constructing the PMFG. We tried implementing the same PMFG construction algorithm in C++, but found it slower than the MATLAB implementation, and so we report the MATLAB runtime.
- **SEQ-TDBHT** is the state-of-art implementation of sequential TMFG and DBHT in MATLAB, which we obtained online [45].
- **PAR-TDBHT** is our implementation of parallel TMFG and DBHT in C++. We tested prefixes of size 1, 2, 5, 10, 30, 50, and 200.
- **COMP** and **AVG** are the parallel C++ complete and average linkage implementations, respectively, by Yu et al. [22].
- **K-MEANS** is the scalable  $k$ -means++ [72] implementation [73] in C, parallelized using OpenMPI.
- **K-MEANS-S** is the parallel  $k$ -means++ implementation from scikit-learn with a preprocessing step that computes a spectral embedding, which constructs its affinity matrix using a nearest-neighbors graph [74]. We choose the number of nearest neighbors that gives the best clustering quality. If the true number of clusters is  $c$ , then we project the original data onto the  $c$ -dimensional space, which we find to be a good heuristic for finding good clusters. We also tried a C++ implementation using Eigen [75], but found the scikit-learn version to be faster.

**Evaluation.** We evaluate the clustering quality using the Adjusted Rand Index (ARI) [11] and Adjusted Mutual Information (AMI) [76] scores. In our experiments, AMI showed similar trends as ARI, and so we only show the plots for ARI. Let  $n_{ij}$  be the number of objects in the ground truth cluster  $i$  and the cluster generated by the algorithm  $j$ ,  $n_{i*}$  be  $\sum_j n_{ij}$ ,  $n_{*j}$  be  $\sum_i n_{ij}$ , and  $n$  be  $\sum_i n_{i*}$ . The ARI is computed as

$$\frac{\sum_{i,j} \binom{n_{ij}}{2} - [\sum_i \binom{n_{i*}}{2} \sum_j \binom{n_{*j}}{2}] / \binom{n}{2}}{\frac{1}{2} [\sum_i \binom{n_{i*}}{2} + \sum_j \binom{n_{*j}}{2}] - [\sum_i \binom{n_{i*}}{2} \sum_j \binom{n_{*j}}{2}] / \binom{n}{2}}.$$

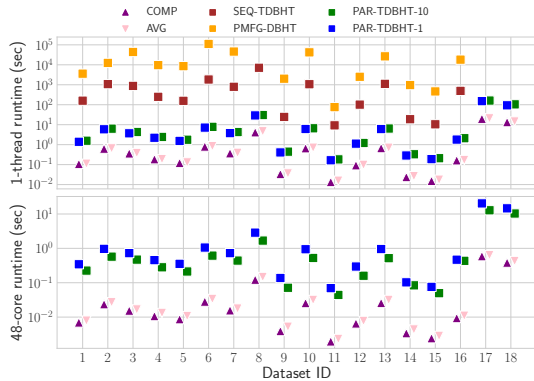
The ARI score is 1 for a perfect match, and its expected value is 0 for random assignments.

When computing the ARI for the hierarchical methods, we cut the dendrogram such that the number of resulting clusters is the same as the number of ground truth clusters. For the  $k$ -means methods, we set  $k$  to be equal to the number of ground truth clusters.

**Data sets.** We show results on 18 data sets from the UCR Time Series Classification Archive [77], including the three

**TABLE II:** Summary of UCR data sets used in the experiments.  $n$  is the number of objects, and  $L$  is the length or size of the object.

ID	Name	$n$	$L$	# of classes
1	Mallat	2400	1024	8
2	UWaveGestureLibraryAll	4478	945	8
3	NonInvasiveFetalECGThorax2	3765	750	42
4	MixedShapesRegularTrain	2925	1024	5
5	MixedShapesSmallTrain	2525	1024	5
6	ECG5000	5000	140	5
7	NonInvasiveFetalECGThorax1	3765	750	42
8	StarLightCurves	9236	84	2
9	HandOutlines	1370	2709	2
10	UWaveGestureLibraryX	4478	315	8
11	CBF	930	128	3
12	InsectWingbeatSound	2200	256	11
13	UWaveGestureLibraryY	4478	315	8
14	ShapesAll	1200	512	60
15	SonyAIBORobotSurface2	980	65	2
16	FreezerSmallTrain	2878	301	2
17	Crop	19412	46	24
18	ElectricDevices	16160	96	7



**Fig. 3:** Time (seconds) of different methods on UCR data sets. The top plot shows run times on a single thread and the bottom plot shows run times on 48 cores with hyper-threading. PMFG-DBHT and SEQ-TDBHT timed out for data sets 17 and 18. PMFG-DBHT also timed out for data set 8.

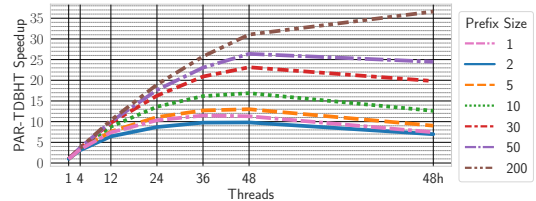
largest data sets Crop, ElectricDevices, and StarLightCurves (we removed duplicate points from these three data sets). The data sets are summarized in Table II. Since the UCR Archive is designed for classification tasks, not all data sets are suitable for clustering. We choose the two largest data sets as well as 16 data sets which have ARI scores of at least 0.2 for K-MEANS.

We also collected the closing daily prices of 1614 US stocks between Jan. 1, 2013 and Jan. 1, 2019 (1761 trading days) using the Yahoo Finance API. We used the Industry Classification Benchmark (ICB) to obtain ground truth clusters.

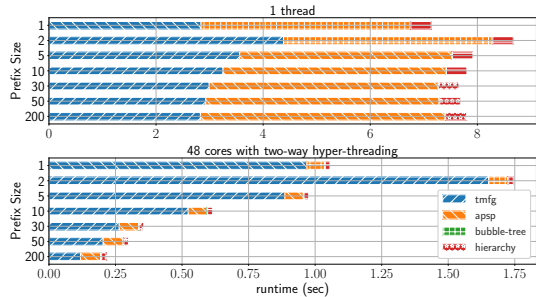
We used the Pearson correlation coefficient  $p$  for the similarity measure and  $d = \sqrt{2(1-p)}$  for the dissimilarity measure [78]. For normalized and zero-mean vectors,  $d$  is the same as the squared Euclidean distance.

### A. Runtime

We show the running times of all hierarchical clustering algorithms and data sets in Figure 3 (sequential times are in the top plot and parallel times are in the bottom plot). PAR-TDBHT-1 is PAR-TDBHT with a prefix of size 1 and PAR-TDBHT-10 is PAR-TDBHT with a prefix of size 10 (we chose this prefix size for most experiments as it gives a good tradeoff between speed and cluster quality). Since PMFG-DBHT and SEQ-TDBHT are



**Fig. 4:** Self-relative parallel speedup vs. thread counts for PAR-TDBHT with different prefix sizes on the Crop data set. "48h" indicates 48 cores with two-way hyper-threading. The speedups of some data sets decrease when hyper-threading because there is not enough work and the overhead of hyper-threading is high relative to the work of the algorithm.



**Fig. 5:** Breakdown of runtime across different steps of our algorithm on ECG5000. "tmfg" corresponds to TMFG construction (Algorithm 1); "aspp" corresponds to the all-pairs shortest paths computation; "bubble-tree" corresponds to computing the direction of bubble tree edges and assigning vertices to bubbles; and "hierarchy" corresponds to running the complete-linkage subroutine.

sequential, we only include it in the top plot.

We see that PMFG-DBHT and SEQ-TDBHT are both orders of magnitude slower than all other methods. PMFG-DBHT is 458–15586x slower than PAR-TDBHT-1 and 414–14254x slower than PAR-TDBHT-10 on a single thread. SEQ-TDBHT is 56–276x slower than PAR-TDBHT-1 and 50–235x slower than PAR-TDBHT-10 on a single thread. On 48 cores with hyper-threading, SEQ-TDBHT is 136–2483x slower than PAR-TDBHT-1 and 226–4487x slower than PAR-TDBHT-10. We discuss the reasoning for this speedup in the Runtime Decomposition section below. PAR-TDBHT-1 and PAR-TDBHT-10 are slower than AVG and COMP, which is expected because DBHT uses complete-linkage clustering as a subroutine. However, we see in Section VII-B that PAR-TDBHT-1 and PAR-TDBHT-10 give significantly better clusters than AVG and COMP on most data sets.

K-MEANS-S and K-MEANS are not included because they do not generate a dendrogram, and one would need to run the algorithm multiple times to obtain clusters of different scales. On average across the data sets, one run of K-MEANS is 1.82x faster than PAR-TDBHT-1 and 1.31x faster than PAR-TDBHT-10 on 48 cores with hyper-threading; and one run of K-MEANS-S is 10.33x slower than PAR-TDBHT-1 and 16.23x slower than PAR-TDBHT-10. We used the 12-thread times for K-MEANS-S because it became slower past 12 threads due its parallel overheads. Though we are slower than K-MEANS, K-MEANS is not hierarchical and also not deterministic.

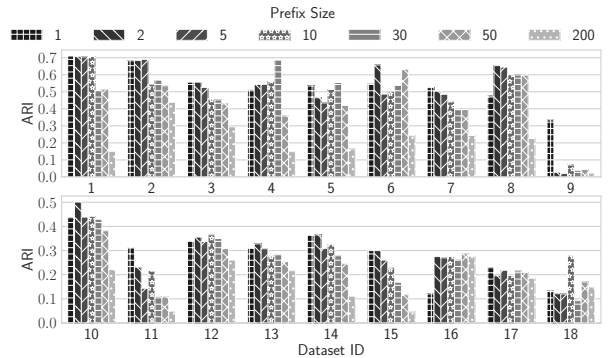
**Scalability with Thread Count.** In Figure 4, we show the scalability of our algorithm vs. thread count on the largest

Crop data set. PAR-TDBHT with a prefix size of 200 achieves a self-relative speedup of 36.6x on 48 cores with two-way hyper-threading. In general, a larger prefix size results in higher scalability, because more insertions in the TMFG construction can be processed in parallel. However, using a prefix of size 2 is actually slower than using a prefix of size 1, which corresponds to the exact TMFG, for the following reason. When we only have a prefix of size 1, our implementation only needs to find the best vertex-face pair to insert using a parallel maximum, but when the prefix size is larger than 1, the algorithm needs to first sort all vertex-face pairs, and then find a prefix of the sorted pairs to insert. A prefix of size 2 is small, and so there is not enough additional parallelism to offset the overheads of sorting. In our experiments, using a prefix of size 5 or larger gives similar or better runtimes than using exact TMFG. Our largest self-relative speedup is on StarLightCurves (data set 8), where PAR-TDBHT with a prefix size of 200 achieves a speedup of 41.57x.

**Scalability with Data Size.** We observe that on our data sets, the PAR-TDBHT runtimes scale with the data size  $n$  approximately as a function of  $O(n^{2.22})$  on a single thread and  $O(n^{1.79})$  on 48 cores with two-way hyper-threading. The scaling for parallel runtimes is better than for the single-threaded runtimes because we get more parallel speedups for larger values of  $n$ .

**Runtime Decomposition.** We show the breakdown of runtime across different steps of our algorithm in Figure 5 on the ECG5000 data set (the different steps are described in the figure caption). Sequentially, the majority of the runtime is in the "tmfg" and "apsp" steps, while in parallel the majority of the runtime is in the "tmfg" step, especially for small prefix sizes where TMFG construction has limited parallelism. When the prefix size is larger, the runtime of the "tmfg" step is significantly shorter. Our "bubble-tree" for TMFG is very fast and the runtime is too small to be visible in the plot.

For comparison, SEQ-TDBHT requires 628s for "tmfg", 9s for "apsp", 69s for "bubble-tree", and 1136s for "hierarchy" on the same data set, and thus all steps of PAR-TDBHT, even on one thread, are faster than SEQ-TDBHT. Our "tmfg" step is faster because we use optimized data structures to update the gain table that do not require looping over all faces. The "apsp" step is faster because we use a different graph data structure. SEQ-TDBHT uses Johnson’s algorithm from the Boost Graph Library, while we used the faster Dijkstra’s algorithm from the same library. We also tried Dijkstra’s algorithm for SEQ-TDBHT and found that in MATLAB’s Boost Graph Library, Dijkstra’s algorithm is on average 16% slower than Johnson’s algorithm (possibly due to having to convert the MATLAB graph to C++ format many more times). Our "bubble-tree" step is faster because we optimized it for the TMFG topology. While this step is much slower than "apsp" in SEQ-TDBHT, its time becomes negligible in PAR-TDBHT. Our "hierarchy" step is faster because we use the optimized complete-linkage algorithm by Yu et al. [22], whereas SEQ-TDBHT uses a less efficient implementation. Across all data sets, our "tmfg" step is 15.62–9629x faster than SEQ-TDBHT and the remaining



**Fig. 6:** Clustering quality (ARI score) of PAR-TDBHT. Different shades represent different prefix sizes.

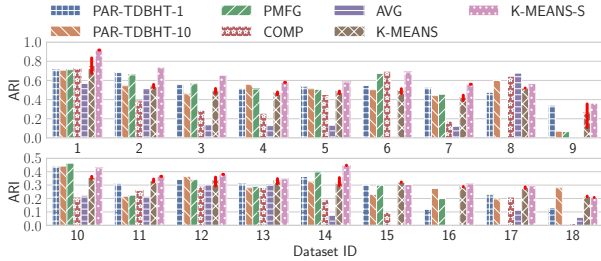
steps are together 1278–14049x faster than SEQ-TDBHT on 48 cores with hyper-threading. On a single thread, our "tmfg" step is 12.19–319.56x faster than SEQ-TDBHT and the remaining steps are together 59.42–431.17x faster than SEQ-TDBHT. In PAR-TDBHT, TMFG and APSP take most of the execution time when running sequentially, which means that the running time could potentially be improved by using a more sophisticated APSP implementation. However, in the baseline implementation [13], APSP takes only a small fraction of running time, and the bottlenecks are in other steps. This means that our algorithm has reduced the bottleneck of the original algorithm.

### B. Clustering Quality

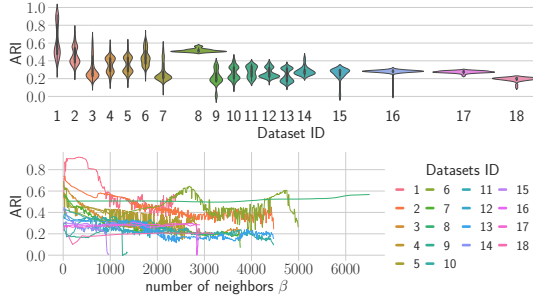
**Prefix Size and Clustering Quality.** Our prefix-based TMFG algorithm is able to produce filtered graphs with weight very close to, and sometimes even higher than, the sequential TMFG and PMFG algorithms. In our experiments, our prefix-based TMFG algorithm produces graphs with edge weight sums that are 92.1–100.3% of the edge weight sums produced by PMFG. If we only consider prefix sizes up to 50, then the edge weight sums are 97.1–100.3% of the edge weight sums produced by PMFG.

We present the ARI of using varying prefix sizes in Figure 6. We find that PAR-TDBHT with a prefix of size greater than 1 gives similar, and sometimes even better ARI than using a prefix of size 1 (which corresponds to using the exact TMFG). Generally, using a larger prefix size results in lower ARI, but sometimes a larger prefix size can result in better quality as the clusters could become less sensitive to noise. For smaller data sets (e.g., data sets 9, 11, and 15), the ARI degradation is larger. This is because the prefix is a larger percentage of all edges in the filtered graph. For larger data sets (e.g., data sets 2, 6, 8, 10, 13, 17, and 18), the ARI degradation is smaller.

**Hierarchical Methods.** We show the clustering quality of all of the methods and data sets in Figure 7. We see that PAR-TDBHT-1 and PAR-TDBHT-10 often generate higher-quality clusters than both of the other hierarchical clustering algorithms, COMP and AVG. On data sets where the number of ground truth clusters is very small, such as data sets 9, 15, and 16, COMP and AVG have low ARI scores. This is because COMP and AVG are sensitive to agglomeration decisions, which only use local information, and when these decisions are



**Fig. 7:** Clustering quality of different methods on UCR data sets. A few bars for COMP and AVG are hard to observe because their ARIs are close to 0. PMFG timed out for data sets 8, 17, and 18. The error bar on K-MEANS shows the range of ARI scores obtained when running it with different numbers of threads.



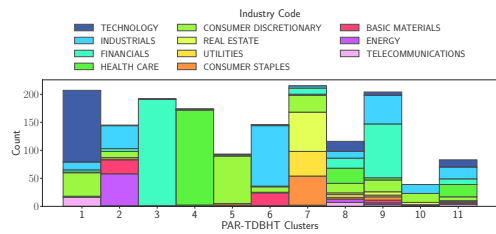
**Fig. 8:** ARI of K-MEANS-S with different number of nearest neighbors on each data set. The top plot shows the distribution of ARI scores for different values of  $\beta$ . The bottom plot shows the ARI scores vs. the value of  $\beta$ , demonstrating the oscillating behavior.

wrong, they lead to poor ARI scores. On the other hand, DBHT’s topological constraints (bubble and converging bubble) uses global information, which makes them less sensitive to wrong agglomeration decisions. DBHT can also suffer from the sensitivity of agglomeration decisions when the global information is not captured correctly (e.g., in data set 9).

**k-Means.** PAR-TDBHT-1 and PAR-TDBHT-10 generate clusters of similar quality to K-MEANS across all data sets. However, K-MEANS does not produce a dendrogram. To find clusters at different scales, we would have to run the algorithm multiple times, which would result in longer running times.

Using the spectral embedding preprocessing step boosts the K-MEANS method to achieve the best quality on most data sets. However, we show in Figure 8 that the quality of K-MEANS-S is highly sensitive to the choice of parameter  $\beta$  (the number of nearest neighbors) in many data sets. On the top plot of Figure 8, we see that for most data sets, using different values of  $\beta$  gives a wide range of ARI scores. On the bottom plot, we see that the ARI oscillates with different  $\beta$  for many data sets, and that the best  $\beta$  is very different for each data set. As a result, this parameter is hard to choose apriori. We also tried PAR-TDBHT on the embedded data sets, and found the ARI scores to be similar to those of the non-embedded data sets. For data sets 8, 17, and 18, we run out of memory when running spectral embedding with  $\beta$  values of 6600, 3000, and 4000 respectively. For all other data sets, we tested  $\beta$  ranging from 10 to  $n$ .

**Example: Clustering Stocks.** Figure 9 shows the clusters produced by PAR-TDBHT with a prefix size of 30 on the US stock data set. We preprocess the daily stock prices by



**Fig. 9:** Clustering result on the US stock data set compared to ground truth using PAR-TDBHT with a prefix of size 30.

computing the detrended daily log-return using the method by Musmeci et al. [6]. We then compute a spectral embedding of the normalized detrended daily log-returns. Finally, we compute the Pearson correlation of the embedded data and run PAR-TDBHT on the correlation matrix. We see that PAR-TDBHT is able to find structure and patterns in the stock data. It accurately clusters the “financial” stocks, “health care” stocks, and “consumer discretionary” stocks. There is also a cluster with mostly “technology” stocks, and a cluster with mostly “industrials” stocks. We also see that almost all of the “energy”, “utilities”, “consumer staples”, and “real estate” stocks are clustered together. Though for some clusters, the companies in the cluster do not come from a single industry in the ICB classification, they have commonalities between them. For example, cluster 7 contains both consumer discretionary and consumer staple stocks—many companies in the consumer discretionary industry are restaurants and wholesale, and many companies in the consumer staple industry are food and drink companies, which are related. In practice, stock clustering can be used in portfolio optimization and risk hedging [6], [79].

The ARI score of this clustering is 0.36. In comparison, the ARI score of the exact TMFG clustering is 0.28.

**Time and Quality Trade-off.** PAR-TDBHT provides a good trade-off between runtime and clustering quality. It is faster than the K-MEANS-S method, and is able to produce clusters of similar quality. Although PAR-TDBHT is slower than AVG and COMP, its clustering quality is more stable and in most cases better. Furthermore, it is able to finish within 30 seconds for all of the data sets.

## VIII. CONCLUSION

We designed new parallel algorithms for constructing TMFGs and DBHTs. We showed that our algorithms are usually able to produce better clusters than complete and average linkage clustering. We believe that our implementations will be a useful addition to the toolkit of data scientists who need to perform clustering efficiently and accurately.

## ACKNOWLEDGEMENTS

This research is supported by DOE Early Career Award #DE-SC0018947, NSF CAREER Award #CCF-1845763, Google Faculty Research Award, Google Research Scholar Award, FinTech@CSAIL Initiative, DARPA SDH Award #HR0011-18-3-0007, and Applications Driving Architectures (ADA) Research Center, a JUMP Center co-sponsored by SRC and DARPA.

## REFERENCES

- [1] J. Ruan, A. K. Dean, and W. Zhang, *BMC Systems Biology*, vol. 4, no. 1, p. 8, 2010.
- [2] R. N. Mantegna, "Hierarchical structure in financial markets," *The European Physical Journal B - Condensed Matter and Complex Systems*, vol. 11, no. 1, pp. 193–197, 1999.
- [3] M. Tumminello, C. Coronnello, F. Lillo, S. Micciche, and R. Mantegna, "Spanning trees and bootstrap reliability estimation in correlation-based networks," *I. J. Bifurcation and Chaos*, vol. 17, pp. 2319–2329, 07 2007.
- [4] M. Tumminello, T. Aste, T. Di Matteo, and R. Mantegna, "A tool for filtering information in complex systems," *PNAS*, vol. 102, pp. 10421–10426, 08 2005.
- [5] G. P. Massara, T. di Matteo, and T. Aste, "Network filtering for big data: Triangulated maximally filtered graph," *J. Complex Networks*, vol. 5, pp. 161–178, 2017.
- [6] N. Musmeci, T. Aste, and T. Di Matteo, "Relation between financial market structure and the real economy: comparison between clustering methods," *PLoS One*, vol. 10, no. 3, p. e0116201, 2015.
- [7] G.-J. Wang, C. Xie, and S. Chen, "Multiscale correlation networks analysis of the us stock market: a wavelet analysis," *Journal of Economic Interaction and Coordination*, vol. 12, no. 3, pp. 561–594, 2017.
- [8] P. T.-W. Yen and S. A. Cheong, "Using topological data analysis (tda) and persistent homology to analyze the stock markets in singapore and taiwan," *Frontiers in Physics*, p. 20, 2021.
- [9] W.-M. Song and B. Zhang, "Multiscale embedded gene co-expression network analysis," *PLoS Computational Biology*, vol. 11, no. 11, p. e1004574, 2015.
- [10] M. J. Burton *et al.*, "Pathogenesis of progressive scarring trachoma in ethiopia and tanzania and its implications for disease control: two cohort studies," *PLoS Neglected Tropical Diseases*, vol. 9, no. 5, p. e0003763, 2015.
- [11] L. Hubert and P. Arabie, "Comparing partitions," *Journal of Classification*, vol. 2, no. 1, pp. 193–218, 1985.
- [12] J. Giffin, "Graph theoretic techniques for facilities layout," 1984.
- [13] W.-M. Song, T. Di Matteo, and T. Aste, "Hierarchical information clustering by means of topologically embedded graphs," *PLoS One*, vol. 7, p. e31929, 03 2012.
- [14] W.-M. Song, T. Di Matteo, and T. Aste, "Nested hierarchies in planar graphs," *Discrete Applied Mathematics*, vol. 159, no. 17, pp. 2135–2146, 2011.
- [15] J. Jaja, *Introduction to Parallel Algorithms*, 1992.
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms (4. ed.)*. MIT Press, 2022.
- [17] R. D. Blumofe and C. E. Leiserson, "Scheduling multithreaded computations by work stealing," *J. ACM*, vol. 46, no. 5, pp. 720–748, Sep. 1999.
- [18] U. Vishkin, "Thinking in parallel: Some basic data-parallel algorithms and techniques," 2010.
- [19] S. Rajasekaran and J. H. Reif, "Optimal and sublogarithmic time randomized parallel sorting algorithms," *SIAM J. Comput.*, vol. 18, no. 3, pp. 594–607, 1989.
- [20] J. Shun, G. E. Blelloch, J. T. Fineman, and P. B. Gibbons, "Reducing contention through priority updates," in *ACM Symposium on Parallelism in Algorithms and Architectures*, 2013, pp. 152–163.
- [21] D. Müllner, "fastcluster: Fast hierarchical, agglomerative clustering routines for R and Python," *Journal of Statistical Software*, vol. 53, no. 9, pp. 1–18, 2013.
- [22] S. Yu, Y. Wang, Y. Gu, L. Dhulipala, and J. Shun, "ParChain: A framework for parallel hierarchical agglomerative clustering using nearest-neighbor chain," *Proc. VLDB Endow.*, vol. 15, no. 2, p. 285–298, Oct 2021.
- [23] Y. Wang, S. Yu, Y. Gu, and J. Shun, "Fast parallel algorithms for euclidean minimum spanning tree and hierarchical spatial clustering," in *SIGMOD*, 2021, pp. 1982–1995.
- [24] L. Dhulipala, D. Eisenstat, J. Lacki, V. S. Mirrokni, and J. Shi, "Hierarchical agglomerative graph clustering in nearly-linear time," in *ICML*, 2021, pp. 2676–2686.
- [25] S. Rajasekaran, "Efficient parallel hierarchical clustering algorithms," *IEEE Transactions on Parallel and Distributed Systems*, vol. 16, no. 6, pp. 497–502, 2005.
- [26] C. F. Olson, "Parallel algorithms for hierarchical clustering," *Parallel Computing*, vol. 21, no. 8, pp. 1313–1325, 1995.
- [27] W. Hendrix, M. M. A. Patwary, A. Agrawal, W.-k. Liao, and A. Choudhary, "Parallel hierarchical clustering on shared memory platforms," in *International Conference on High Performance Computing*, 2012, pp. 1–9.
- [28] E. Dahlhaus, "Parallel algorithms for hierarchical clustering and applications to split decomposition and parity graph recognition," *Journal of Algorithms*, vol. 36, no. 2, pp. 205–240, 2000.
- [29] M. Dash, S. Petrutiu, and P. Scheuermann, "Efficient parallel hierarchical clustering," in *European Conference on Parallel Processing*, 2004, pp. 363–371.
- [30] Q. Mao, W. Zheng, L. Wang, Y. Cai, V. Mai, and Y. Sun, "Parallel hierarchical clustering in linearithmic time for large-scale sequence analysis," in *IEEE International Conference on Data Mining*, 2015, pp. 310–319.
- [31] R. Campello, D. Moulavi, A. Zimek, and J. Sander, "Hierarchical density estimates for data clustering, visualization, and outlier detection," *TKDD*, pp. 5:1–5:51, 2015.
- [32] V. Cohen-Addad, V. Kanade, F. Mallmann-Trenn, and C. Mathieu, "Hierarchical clustering: Objective functions and algorithms," *Journal of the ACM*, vol. 66, no. 4, pp. 1–42, 2019.
- [33] B. Wang, I. Rahal, and A. Dong, "Parallel hierarchical clustering using weighted confidence affinity," *International Journal of Data Mining, Modelling and Management*, vol. 3, no. 2, pp. 110–129, 2011.
- [34] S. Lattanzi, T. Lavastida, K. Lu, and B. Moseley, "A framework for parallelizing hierarchical clustering methods," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2020, pp. 73–89.
- [35] P. A. Katakai, M. R. Cappelle, L. R. Foulds, and H. J. Longo, "A new algorithm for the maximum-weight planar subgraph problem," *Anais do LII Simpósio Brasileiro de Pesquisa Operacional, João Pessoa-PB*, 2020.
- [36] I. H. Osman, B. Al-Ayoubi, and M. Barake, "A greedy random adaptive search procedure for the weighted maximal planar graph problem," *Computers & Industrial Engineering*, vol. 45, no. 4, pp. 635–651, 2003.
- [37] M. E. Dyer, L. R. Foulds, and A. M. Frieze, "Analysis of heuristics for finding a maximum weight planar subgraph," *European Journal of Operational Research*, vol. 20, no. 1, pp. 102–114, 1985.
- [38] R. Cimikowski, "An analysis of some heuristics for the maximum planar subgraph problem," in *ACM-SIAM Symposium on Discrete Algorithms*, 1995, p. 322–331.
- [39] P. Eades, L. Foulds, and J. Giffin, "An efficient heuristic for identifying a maximum weight planar subgraph," in *Combinatorial Mathematics IX*, 1982, pp. 239–251.
- [40] Z. Ma, Z. Kang, G. Luo, L. Tian, and W. Chen, "Towards clustering-friendly representations: Subspace clustering via graph filtering," in *ACM MM*, 2020, pp. 3081–3089.
- [41] N. Tremblay, G. Puy, P. Borgnat, R. Gribonval, and P. Vandergheynst, "Accelerated spectral clustering using graph filtering of random signals," in *ICASSP*, 2016, pp. 4094–4098.
- [42] T. Jebara and V. Shchogolev, "B-matching for spectral clustering," in *European Conference on Machine Learning*, 2006, pp. 679–686.
- [43] T. Jebara, J. Wang, and S.-F. Chang, "Graph construction and b-matching for semi-supervised learning," in *Proceedings of the International Conference on Machine Learning*, 2009, pp. 441–448.
- [44] Y. Emek, S. Kutten, M. Shalom, and S. Zaks, "Hierarchical b-matching," in *International Conference on Current Trends in Theory and Practice of Informatics*, 2021, pp. 189–202.
- [45] T. Aste, "DBHT," 2022, [MATLAB Central File Exchange]<https://www.mathworks.com/matlabcentral/fileexchange/46750-dbht>.
- [46] G. E. Blelloch and B. M. Maggs, "Parallel algorithms," in *The Computer Science and Engineering Handbook*, 1997, pp. 277–315.
- [47] U. Meyer and P. Sanders, "Δ-stepping: a parallelizable shortest path algorithm," *Journal of Algorithms*, vol. 49, no. 1, pp. 114–152, 2003.
- [48] X. Dong, Y. Gu, Y. Sun, and Y. Zhang, "Efficient stepping algorithms and implementations for parallel shortest paths," in *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2021, p. 184–197.
- [49] L. Dhulipala, G. Blelloch, and J. Shun, "Julienne: A framework for parallel graph algorithms using work-efficient bucketing," in *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2017.
- [50] P. N. Klein and S. Subramanian, "A randomized parallel algorithm for single-source shortest paths," *J. Algorithms*, vol. 25, no. 2, Nov. 1997.
- [51] E. Cohen, "Polylog-time and near-linear work approximation scheme for undirected shortest paths," *J. ACM*, vol. 47, no. 1, Jan. 2000.



- [52] B. Sinha, B. Bhattacharya, S. Ghose, and P. Srimani, "A parallel algorithm to compute the shortest paths and diameter of a graph and its VLSI implementation," *IEEE Transactions on Computers*, vol. C-35, no. 11, nov. 1986.
- [53] R. Seidel, "On the all-pairs-shortest-path problem," in *ACM Symposium on Theory of Computing*, 1992.
- [54] M. Papaefthymiou and J. Rodrigue, "Implementing parallel shortest-paths algorithms," in *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 1994.
- [55] K. Madduri, D. A. Bader, J. W. Berry, and J. R. Crobak, "An experimental study of a parallel shortest path algorithm for solving large-scale graph instances," in *Algorithms Engineering and Experiments (ALENEX)*, 2007, pp. 23–35.
- [56] A. Karczmarz and P. Sankowski, "A deterministic parallel apsp algorithm and its applications," in *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 255–272.
- [57] M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian, "Faster shortest-path algorithms for planar graphs," *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 3–23, 1997.
- [58] G. E. Pantziou, P. G. Spirakis, and C. D. Zaroliagis, "Efficient parallel algorithms for shortest paths in planar graphs," in *Scandinavian Workshop on Algorithm Theory*, 1990, pp. 288–300.
- [59] G. N. Frederickson, "Fast algorithms for shortest paths in planar graphs, with applications," *SIAM Journal on Computing*, vol. 16, no. 6, pp. 1004–1022, 1987.
- [60] J. L. Träff and C. D. Zaroliagis, "A simple parallel algorithm for the single-source shortest path problem on planar digraphs," in *International Workshop on Parallel Algorithms for Irregularly Structured Problems*, 1996, pp. 183–194.
- [61] P. Klein and S. Subramanian, "A linear-processor polylog-time algorithm for shortest paths in planar graphs," in *IEEE Symposium on Foundations of Computer Science*, 1993, pp. 259–270.
- [62] D. Z. Chen and J. Xu, "Shortest path queries in planar graphs," in *ACM Symposium on Theory of Computing*, 2000, pp. 469–478.
- [63] H. N. Djidjev, G. E. Pantziou, and C. D. Zaroliagis, "Computing shortest paths and distances in planar graphs," in *International Colloquium on Automata, Languages, and Programming*, 1991, pp. 327–338.
- [64] H. N. Djidjev, "Efficient algorithms for shortest path queries in planar digraphs," in *International Workshop on Graph-Theoretic Concepts in Computer Science*, 1996, pp. 151–165.
- [65] J. Fakcharoenphol and S. Rao, "Planar graphs, negative weight edges, shortest paths, and near linear time," *Journal of Computer and System Sciences*, vol. 72, no. 5, pp. 868–889, 2006.
- [66] P. Klein, S. Rao, M. Rauch, and S. Subramanian, "Faster shortest-path algorithms for planar graphs," in *ACM Symposium on Theory of Computing*, 1994, pp. 27–37.
- [67] P. N. Klein, S. Mozes, and O. Weimann, "Shortest paths in directed planar graphs with negative lengths: A linear-space  $O(n \log^2 n)$ -time algorithm," *ACM Trans. Algorithms*, vol. 6, no. 2, apr 2010.
- [68] S. Mozes and C. Wulff-Nilsen, "Shortest paths in planar graphs with real lengths in  $O(n \log^2 n / \log)$  time," in *European Symposium on Algorithms*, 2010, pp. 206–217.
- [69] P. Gawrychowski, S. Mozes, O. Weimann, and C. Wulff-Nilsen, "Better tradeoffs for exact distance oracles in planar graphs," in *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 515–529.
- [70] V. Pan and J. Reif, "Fast and efficient solution of path algebra problems," *Journal of Computer and System Sciences*, vol. 38, no. 3, pp. 494–510, 1989.
- [71] G. E. Blelloch, D. Anderson, and L. Dhulipala, "ParlayLib—a toolkit for parallel algorithms on shared-memory multicore machines," in *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2020, pp. 507–509.
- [72] B. Bahmani, B. Moseley, A. Vattani, R. Kumar, and S. Vassilvitskii, "Scalable k-means++," *Proc. VLDB Endow.*, vol. 5, no. 7, p. 622–633, mar 2012.
- [73] G. Thompson, "mpi-scalablekmeanspp," 2018, <https://github.com/gzt/mpi-scalablekmeanspp>.
- [74] M. Lucińska and S. T. Wierzchoń, "Spectral clustering based on k-nearest neighbor graph," in *IFIP CISIM*, 2012, pp. 254–265.
- [75] G. Guennebaud, B. Jacob *et al.*, "Eigen," 2010, <http://eigen.tuxfamily.org>.
- [76] N. X. Vinh, J. Epps, and J. Bailey, "Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance," *J. Mach. Learn. Res.*, vol. 11, p. 2837–2854, dec 2010.
- [77] H. A. Dau *et al.*, "The UCR time series classification archive," October 2018, [https://www.cs.ucr.edu/~eamonn/time\\_series\\_data\\_2018/](https://www.cs.ucr.edu/~eamonn/time_series_data_2018/).
- [78] G. Marti, F. Nielsen, M. Bińkowski, and P. Donnat, "A review of two decades of correlations, hierarchies, networks and clustering in financial markets," *Progress in Information Geometry*, pp. 245–274, 2021.
- [79] V. Tola, F. Lillo, M. Gallegati, and R. N. Mantegna, "Cluster analysis for portfolio optimization," *Journal of Economic Dynamics and Control*, vol. 32, no. 1, pp. 235–258, 2008.