DISTRIBUTED MAPPER

MUSTAFA HAJIJ, BASEM ASSIRI, AND PAUL ROSEN

ABSTRACT. The construction of Mapper has emerged in the last decade as a powerful and effective topological data analysis tool that approximates and generalizes other topological summaries, such as the Reeb graph, the contour tree, split, and joint trees. In this paper we study the parallel analysis of the construction of Mapper. We give a provably correct algorithm to distribute Mapper on a set of processors and discuss the performance results that compare our approach to a reference sequential Mapper implementation. We report the performance experiments that demonstrate the efficiency of our method.

1. INTRODUCTION AND MOTIVATION

The topology of data is one of the fundamental originating principle in studying data. Consider the classical problem of fitting data set of point in $\mathbb{R}^n$ using linear regression. In linear regression one usually assumes that data is almost distributed near a hyperplane in $\mathbb{R}^n$. See Figure 1 (a). If the data does not meet this assumption then the model chosen to fit the data may not work very well.

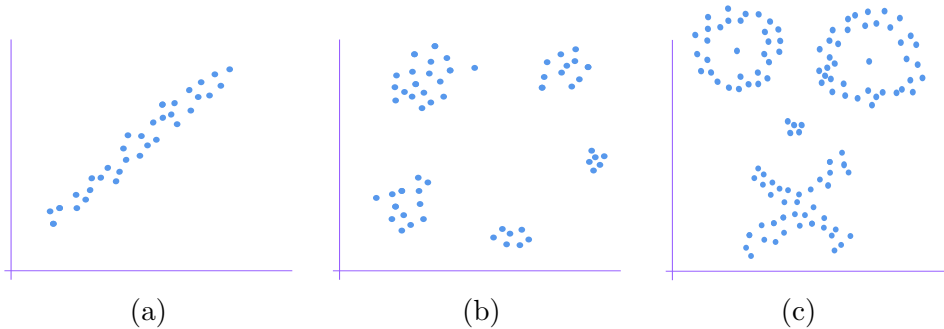(a)   (b)   (c)

FIGURE 1. (a) The linear shape of the data is a fundamental assumption underlying the linear regression method. (b) Clustering algorithms assume that the data is clustered in a certain way. (c) Data can come in many other forms and shapes.

On the hand if the clustering algorithms make the shape assumption that the data forms into clusters. See Figure 1 (b). Data can come in many other forms and shapes, see Figure 1 (c). It is the *shape of data* [9] that drives the meaning of the analytical methods and underlines the fundamental assumptions.

Topology is the field in Mathematics that rigorously defines and studies the notion of shape. Over the past two decades, topology has found enormous applications in data analysis and the application of topological techniques to study data has become is now considered a vibrant area of research called as *Topological Data Analysis* (TDA) [8–13, 17]. Many popular tools have been invented in the last two decades to study the shape of data, most notably *Persistence Homology* [20, 41] and the *construction of Mapper* [45]. Persistent Homology has been successfully used to study a wide range of data problems including three-dimensional structure of the DNA [21], financial networks [23], material science [28] and many other applications [39]. The construction of Mapper has emerged recently as a powerful and effective topological data analysis tool to solve

a wide variety of problems[31, 38] and it has been studied from multiple points of view [14, 19, 35]. Mapper works as a tool of approximation of a topological space by mapping this space via a "lens", or a sometimes called a filter, to another domain. One uses properties of the lens and the range to then extract a topological approximation of the original space. We give the precious notion in section 3. Mapper generalizes other topological summaries such as the Reeb graph, the contour tree, split, and joint trees. Moreover, Mapper is the core software developed by Ayasdi, a data analytic company whose main interest is promoting the usage of methods inspired by topological constructions in data science applications.

As the demand of analyzing larger data sets grows, it is natural to consider parallelization of topological computations. While there are numerous parallel algorithms that tackle the less general topological constructions, such as Reeb graph and contour tree, we are not aware of similar attempts targeting the parallel computation of Mapper in the literature. Our work here is an attempt to fill in this gap.

This article addresses the parallel analysis of the construction of Mapper. We give a provably correct algorithm to distribute Mapper on a set of processors and discuss the performance results that compare our approach to a reference sequential implementation for the computation of Mapper. Finally, we report the performance analysis experiments that demonstrate the efficiency of our method.

2. Prior Work

While there are numerous algorithms to compute topological constructions sequentially, the literature of parallel computing in topology is relatively young. One notable exception is parallelizing Morse-Smale complex computations [26, 44]. Parallelization of merge trees is studied in [24, 32, 40, 43]. Other parallel algorithms in topology include multicore homology computation [29] spectral sequence parallelization [30], distributed contour tree [33]. There are several other attempts to speed up the serial computation of topological constructions including an optimized Mapper sequential algorithm for large data [46], a memory efficient method to compute persistent cohomology [6], efficient data structure for simplicial complexes [2], optimized computation of persistent homology [15] and Morse-Smale complexes [25].

3. Preliminaries and Definitions

We start this section by recall basic notions from topology. For more details the reader is referred to standard texts in topology. See for instance [36]. All topological spaces we consider in this paper will be compact unless otherwise specified. An *open cover* of a topological space is a collection of open sets $\mathcal{U} = \{A_\alpha\}_{\alpha \in \mathcal{I}}$ such that $\cup_{\alpha \in \mathcal{I}} A_\alpha = X$. All covers in this article will consist of a finite number of sets unless otherwise specified. Given a topological space X with a cover $\mathcal{U}$, one may approximate this space via an abstract simplicial complex construction called the *nerve* of the cover $\mathcal{U}$. The nerve of a cover is a simplicial complex whose vertices are represented by the open sets the cover. Each non-empty intersection between two sets in the cover defines an edge in the nerve and each non-empty intersection between multiple sets defines higher order simplices. See Figure 4 for an illustrative example. Under mild conditions the nerve of a cover can be considered as an approximation of the underlying topological space. This is usually called the *Nerve Theorem* [22]. The Nerve Theorem plays an essential role in TDA: it gives a mathematically justified approximation of the topological space, being thought as the data under study, via simplicial complexes which are suitable for data structures and algorithms. In [45] Singh et al proposed using a continuous map $f : X \longrightarrow Z$ to construct a nerve of the space X. Instead of covering X directly, Singh et al suggested covering the codomain Z and then use the map f to pull back this cover to X. This perspective has multiple useful points of view. On one hand, choosing different maps on X can be used to capture different aspects of the space X. In this sense

the function f is thought of as a "lens" or a "filter" in which we view the space X. On the other hand, fixing the map f and choosing different covers for the codomain Z can be used to obtain multi-level resolution of the Mapper structure. This has been recently studied in details in [18, 19] and utilized to obtain a notion of persistence-based signature based on the definition of Mapper.

As we mentioned in the introduction, Mapper is related to Reeb graphs. In fact, it has been proven that under certain conditions Mapper and the Reeb graph give essentially the same topological summery for the underlying space [14, 35]. While Mapper is a more general topological construction than Reeb graph, we choose to motivate the main idea behind Mapper by illustrating its relationship with Reeb graph because the latter has found numerous applications. See for instance [5] and the references therein.

For the sake of illustration X will be considered to be a surface. Let $f : X \longrightarrow [a, b]$ be a Morse scalar function on X. The Reeb graph $R(X, f)$ of the Morse function f defined on X is constructed as follows. Define the equivalence relation $\sim$ on X by $x \sim y$ if and only if x and y belong to the same connected component of a level set $f^{-1}(c)$ for the same $c \in [a, b]$. The set $X/\sim$ with the standard quotient topology induced by the function $\pi : X \longrightarrow X/\sim$ is called the *Reeb graph* of f. See Figure 2 for an illustrative example of a Reeb graph on a surface. If X is an embedded surface, orientable and without a boundary then X can be recovered up to a homeomorphism from $R(X, f)$ as the boundary of an oriented 3-dimensional regular neighborhood of the graph $R(X, f)$ [27]. This is a perspective from which a Reeb graph may be considered as a topological summary of the original space. Other properties of Reeb graph can be found in [5, 16].
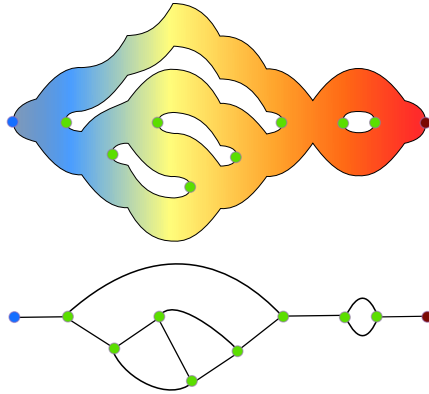

FIGURE 2. The Reeb graph $R(X, f)$ of surface X and a scalar function f. The scalar function is indicated by the color on the surface. The highlighted points on the surface are illustration of the critical points of f. These points correspond to the vertices of the graph $R(X, f)$.

Due to the successfulness of Reeb graph in various applications on meshes, there have been several attempts to give similar constructions on point clouds. In [37] Natali et el gave a method to mimic the Reeb graph construction for a point cloud in $\mathbb{R}^3$. Similar attempts were made in [48]. See also [4] and the references within. The problem with most of these methods is the lack of theoretical justification. Moreover it is not clear how these constructions can be generalized to a higher dimensional data. Besides being theoretically justified and applicable to data of any dimension, the Mapper construction gives a data structure that generalizes both Reeb graphs and merge/split trees and operates in practice on a metric space. We illustrate here how the construction of Mapper can realize in the special case of the Reeb graph of a space. We need first the following definition.
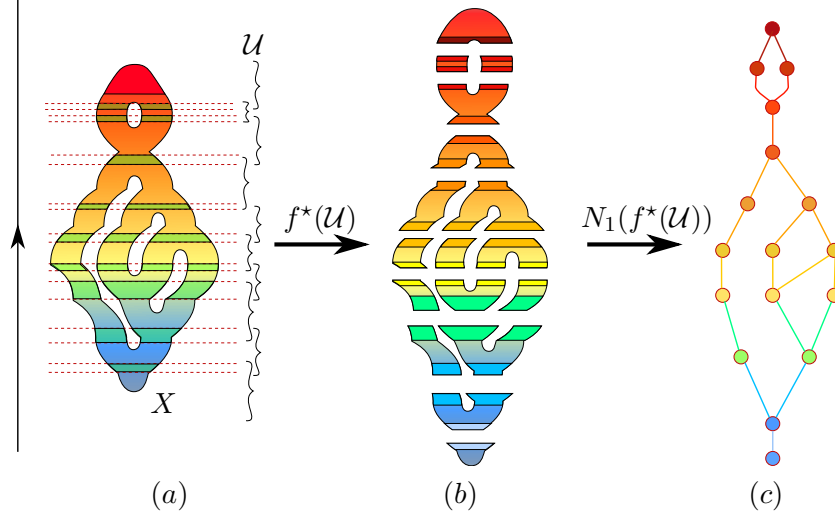
FIGURE 3. (a) Given a scalar function $f : X \longrightarrow [a, b]$ and an open cover $\mathcal{U}$ for $[a, b]$ we obtain an open cover $f^\star(\mathcal{U})$ for the space X by considering the inverse images of the elements of $\mathcal{U}$ under f. (b) The connected-components of the inverse images are identified as well as the intersection between these sets. (c) Mapper is defined as a graph whose vertices represent the connected component and whose edge represent the intersection between these components.

Definition 3.1. Let X be a topological space and let $\mathcal{U}$ be an open cover for X. The 1-**nerve** $N_1(\mathcal{U})$ of $\mathcal{U}$ is a graph whose nodes are represented by the elements of $\mathcal{U}$ and whose edges are the pairs A, B of $\mathcal{U}$ such that $A \cap B \neq \emptyset$.
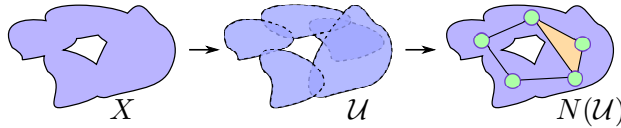

FIGURE 4. Each open set defines a vertex in the nerve simplicial complex. Each non-empty intersection between two sets define an edge and each non-empty intersection between multiple sets define higher order simplicies.

A scalar function f on X and a cover for the codomain $[a, b]$ of f give rise to a natural cover of X in the following way. Start by defining an open cover for the interval $[a, b]$ and take the inverse image of each open set to obtain an open cover for X. This is illustrated in Figure 3 (a). In other words if $\mathcal{U} = \{(a_1, b_1), ..., (a_n, b_n)\}$ is a finite collection of open sets that cover the interval $[a, b]$ then $f^\star(\mathcal{U}) := \{f^{-1}((a_1, b_1)), ..., f^{-1}((a_n, b_n))\}$ is an open cover for the space X. The open cover $f^\star(\mathcal{U})$ can now be used to obtain the 1-nerve graph $N_1(f^\star(\mathcal{U}))$. With an appropriate choice of the cover $\mathcal{U}$, the graph $N_1(f^\star(\mathcal{U}))$ is a version of the Reeb graph $R(X, f)$ [14, 35]. This is illustrated in Figure 3. Clearly, for a fixed function f the graph $N_1(f^\star(\mathcal{U}))$ depends on the choice of the cover $\mathcal{U}$ of the interval $[a, b]$. Figure 5 shows how different choices affect the graph G.

Observe that the different covers for $[a, b]$ give various "resolution" of the graph $N_1(f^\star(\mathcal{U}))$.

The idea of mapper presented in Definition 3.1 can be generalized to encompass a larger set of problems. One can replace the interval $[a, b]$ in Definition 3.1 by any parametization domain Z to obtain more sophisticated insights on the data X. This requires introducing the definition of a nerve of a cover of a topological space.
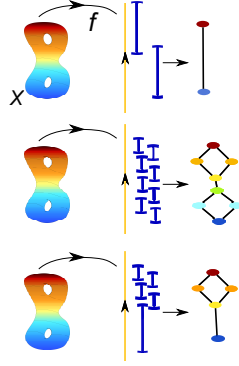
FIGURE 5. For the same function $f : X \longrightarrow [a,b]$, different choices of the cover $\mathcal{U}$ for $[a,b]$ induce different graphs. In other words, the 1-nerve graph depends on the choice of the cover of $[a,b]$.

Definition 3.2. Let X be a topological space and let $\mathcal{U}$ be a finite cover for X. The nerve of $\mathcal{U}$ is the abstract simplicial complex $N(\mathcal{U})$ whose vertices are the elements of $\mathcal{U}$ and whose simplicies are the finite subcollections $A_1, ..., A_k$ of $\mathcal{U}$ such that : $A_1 \cap ... \cap A_k \neq \emptyset$.

In this paper we will deal with nerves of multiple topological spaces simultaneously. For this reason we will sometimes refer to the nerve of a cover $\mathcal{U}$ of a space X by $N(X, \mathcal{U})$.

Figure 4 shows an illustrative example of nerve on a topological space X. We will denote the vertex in $N(\mathcal{U})$ that corresponds to an open set A in $\mathcal{U}$ by v_A.

Let $f : X \longrightarrow Z$ be a continuous map between two topological spaces X and Z. Let $\mathcal{U}$ be a finite cover of Z. The cover that consists of $f^{-1}(U)$ for all open sets $U \in \mathcal{U}$ will be called the *pullback* of $\mathcal{U}$ under f and will be denoted by $f^*(\mathcal{U})$. A continuous map $f : X \longrightarrow Z$ is said to be *well-behaved* if the inverse image of any path-connected set U in Z, consists of finitely many path-connected sets in X [18]. All maps in this paper will be assumed to be well-behaved.

Definition 3.3. Let $f : X \longrightarrow Z$ be a continuous map between two topological space X and Z. Let $\mathcal{U}$ be a finite cover for Z. The Mapper of f and $\mathcal{U}$, denoted by $M(f, \mathcal{U})$, is the nerve $N(f^*\mathcal{U})$.

3.1. **Some Graph Theory Notions.** Our construction requires a few definitions from graph theory. We include these notions here for completeness. See [3] for a more thorough treatment.

Definition 3.4. Let $G = (V, E)$ be a graph. Let $\sim$ be an equivalence relation defined on the node set V. The quotient graph of G with respect to the equivalence relation is a graph $G/\sim$ whose node set is the quotient set $V/\sim$ and whose edge set is $\{([u], [v]) | (u, v) \in E\}$.

For example consider the cyclic graph C_6 with $V = \{1, 2, 3, 4, 5, 6\}$ and edges $(1, 2), (2,3),...,(6, 1)$. Define the partition $\sim$ on V by $p_1 = \{1, 2\}$, $p_2 = \{3, 4\}$ and $p_3 = \{5, 6\}$. The quotient graph induced by $\sim$ is the cyclic graph C_3. See Figure 6.
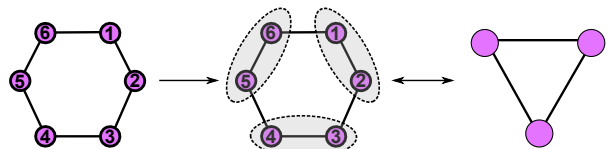

FIGURE 6. An example of a quotient graph.

We will also need the definition of disjoint union of two graphs. We will denote to the disjoint union of two sets A and B by $A \sqcup B$.

Definition 3.5. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs. The disjoint union of G_1 and G_2 is the graph $G_1 \sqcup G_2$ defined by $(V_1 \sqcup V_2, E_1 \sqcup E_2)$.

4. Distributed Computing of Topological Constructions

The idea of parallelizing the computation of Mapper lies in decomposing the space of interest into multiple smaller subspaces. The subspaces will be chosen to overlap on a smaller portion to insure a meaningful merging for the individual pieces. A cover of each space is then chosen. Each subspace along with its cover is then processed independently by a processing unit. The final stage consists of gathering the individual pieces and merging them together to produce the final correct Mapper construction on the entire space. Since Mapper is essentially a nerve of a certain cover we will start our discussion by laying our some theorems that will facilitate the computation of nerve of a topological space on two units. Our discussion here operates on a general topological space X. The algorithmic aspects will be discussed in later sections.

4.1. **Distributed Nerve Computation.** Let X be a topological space and let Y and Z be two open sets of X with $X = Y \cup Z$. See Figure 7. Let $\mathcal{U}$ and $\mathcal{V}$ be covers for the spaces Y and Z respectively. The cover $\mathcal{U} \cup \mathcal{V}$ is clearly a cover for X. Under which conditions can we build the 1-nerve $N_1(X, \mathcal{U} \cup \mathcal{V})$ from the 1-nerves $N_1(Y, \mathcal{U})$ and $N_1(Z, \mathcal{V})$? This is an important question for us because we want to be able to build the 1-nerve $N_1(X, \mathcal{U} \cup \mathcal{V})$ from the "building blocks" $N_1(Y, \mathcal{U})$ and $N_1(Z, \mathcal{V})$.
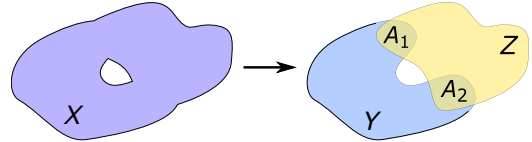

FIGURE 7. Decomposing a space X into two overlapping open sets Y and Z such that $X = Y \cup Z$ in order to process each space individually.

A natural initial choice is to consider the disjoint union of $N_1(Y, \mathcal{U})$ and $N_1(Z, \mathcal{V})$. The disjoint union $N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V})$ shares some vertices and edges with $N_1(X, \mathcal{U} \cup \mathcal{V})$ but these two graphs are not identical in general. The problems that could arise when trying to rebuild $N_1(X, \mathcal{U} \cup \mathcal{V})$ from the disjoint union of $N_1(Y, \mathcal{U})$ and $N_1(Z, \mathcal{V})$ are the following :

(1) There may exist some duplicate nodes in $N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V})$. This could occur when the same open set occur in both covers $\mathcal{V}$ and $\mathcal{U}$.

(2) The same problem can occur for edges as well: the same edge might exist in both $N_1(Y, \mathcal{U})$ and $N_1(Z, \mathcal{V})$.

(3) Finally, $N_1(X, \mathcal{U} \cup \mathcal{V})$ might contain an edge (v_A, v_B) that is coming from two open sets A, B with $A \in \mathcal{U}$ but $A \notin \mathcal{V}$ and $B \in \mathcal{V}$ but $B \notin \mathcal{U}$. Such an edge will not exist in $N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V})$.

The following definition addresses the previous problems and allows a correct reconstruction of $N_1(X, \mathcal{U} \cup \mathcal{V})$ from $N_1(X, \mathcal{U})$ and $N_1(X, \mathcal{V})$.

Definition 4.1. Let X be a topological space. Let Y, Z be two subspaces of X such that $X = Y \cup Z$ and $Y \cap Z \neq \emptyset$. Suppose that the intersection $Y \cap Z$ consists of a finite collection of maximal path-connected open sets $A_1, \cdots, A_n$. Let $\mathcal{U}, \mathcal{V}$ be open covers for Y and Z respectively. We say that the covers $\mathcal{U}$ and $\mathcal{V}$ are *nerve-consistent* if the following two conditions are satisfied :

(1) $A \in \mathcal{U} \cap \mathcal{V}$ for every A in $\{A_1, \cdots, A_n\}$.

(2) If $U \in \mathcal{U}$ and $V \in \mathcal{V}$ such that $U \cap V \neq \emptyset$, then $U = V \in \{A_1, \cdots, A_n\}$.

Note that it is necessary to impose the condition of maximal path-connectedness on the open sets of $\mathcal{A} := \{A_1, \cdots, A_n\}$ in the definition above. Otherwise, the opens sets that form $\mathcal{A}$ may not be unique and the nerve-consistency definition would have to depend on the open sets that form $\mathcal{A}$. In the definition above nerve-consistency depends only on the spaces Y and Z. Moreover, this condition implies $A \cap B = \emptyset$ for every $A, B \in \mathcal{A}$ with $A \neq B$ because otherwise that would contradict the fact that A and B are maximal path-connected. We now have the following theorem.

Theorem 4.2. Let X be a topological space. Let Y and Z be two open sets of X such that $X = Y \cup Z$ and $Y \cap Z \neq \emptyset$ intersecting on a finite union of maximal path-connected open sets $A_1 \cup \cdots \cup A_n$. Let $\mathcal{U}, \mathcal{V}$ be open covers for Y and Z that are nerve-consistent. Then $N_1(X, \mathcal{U} \cup \mathcal{V}) = (N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V}))/\sim$ where $v_A \sim v_B$ for all $A \in \mathcal{U}$ and $B \in \mathcal{V}$ such that $A = B$ and $A \in \mathcal{U} \cap \mathcal{V}$.

Proof. We show that both $N_1(X, \mathcal{U} \cup \mathcal{V})$ and $(N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V}))/\sim$ have essentially the same node and edge sets. Let $\mathcal{A} = \{A_i\}_{i=1}^n$. Every node v_A in $N_1(X, \mathcal{U} \cup \mathcal{V})$ corresponds to an open set A in $\mathcal{U} \cup \mathcal{V}$. If A is not in $\mathcal{A}$ then there is a unique corresponding node of v_A in $(N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V}))/\sim$. Otherwise $A \in \mathcal{A}$. In this case A must be in both $\mathcal{U}$ and $\mathcal{V}$, since the covers $\mathcal{U}$ and $\mathcal{V}$ are nerve-consistent. Hence there are two copies of the vertex v_A inside $N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V})$. By the definition of the equivalence relation $\sim$ these two nodes are identified in $N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V})/\sim$. Hence there is a one to one correspondence between the node sets of $N_1(X, \mathcal{U} \cup \mathcal{V})$ and $N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V})/\sim$. On the other hand let (v_A, v_B) be an edge in $N_1(X, \mathcal{U} \cup \mathcal{V})$. We distinguish four cases:

- If A and B are in $\mathcal{U}$ but not in $\mathcal{V}$ then there is a unique edge that corresponds to (v_A, v_B) in $N_1(Y, \mathcal{U}) \sqcup N_1(Z, \mathcal{V})/\sim$.
- The case when A and B are in $\mathcal{V}$ but not in $\mathcal{U}$ is similar to case (2).
- The case when $A \in \mathcal{U}$ and $B \in \mathcal{V}$. We show that no such edge can exist in $N_1(X, \mathcal{U} \cup \mathcal{V})$. Assume otherwise that (v_A, v_B) an edge in $N_1(X, \mathcal{U} \cup \mathcal{V})$ then we must have $A \cap B \neq \emptyset$. Since $\mathcal{U}$ and $\mathcal{V}$ are nerve-consistent then we have $A = B$.
- The case when when $A \in \mathcal{V}$ and $B \in \mathcal{U}$ is similar to the previous case.

$\square$

With this in mind, we now address the problem of distributed Mapper on two concurrent units.

4.2. **Distributed Mapper Computation.** Let $f : X \longrightarrow [a, b]$ be a continuous function. The construction of distributed Mapper on two units goes as follows:

(1) Choose an open cover for the interval $[a, b]$ that consists of exactly two sub-intervals A_1 and A_2 such that $A := A_1 \cap A_2 \neq \emptyset$. See Figure 8 (a).

(2) Choose open covers $\mathcal{U}_1$ and $\mathcal{U}_2$ for A_1 and A_2 respectively that satisfy the following conditions. First we want the intersection of the two coverings $\mathcal{U}_1$ and $\mathcal{U}_1$ to have only of the set A. Furthermore we do not want the covers $\mathcal{U}_1$ and $\mathcal{U}_2$ to overlap in anyway on any open set other than A. For this reason also have the following condition : if $U_1 \in \mathcal{U}_1$ and $U_2 \in \mathcal{U}_2$ such that $U_1 \cap U_2 \neq \emptyset$ then $U_1 \cap U_2 = A$. See Remark 4.4 below.

(3) We compute the Mapper construction on the covers $f^*(\mathcal{U}_i)$ for $i = 1, 2$. We obtain two graphs G_1 and G_2. See Figure 8 (b).

(4) We merge the graphs G_1, G_2 as follows. By the construction of A, $\mathcal{U}_1$ and $\mathcal{U}_2$, the set A presents in both covers $\mathcal{U}_i$, i=1,2. Let $C_1, ..., C_n$ be the path-connected components of $f^{-1}(A)$. Since A appears in both of the covers every connected component C_i in $f^{-1}(A)$ occurs in both graphs G_1 and G_2. In other words, the nodes $v_1, ..., v_n$ that corresponds to the components $C_1, ..., C_n$ occur in both G_1 and G_2 where each vertex v_i corresponds to the set C_i. The merge of the graph is done by considering the disjoint union $G_1 \sqcup G_2$ and then take the quotient of this graph by identifying the duplicate nodes $v_1, ..., v_k$ presenting in both G_1 and G_2. See Figure 8 (c).

The steps of the previous algorithm are summarized in Figure 8.

Remark 4.3. Note that the interval $[a, b]$ in the construction above can be replaced by any domain Y and the construction above remains valid. However for the purpose of this paper we restrict ourselves to the simplest case when $Y = [a, b]$.

Remark 4.4. Note that the covers $\mathcal{U}_1$ and $\mathcal{U}_2$ we chose for the sets A_1 and A_2 are nerve-consistent in the sense of Definition 4.1.
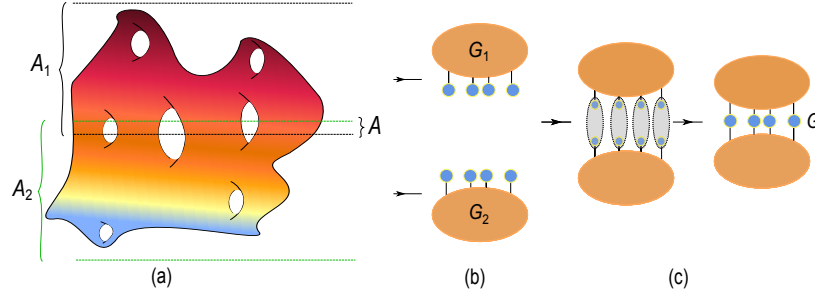
FIGURE 8. The steps of the distributed Mapper on two units. (a) The space X is decomposition based on a decomposition of the codomain (b) Each part is sent to a processing unit and the Mapper graphs are computed on the subspaces (c) The graphs are merged by identifying the corresponding the nodes.

It is immediate from the construction of the distributed Mapper on two units above that the covers $f^*(\mathcal{U}_1)$ and $f^*(\mathcal{U}_2)$ are nerve-consistent. We record this fact in the following Lemma.

Lemma 4.5. The covers $f^*(\mathcal{U}_1)$ and $f^*(\mathcal{U}_2)$ given in the construction above are nerve-consistent.

Lemma 4.5 and Theorem 4.2 imply immediately the following Corollary.

Corollary 4.6. $N_1(f^*(\mathcal{U}_1) \cup f^*(\mathcal{U}_2)) = (N_1(f^*(\mathcal{U}_1) \sqcup N_1(f^*(\mathcal{U}_2))/ \sim$ where $v_C \sim v_D$ for all $C \in f^*(\mathcal{U}_1)$ and $D \in f^*(\mathcal{U}_1)$ such that $C = D$ and $C \in f^{-1}(A_1 \cap A_2)$.

Now define an N-chain cover of $[a, b]$ to be a cover $\mathcal{U}$ of $[a, b]$ that consists of N open intervals $A_1, ..., A_N$ such that $A_{i,j} := A_i \cap A_j \neq \emptyset$ when $|i - j| = 1$ and empty otherwise. By convention, a 1-chain cover for an interval $[a, b]$ is any open interval that contains $[a, b]$.

5. THE DESIGN OF THE ALGORITHM

In this section we discuss the computational details of the distributed Mapper algorithm we discussed in the previous section from the topological perspective. Before we give our algorithm we recall quickly the reference sequential version.

5.1. **The Sequential Mapper Algorithm.** The serial Mapper algorithm can be obtained by a straightforward change of terminology of the topological mapper introduced in Section 3. To this end, the topological space X is replaced by the data under investigation. The lens, or the filter, f is chosen to reflect a certain property of the data. Finally, the notion of path-connectedness is replaced by an appropriate notion of clustering. This is algorithm is summarized in the Algorithm 1.

Algorithm 1: Sequential Mapper [45]

Input: A dataset X with a notion of metric between the data points;
a scalar function $f : X \longrightarrow \mathbb{R}^n$;
a finite cover $\mathcal{U} = \{U_1, ..., U_k\}$ of $f(X)$;
Output: A graph that represents $N_1(f^\star(\mathcal{U}))$.

1 For each set $X_i := f^{-1}(U_i)$, its clusters $X_{ij} \subset X_i$ are computed using the chosen clustering algorithm.;

2 Each cluster is considered as a vertex in the Mapper graph. Moreover we insert an edge between two nodes X_{ij} and X_{kl} whenever $X_{ij} \cap X_{ij} \neq \emptyset$;

We will refer the mapper graph obtained using Algorithm 1 by the *sequential Mapper*.

5.2. **The Main Algorithm.** We now give the details of the Distributed Mapper algorithm. To guarantee that the output of the distributed Mapper is identical to the sequential Mapper we need to do some processing on the cover that induces the final Mapper output. In distributed Mapper, we consider an N-chain cover $A_1, \cdots, A_N$ of the interval $[a, b]$ along with the their cover $\mathcal{U}_1, ..., \mathcal{U}_N$. The justification of the previous choice of covers will be given when we discuss the correctness of Algorithm 3. The details of the cover preprocessing are described in Algorithm 2.

Algorithm 2: Cover Preprocessing

Input: A point cloud X;
a scalar function $f : X \longrightarrow [a, b]$;
a set of N processors $(\mathcal{P})$;
Output: A collection of pairs $\{(A_i, \mathcal{U}_i)\}_{i=1}^N$ where $\{A_i\}_{i=1}^N$ is an N-chain cover of $[a, b]$ and $\mathcal{U}_i$ is a cover of A_i.

1 Construct an N-chain cover of $[a, b]$. That is, cover $[a, b]$ by N open intervals $A_1, \cdots, A_N$ such that $A_{i,j} := A_i \cap A_j \neq \emptyset$ when $|i - j| = 1$ and empty otherwise;

2 For each open set A_i construct an open cover $\mathcal{U}_i$. The covers $\{\mathcal{U}_i\}_{i=1}^N$ satisfy the following conditions: (1) $A_{i,i+1}$ is an open set in both coverings $\mathcal{U}_i$ and $\mathcal{U}_{i+1}$. In other words $\mathcal{U}_i \cap \mathcal{U}_{i+1} = \{A_{i,i+1}\}$ and (2) if $U_i \in \mathcal{U}_i$ and $U_{i+1} \in \mathcal{U}_{i+1}$ such that $U_i \cap U_{i+1} \neq \emptyset$ then $U_i \cap U_{i+1} = A_{i,i+1}$ for each $i = 1, ..., N - 1$;

After doing the preprocessing of the cover and obtaining the collection $\{(A_i, \mathcal{U}_i)\}_{i=1}^N$, every pair $(A_i, \mathcal{U}_i)$ is then mapped to a specific processor P_i which performs some calculations to produce a subgraph G_i. At the end, we merge the subgraphs into one graph G. The details of the algorithm are presented in Algorithm 3.

Algorithm 3: Distributed Mapper

Input: A point cloud X;
a scalar function $f : X \longrightarrow [a, b]$;
a set of N processors $(\mathcal{P})$;
a collection of pairs $\{(A_i, \mathcal{U}_i)\}_{i=1}^N$ obtained from the cover preprocessing algorithm;
Output: Distributed Mapper Graph.

1 **for** ($i \leftarrow 1$ *to* $i = N$) **do**
2 $\quad$ $P_i \leftarrow (A_i, \mathcal{U}_i)$; //Map each A_i, and its cover $\mathcal{U}_i$ to the processor P_i.
3 Determine the set of point $X_i \subset X$ that maps to A_i via f and run the sequential Mapper construction concurrently on the covers $(f|_{X_i})^*(\mathcal{U}_i)$ for $i = 1, .., N$. We obtain N graphs $G_1, ... G_N$. If $N = 1$, return the graph G_1;
4 Let $C_{j_1}^i, ..., C_{j_i}^i$ be the clusters obtained from $f^{-1}(A_{i,i+1})$. These clusters are represented by the vertices $v_{j_1}^i, ..., v_{j_i}^i$ in both G_i and G_{i+1} (each vertex v_k^i corresponds to the cluster C_k^i) by the choice of the coverings $\mathcal{U}_i$ and $\mathcal{U}_{i+1}$;
5 Merge the graphs $G_1, ..., G_N$ as follows. By the construction of $A_{i,i+1}$, $\mathcal{U}_i$ and $\mathcal{U}_{i+1}$, each one of the sets $f^*(\mathcal{U}_i)$ and $f^*(\mathcal{U}_{i+1})$ share the clusters $C_{j_k}^i$ in $f^*(A_{i,i+1})$. Hence $C_{j_k}^i$ is represented by a vertex in both graphs G_i and G_{i+1}. The merging is done by considering the of the disjoint union graph $G_1 \sqcup ... \sqcup G_N$ and then take the quotient of this graph that identifies the corresponding vertices in G_i and G_{i+1} for $1 \leq i \leq N - 1$.

5.3. **Correctness of the Algorithm.** The correctness of the previous algorithm follows basically from Corollary 4.6. That being said, we give here a detailed proof of the correctness that discusses the steps of the algorithms in details.

Proposition 5.1. The Distributed Mapper algorithm returns a graph identical to the sequential Mapper.

Proof. We will prove that the Distributed Mapper performs the computations on X and correctly produces a graph G that is identical to the graph obtained by the sequential Mapper algorithm using induction.

Denote by N to the number of units of initial partitions of interval I, which is the same number of processing units. If $N = 1$, then the Distributed Mapper works exactly like the sequential Mapper. In this case $A_1 = X$ and the single cover $\mathcal{U}_1$ for X is used to produce the final graph which Algorithm 3 returns at step (3).

Now assume the hypothesis is true on k unit. We show that it holds on $k+1$ units. In step (1) and (2) Algorithm 3 constructs a $k + 1$-chain cover for $[a, b]$ consisting of the open sets $A_1, ..., A_k, A_{k+1}$. Denote by $\mathcal{U}_i$ to the cover of A_i for $1 \leq i \leq k + 1$. We can run Algorithm 3 on the collection $\{(A_i, \mathcal{U}_i)\}_{i=1}^k$ and produce a sequential Mapper graphs G_i $1 \leq i \leq k$ in step (3). By the induction hypothesis, Algorithm 3 produces correctly a graph G' obtained by merging the sequential Mapper graphs $G_1, ..., G_k$. In other words the graph G' obtained from Algorithm 3 is identical to the graph obtain by running the sequential Mapper construction on the cover $\cup_i^k \mathcal{U}i$.

We want to show that combining G' and G_{k+1} using our algorithm produces a graph G that is identical to running the sequential Mapper on the covering consists of $\cup_i^{k+1}\mathcal{U}i$. Let $\mathcal{U}'$ be the union $\cup_i^k \mathcal{U}i$ and a denote by A' to the union $\cup_{i=1}^k A_i$. By the the construction of the covers $\{\mathcal{U}_i\}_{i=1}^{k+1}$ in step (2), $\mathcal{U}'$ covers A'. Moreover, the covers $\mathcal{U}'$ and $\mathcal{U}_{k+1}$ only share the open set $A' \cap A_{k+1}$. Moreover there are no intersections between the open sets of the cover $\mathcal{U}'$ and the open sets of the cover $\mathcal{U}_{k+1}$ except for $A' \cap A_{k+1}$. Since there is no intersection between the open sets of $\mathcal{U}'$ and $\mathcal{U}_{k+1}$ then there will be no creation of edges between the nodes induced from them and hence the computation of edges done on the first k processors are independent from the computation of edges done on the $k + 1$ possessor. Now we compare the node sets of the graphs G', G_{k+1} and the graph G. Recall that each node in a sequential Mapper is obtained by a connected component of an inverse image of an open set in the cover that defines the Mapper construction. Since the covers $\mathcal{U}'$ and $\mathcal{U}_{k+1}$ intersect at the open set $f^{-1}(A' \cap A_{k+1})$ then each connected component of $f^{-1}(A' \cap A_{k+1})$ corresponds to a node that exists in both graphs G' and G_{k+1}. This means that each connected component of $f^{-1}(A' \cap A_{k+1})$ is processed twice : one time on the first k processor and one time on the $k + 1$ processor and for each such component corresponds to a node in both G' and G_{k+1}. In step (5) the algorithm checks the graphs G' and G_{k+1} for node duplication and merge them according to their correspondence to produce the graph G. □

6. Performance of the Algorithm

In this section, we discuss the performance improvement obtained using the Distributed Mapper algorithm. Generally there are some systems and applications where parallelism cannot be applied on all data or processes. In this case, part of data can be processed in parallel, while the other should be sequential. This may happen because of the nature of data (e.g. dependencies), the natures of processors (e.g. heterogeneity) or some other factors. In this case we can rely on a well-known formula which is the *Amedahl's law* to calculate the *speedup* ratio upper bound that comes from parallelism and the improvement percentage [1]. The Amedahl's law is formulated as follows:

$$S = \frac{1}{(1 - part) + part/N},$$

where S is the theoretical speedup ratio, $part$ is the proportion of system or program that can be made in parallel, $1 - part$ is the proportion that remains sequential, and N is the number of processors.

Speedup				
N	part =0.25	part =0.50	(Distributed Mapper) part =0.94	part =0.99
10	1.29	1.82	6.49	9.17
100	1.33	1.98	14.40	50.25
1000	1.33	1.99	16.40	90.99

TABLE 1. Speedup calculations based on Amedahl's law, using different numbers of processors and different ratios of parallelism. Column four shows the speedup of the distributed Mapper with respect to the sequential Mapper

In Table 1 we use Amedahl's law to calculate the theoretic speedup ratios using different numbers of processors and different ratios of parallelism. The table shows that the speedup increases as a response of the increase in the number of processors and/or in the proportion of the program that can be made parallel. For instnace, based on the table, when $part = 0.25$ the speedup of the program increases by 1.33. Notice however that at some points the performance stops improving even if we increase the number of processors.

In the distributed Mapper, there are two computational pieces which are the clustering piece that is about 94% of the execution, and cover construction/merging subgraphs piece that is about 4% of the execution. Our algorithm makes the clustering piece completely in parallel while the cover construction/merging subgraphs piece is processed sequentially. Table 1 shows that the speedup of the Distributed Mapper achieves 6.49 when $N = 10$ and it goes up to 16.40 when $N = 1000$.

6.1. **Experimentations.** In this section we present practical results obtained using a Python implementation. We ran our experimentations on a Dell OptiPlex 7010 machine with 4-core i7-3770 Intel CPU @ 3.40GHz and with a 24GiB System Memory. The distributed Mapper algorithm was tested on different models and compared their run-time with a publicly available data available at [47]. The size of the point cloud data are shown in Table 2.

Data	Size
camel	21887
cat	7277
elephant	42321
horse	8431
face	29299
head	15941

TABLE 2. The size of the datasets used in our experiments.

The sequential Mapper algorithm relies on 3 inputs: the data X, the scalar function $f : X \longrightarrow [a, b]$ and the choice of cover $\mathcal{U}$ of $[a, b]$. The existing publicly available Mapper implementations, see for instance [34], do not satisfy the level of control that we require for the cover choice and so we relied on our own Mapper implementation. The clustering algorithm that we used to specify the Mapper nodes is a modified version of the DBSCAN [7].

Using our implementation on the data given in Table 2 we obtained a minimum speedup of 2.41, a maximum one of 3.35 and an average speedup of 2.95 on 4 cores. This gives 73% average parallel efficiency. This is detailed in Figure 9. The figure shows the results of distributed Mapper speedup

obtained using our experiments. The x-axis represents the number of processes while the y-axis shows the speedup. It is clear from the figure that the curves are increasing in a monotonic fashion as we increase the number of processes. Moreover, the experimental results matches the theoretical calculations that appears in Table 1.
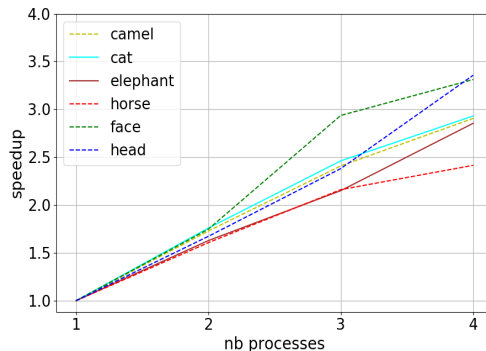

FIGURE 9. Speedups obtained by the distributed Mapper using number of processes that run concurrently.

7. CONCLUSION AND FUTURE WORK

In this work we gave a provably correct algorithm to distribute Mapper on a set of processors. Our algorithm relies on a divide an conquer strategy for the codomain cover which gets pulled back to the domain cover. This work has several potential directions of the work that we have not discussed here. For instance, the recursive nature of the main algorithm was implied throughout the paper but never discussed explicitly. On the other hand the algorithm can be utilized to obtain a multi-resolution Mapper construction. In other words, using this algorithm we have the ability to increase the resolution of Mapper for certain subsets of the data and decrease at others. This is potentially useful for interactive Mapper applications.

REFERENCES

[1] Gene M Amdahl, *Validity of the single processor approach to achieving large scale computing capabilities*, Proceedings of the april 18-20, 1967, spring joint computer conference, 1967, pp. 483–485.

[2] Ulrich Bauer, Michael Kerber, Jan Reininghaus, and Hubert Wagner, *Phat–persistent homology algorithms toolbox*, Journal of Symbolic Computation **78** (2017), 76–90.

[3] Lowell W Beineke and Robin J Wilson, *Topics in algebraic graph theory*, Vol. 102, Cambridge University Press, 2004.

[4] Silvia Biasotti, Leila De Floriani, Bianca Falcidieno, Patrizio Frosini, Daniela Giorgi, Claudia Landi, Laura Papaleo, and Michela Spagnuolo, *Describing shapes by geometrical-topological properties of real functions*, ACM Computing Surveys (CSUR) **40** (2008), no. 4, 12.

[5] Silvia Biasotti, Daniela Giorgi, Michela Spagnuolo, and Bianca Falcidieno, *Reeb graphs for shape analysis and applications*, Theoretical Computer Science **392** (2008), no. 1-3, 5–22.

[6] Jean-Daniel Boissonnat, Tamal K Dey, and Clément Maria, *The compressed annotation matrix: An efficient data structure for computing persistent cohomology*, Algorithmica **73** (2015), no. 3, 607–619.

[7] Tadeusz Caliński and Jerzy Harabasz, *A dendrite method for cluster analysis*, Communications in Statistics-theory and Methods **3** (1974), no. 1, 1–27.

[8] Erik Carlsson, Gunnar Carlsson, and Vin De Silva, *An algebraic topological method for feature identification*, International Journal of Computational Geometry & Applications **16** (2006), no. 04, 291–314.

[9] Gunnar Carlsson, *Topology and data*, Bulletin of the American Mathematical Society **46** (2009), no. 2, 255–308.

[10] Gunnar Carlsson, Tigran Ishkhanov, Vin De Silva, and Afra Zomorodian, *On the local behavior of spaces of natural images*, International journal of computer vision **76** (2008), no. 1, 1–12.

[11] Gunnar Carlsson and Facundo Mémoli, *Persistent clustering and a theorem of j. kleinberg*, arXiv preprint arXiv:0808.2241 (2008).

[12] Gunnar Carlsson and Afra Zomorodian, *The theory of multidimensional persistence*, Discrete & Computational Geometry **42** (2009), no. 1, 71–93.

[13] Gunnar Carlsson, Afra Zomorodian, Anne Collins, and Leonidas J Guibas, *Persistence barcodes for shapes*, International Journal of Shape Modeling **11** (2005), no. 02, 149–187.

[14] Mathieu Carrière and Steve Oudot, *Structure and stability of the 1-dimensional mapper*, arXiv preprint arXiv:1511.05823 (2015).

[15] Chao Chen and Michael Kerber, *Persistent homology computation with a twist*, Proceedings 27th european workshop on computational geometry, 2011.

[16] Kree Cole-McLaughlin, Herbert Edelsbrunner, John Harer, Vijay Natarajan, and Valerio Pascucci, *Loops in reeb graphs of 2-manifolds*, Proceedings of the nineteenth annual symposium on computational geometry, 2003, pp. 344–350.

[17] Anne Collins, Afra Zomorodian, Gunnar Carlsson, and Leonidas J Guibas, *A barcode shape descriptor for curve point cloud data*, Computers & Graphics **28** (2004), no. 6, 881–894.

[18] Tamal K Dey, Facundo Mémoli, and Yusu Wang, *Multiscale mapper: topological summarization via codomain covers*, Proceedings of the twenty-seventh annual acm-siam symposium on discrete algorithms, 2016, pp. 997–1013.

[19] Tamal K Dey, Facundo Memoli, and Yusu Wang, *Topological analysis of nerves, reeb spaces, mappers, and multiscale mappers*, arXiv preprint arXiv:1703.07387 (2017).

[20] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian, *Topological persistence and simplification*, Foundations of computer science, 2000. proceedings. 41st annual symposium on, 2000, pp. 454–463.

[21] Kevin Emmett, Benjamin Schweinhart, and Raul Rabadan, *Multiscale topology of chromatin folding*, Proceedings of the 9th eai international conference on bio-inspired information and communications technologies (formerly bionetics), 2016, pp. 177–180.

[22] Robert Ghrist, *Barcodes: the persistent topology of data*, Bulletin of the American Mathematical Society **45** (2008), no. 1, 61–75.

[23] Marian Gidea, *Topology data analysis of critical transitions in financial networks* (2017).

[24] Charles Gueunet, Pierre Fortin, Julien Jomier, and Julien Tierny, *Task-based augmented merge trees with fibonacci heaps*, Ieee symposium on large data analysis and visualization 2017, 2017.

[25] David Günther, Jan Reininghaus, Hubert Wagner, and Ingrid Hotz, *Efficient computation of 3d morse–smale complexes and persistent homology using discrete morse theory*, The Visual Computer **28** (2012), no. 10, 959–969.

[26] Attila Gyulassy, Valerio Pascucci, Tom Peterka, and Robert Ross, *The parallel computation of morse-smale complexes*, Parallel & distributed processing symposium (ipdps), 2012 ieee 26th international, 2012, pp. 484–495.

[27] Allen Hatcher and William Thurston, *A presentation for the mapping class group of a closed orientable surface*, Topology **19** (1980), no. 3, 221–237.

[28] Yasuaki Hiraoka, Takenobu Nakamura, Akihiko Hirata, Emerson G Escolar, Kaname Matsue, and Yasumasa Nishiura, *Hierarchical structures of amorphous solids characterized by persistent homology*, Proceedings of the National Academy of Sciences **113** (2016), no. 26, 7035–7040.

[29] Ryan H Lewis and Afra Zomorodian, *Multicore homology via mayer vietoris*, arXiv preprint arXiv:1407.2275 (2014).

[30] David Lipsky, Primoz Skraba, and Mikael Vejdemo-Johansson, *A spectral sequence for parallelized persistence*, arXiv preprint arXiv:1112.1245 (2011).

[31] PY Lum, G Singh, A Lehman, T Ishkanov, Mikael Vejdemo-Johansson, M Alagappan, J Carlsson, and G Carlsson, *Extracting insights from the shape of complex data using topology*, Scientific reports **3** (2013), 1236.

[32] Dmitriy Morozov and Gunther Weber, *Distributed merge trees*, Acm sigplan notices, 2013, pp. 93–102.

[33] Dmitriy Morozov and Gunther H Weber, *Distributed contour trees* (2012).

[34] Daniel Müllner and Aravindakshan Babu, *Python mapper: An open-source toolchain for data exploration, analysis, and visualization*, URL http://math. stanford. edu/muellner/mapper (2013).

[35] Elizabeth Munch and Bei Wang, *Convergence between categorical representations of reeb space and mapper*, arXiv preprint arXiv:1512.04108 (2015).

[36] James R Munkres, *Elements of algebraic topology*, Vol. 2, Addison-Wesley Menlo Park, 1984.

[37] Mattia Natali, Silvia Biasotti, Giuseppe Patanè, and Bianca Falcidieno, *Graph-based representations of point clouds*, Graphical Models **73** (2011), no. 5, 151–164.

[38] Monica Nicolau, Arnold J Levine, and Gunnar Carlsson, *Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival*, Proceedings of the National Academy of Sciences **108** (2011), no. 17, 7265–7270.

[39] Nina Otter, Mason A Porter, Ulrike Tillmann, Peter Grindrod, and Heather A Harrington, *A roadmap for the computation of persistent homology*, EPJ Data Science **6** (2017), no. 1, 17.

[40] Valerio Pascucci and Kree Cole-McLaughlin, *Parallel computation of the topology of level sets*, Algorithmica **38** (2004), no. 1, 249–268.

[41] Vanessa Robins, *Towards computing homology from finite approximations*, Topology proceedings, 1999, pp. 503–532.

[42] Alejandro Robles, Mustafa Hajij, and Paul Rosen, *The shape of an image: A study of mapper on images*, Accepted to VISAPP 2018 (2018).

[43] Paul Rosen, Junyi Tu, and L Piegl, *A hybrid solution to calculating augmented join trees of 2d scalar fields in parallel*, Cad conference and exhibition (accepted), 2017.

[44] Nithin Shivashankar, M Senthilnathan, and Vijay Natarajan, *Parallel computation of 2d morse-smale complexes*, IEEE Transactions on Visualization and Computer Graphics **18** (2012), no. 10, 1757–1770.

[45] Gurjeet Singh, Facundo Mémoli, and Gunnar E Carlsson, *Topological methods for the analysis of high dimensional data sets and 3d object recognition.*, Spbg, 2007, pp. 91–100.

[46] Václav Snášel, Jana Nowaková, Fatos Xhafa, and Leonard Barolli, *Geometrical and topological approaches to big data*, Future Generation Computer Systems **67** (2017), 286–296.

[47] Robert W Sumner and Jovan Popović, *Deformation transfer for triangle meshes*, ACM Transactions on Graphics (TOG) **23** (2004), no. 3, 399–405.

[48] Yijun Xiao, Paul Siebert, and Naoufel Werghi, *A discrete reeb graph approach for the segmentation of human body scans*, 3-d digital imaging and modeling, 2003. 3dim 2003. proceedings. fourth international conference on, 2003, pp. 378–385.

University of South Florida, Tampa, Florida
E-mail address: `mhajij@usf.edu`

Jazan University, Jazan City, Saudi Arabia
E-mail address: `babumussmar@jazanu.edu.sa`

University of South Florida, Tampa, Florida
E-mail address: `prosen@usf.edu`