

PARALLEL MAPPER

MUSTAFA HAJJ, 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 parallel algorithm to execute Mapper on a multiple 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. Mapper, Topological Data Analysis

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.

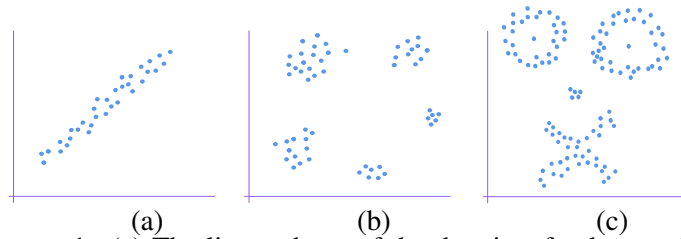


FIGURE 1. (a) The linear shape of the data is a fundamental assumption underlying the linear regression method. (b) Clustering algorithms assume that the data is clustered in a certain way. (c) Data can come in many other forms and shapes.

On the other hand, a clustering algorithm normally makes the shape assumption that the data falls 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* [7] that drives the meaning of these analytical methods and determines the successfulness of application of these methods on the data.

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 is now considered a vibrant area of research called as *Topological Data Analysis* (TDA) [6–11, 14]. Many popular tools have been invented in the last two decades

to study the shape of data, most notably *Persistent Homology* [17, 36] and the *construction of Mapper* [40]. Persistent Homology has been successfully used to study a wide range of data problems including three-dimensional structure of the DNA [18], financial networks [20], material science [24] and many other applications [34]. The construction of Mapper has emerged recently as a powerful and effective topological data analysis tool to solve a wide variety of problems [27, 33, 37] and it has been studied from multiple points of view [12, 16, 31]. 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 codomain 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 [23, 39]. Parallelization of merge trees is studied in [21, 28, 35, 38]. Other parallel algorithms in topology include multicore homology computation [25] spectral sequence parallelization [26], distributed contour tree [29]. There are several other attempts to speed up the serial computation of topological constructions including an optimized Mapper sequential algorithm for large data [41], a memory efficient method to compute persistent cohomology [4], efficient data structure for simplicial complexes [2], optimized computation of persistent homology [13] and Morse-Smale complexes [22].

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 [32]. 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 simplicies. See Figure 3 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* [19]. 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 [40] Singh et al proposed using a continuous map $f : X \rightarrow Z$ 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 [15, 16] and utilized to obtain a notion of persistence-based signature based on the definition of Mapper.

The Mapper construction is related to Reeb graphs. To illustrate relationship, we give the following definition.

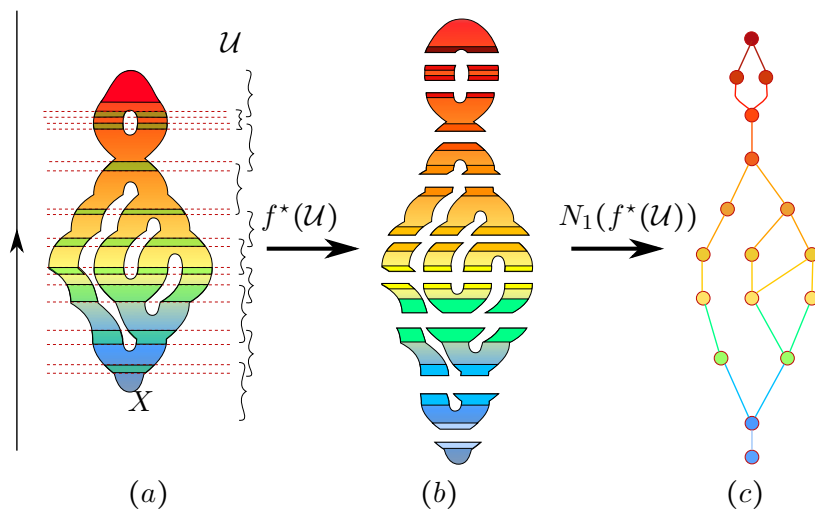


FIGURE 2. (a) Given a scalar function $f : X \rightarrow [a, b]$ and an open cover \mathcal{U} for $[a, b]$ we obtain an open cover $f^*(\mathcal{U})$ for the space X by considering the inverse images of the elements of \mathcal{U} under f . (b) The connected-components of the inverse images are identified 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 $A \cap B \neq \emptyset$.

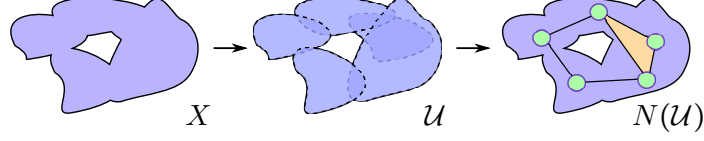


FIGURE 3. Each open set defines a vertex in the nerve simplicial complex. Each intersection between two sets define an edge and intersection between multiple sets define higher order simplices.

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 2 (a). In other words if $\mathcal{U} = \{(a_1, b_1), \dots, (a_n, b_n)\}$ is a finite collection of open sets that covers the interval $[a, b]$ then $f^*(\mathcal{U}) := \{f^{-1}((a_1, b_1)), \dots, f^{-1}((a_n, b_n))\}$ is an open cover for the space X . The open cover $f^*(\mathcal{U})$ can now be used to obtain the 1-nerve graph $N_1(f^*(\mathcal{U}))$. With an appropriate choice of the cover \mathcal{U} , the graph $N_1(f^*(\mathcal{U}))$ is a version of the Reeb graph $R(X, f)$ [12, 31]. This is illustrated in Figure 2.

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

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 simplices are the finite subcollections A_1, \dots, A_k of \mathcal{U} such that $A_1 \cap \dots \cap A_k \neq \emptyset$.

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 3 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 $f : X \rightarrow Z$ 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 $f : X \rightarrow Z$ 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 [15]. All maps in this paper will be assumed to be well-behaved.

Definition 3.3. Let $f : X \rightarrow Z$ 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 $(1, 2), (2, 3), \dots, (6, 1)$. 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 4.

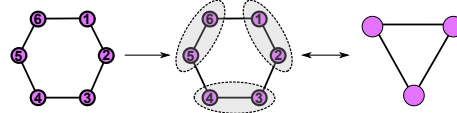


FIGURE 4. An example of a quotient graph.

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. PARALLEL COMPUTING OF MAPPER

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.

Let $f : X \rightarrow [a, b]$ be a continuous function. The construction of parallel 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 5 (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}_2 to have only 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 .
- (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 5 (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 exists in both covers $\mathcal{U}_i, i=1,2$. Let C_1, \dots, C_n be the path-connected components of $f^{-1}(A)$. Since A appears in both of the covers then 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, \dots, v_n that correspond to the components C_1, \dots, 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, \dots, v_k presenting in both G_1 and G_2 . See Figure 5 (c).

The steps of the previous algorithm are summarized in Figure 5.

Remark 4.1. 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 $Y = [a, b]$.

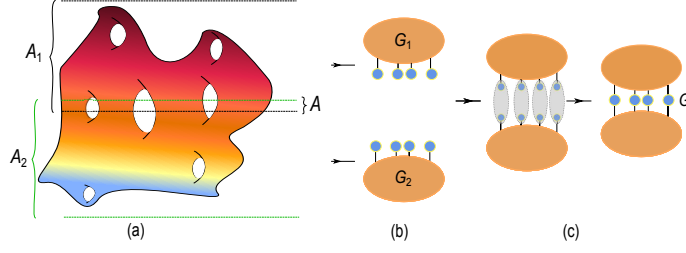


FIGURE 5. The steps of the parallel Mapper on two units. (a) The space X is decomposition based on a decomposition of the codomain (b) Each part is sent to a processing unit and the Mapper graphs are computed on the subspaces (c) The graphs are merged by identifying the corresponding the nodes.

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, \dots, 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 parallel Mapper algorithm that we already explained 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 summarized in the Algorithm 1. Note that we will refer the mapper graph obtained using Algorithm 1 by the *sequential Mapper*.

Algorithm 1: Sequential Mapper [40]

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

- 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 $X_{ij} \cap X_{kl} \neq \emptyset$;
-

5.2. The Main Algorithm. We now give the details of the parallel Mapper algorithm. To guarantee that the output of the parallel Mapper is identical to that of

the sequential Mapper we need to perform some processing on the cover that induces the final parallel Mapper output. In parallel Mapper, we consider an N -chain cover of open intervals A_1, \dots, A_N of the interval $[a, b]$ along with their covers $\mathcal{U}_1, \dots, \mathcal{U}_N$. The details of the cover preprocessing are described in Algorithm 2.

Algorithm 2: Cover Preprocessing

Input: A point cloud X ;

a scalar function $f : X \rightarrow [a, b]$;

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, \dots, 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 $i = 1, \dots, N - 1$;
-

Algorithm 3: Parallel Mapper

Input: A point cloud X ;

a scalar function $f : X \rightarrow [a, b]$;

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: Parallel Mapper Graph.

- 1 **for** ($i \leftarrow 1$ to $i = N$) **do**
 - 2 $P_i \leftarrow (A_i, \mathcal{U}_i)$; //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 $i = 1, \dots, N$. We obtain N graphs G_1, \dots, G_N . If $N = 1$, return the graph G_1 ;
 - 4 Let $C_{j_1}^i, \dots, C_{j_i}^i$ be the clusters obtained from $f^{-1}(A_{i,i+1})$. These clusters are represented by the vertices $v_{j_1}^i, \dots, v_{j_i}^i$ 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, \dots, 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 disjoint union graph $G_1 \sqcup \dots \sqcup G_N$ and then take the quotient of this graph that identifies the corresponding vertices in G_i and G_{i+1} for $1 \leq i \leq N - 1$.
-

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 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.

5.3. Correctness of the Algorithm. In here, we give a detailed proof of the correctness of parallel Mapper that discusses the steps of the algorithm.

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

Proof. We will prove that the parallel 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 parallel 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, and then 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, \dots, 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, \dots, 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 $\cup_i^k \mathcal{U}_i$.

Now we 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 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 $A' \cap A_{k+1}$. This means 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 $k+1$ processors. 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 . \square

6. EXPERIMENTATIONS

In this section, we present practical results obtained using a Python implementation. We ran our experimentation on a Dell OptiPlex 7010 machine with 4-core i7-3770 Intel CPU @ 3.40GHz and with a 24GiB System Memory. The parallel Mapper algorithm was tested on different models and compared their run-time with a publicly available data available at [42]. The size of the point cloud data are

shown in Table 1. The size of datasets given in Table 1 is the number of points in the point cloud data.

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

TABLE 1. The number of points for each dataset used in our tests.

The sequential Mapper algorithm relies on 3 inputs: the data X , the scalar function $f : X \rightarrow [a, b]$ and the choice of cover \mathcal{U} of $[a, b]$. The existing publicly available Mapper implementations, see for instance [30], 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 [5]. Using parallel Mapper on the data given in Table 1, we obtained a remarkable speed up that can be 4 times faster, compared with the sequential Mapper. Figure 6, shows the speedup results of parallel Mapper that are 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. Indeed, at the beginning the speed up increases significantly as we increase the number of processes. However, at some point (when we use more than 10 processes), we increase the number of processes to 30 processes and the speedup does not show significant improvement.

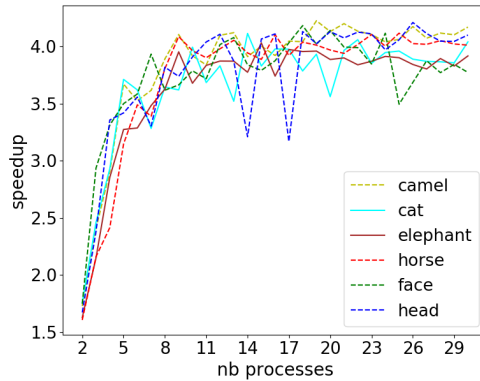


FIGURE 6. Speedups obtained by the parallel Mapper using number of processes that run concurrently.

6.1. Performance of the Algorithm. To verify our experimental results in Figure 6, we use a well-known theoretical formula which is the *Amdahl's law* to calculate the *speedup* ratio upper bound that comes from parallelism and the improvement percentage [1]. The Amdahl'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 processes.

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 processes (e.g. heterogeneity) or some other factors.

In the parallel Mapper, there are two computational pieces which are the clustering piece and the cover construction/merging subgraphs piece. Our algorithm makes the clustering piece completely in parallel while the cover construction/merging subgraphs piece is processed sequentially. Now, we use Amdahl's law to calculate the theoretic speedup ratios to verify the experimental results. Indeed, considering the algorithm, the clustering piece is approximately 75% of execution time, while the cover construction/merging subgraphs piece is about 25%.

In Table 2, we use Amdahl's law to calculate the theoretic speedup ratios using different numbers of processes. The table shows that the speedup increases as a response of the increase in the number of processes. Notice that at some points the performance almost stops improving even if we increase the number of processes. Table 2 shows that the speedup of $part = .75$ (the parallel Mapper) achieves to 3.07 when $N = 10$ and it goes up to 3.99 when $N = 1000$. Therefore, the theoretical calculations clearly matches the experimental results that appears in Figure 6.

Speedup	
N	(Parallel Mapper) part=0.75
10	3.07
100	3.88
1000	3.99
10000	3.99

TABLE 2. Speedup calculations based on Amdahl's law, using different numbers of processes. It shows the speedup of the parallel Mapper with respect to the sequential Mapper

7. CONCLUSION AND FUTURE WORK

In this work we gave a provably correct algorithm to distribute Mapper on a set of processors and run them in parallel. Our algorithm relies on a divide and 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] 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.
- [5] Tadeusz Caliński and Jerzy Harabasz, *A dendrite method for cluster analysis*, Communications in Statistics-theory and Methods **3** (1974), no. 1, 1–27.
- [6] 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.
- [7] Gunnar Carlsson, *Topology and data*, Bulletin of the American Mathematical Society **46** (2009), no. 2, 255–308.
- [8] 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.
- [9] Gunnar Carlsson and Facundo Mémoli, *Persistent clustering and a theorem of j. kleinberg*, arXiv preprint arXiv:0808.2241 (2008).
- [10] Gunnar Carlsson and Afra Zomorodian, *The theory of multidimensional persistence*, Discrete & Computational Geometry **42** (2009), no. 1, 71–93.
- [11] 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.
- [12] Mathieu Carrière and Steve Oudot, *Structure and stability of the 1-dimensional mapper*, arXiv preprint arXiv:1511.05823 (2015).
- [13] Chao Chen and Michael Kerber, *Persistent homology computation with a twist*, Proceedings 27th european workshop on computational geometry, 2011.
- [14] 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.
- [15] 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.
- [16] Tamal K Dey, Facundo Memoli, and Yusu Wang, *Topological analysis of nerves, reeb spaces, mappers, and multiscale mappers*, arXiv preprint arXiv:1703.07387 (2017).
- [17] 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.
- [18] 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.
- [19] Robert Ghrist, *Barcodes: the persistent topology of data*, Bulletin of the American Mathematical Society **45** (2008), no. 1, 61–75.
- [20] Marian Gidea, *Topology data analysis of critical transitions in financial networks* (2017).
- [21] 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.
- [22] 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.
- [23] 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.
- [24] Yasuaki Hiraoka, Takenobu Nakamura, Akihiko Hirata, Emerson G Escobar, 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.
- [25] Ryan H Lewis and Afra Zomorodian, *Multicore homology via mayer vietoris*, arXiv preprint arXiv:1407.2275 (2014).
- [26] David Lipsky, Primoz Skraba, and Mikael Vejdemo-Johansson, *A spectral sequence for parallelized persistence*, arXiv preprint arXiv:1112.1245 (2011).
- [27] 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.
- [28] Dmitriy Morozov and Gunther Weber, *Distributed merge trees*, Acm sigplan notices, 2013, pp. 93–102.
- [29] Dmitriy Morozov and Gunther H Weber, *Distributed contour trees* (2012).
- [30] 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).
- [31] Elizabeth Munch and Bei Wang, *Convergence between categorical representations of reeb space and mapper*, arXiv preprint arXiv:1512.04108 (2015).
- [32] James R Munkres, *Elements of algebraic topology*, Vol. 2, Addison-Wesley Menlo Park, 1984.
- [33] 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.
- [34] 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.
- [35] Valerio Pascucci and Kree Cole-McLaughlin, *Parallel computation of the topology of level sets*, Algorithmica **38** (2004), no. 1, 249–268.
- [36] Vanessa Robins, *Towards computing homology from finite approximations*, Topology proceedings, 1999, pp. 503–532.
- [37] Alejandro Robles, Mustafa Hajij, and Paul Rosen, *The shape of an image: A study of mapper on images*, to appear VISAPP 2018 (2018).
- [38] 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.
- [39] 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.
- [40] Gurjeet Singh, Facundo Mévoli, and Gunnar E Carlsson, *Topological methods for the analysis of high dimensional data sets and 3d object recognition.*, Spbg, 2007, pp. 91–100.
- [41] Václav Snášel, Jana Nowaková, Fatos Xhafa, and Leonard Barolli, *Geometrical and topological approaches to big data*, Future Generation Computer Systems **67** (2017), 286–296.
- [42] Robert W Sumner and Jovan Popović, *Deformation transfer for triangle meshes*, ACM Transactions on Graphics (TOG) **23** (2004), no. 3, 399–405.

KLA CORPORATION, USA

E-mail address: mustafahajij@gmail.com

DEPARTMENT OF COMPUTER SCIENCE, JAZAN UNIVERSITY, JAZAN CITY, SAUDI ARABIA

E-mail address: babumusmar@jazanu.edu.sa

UNIVERSITY OF SOUTH FLORIDA, TAMPA, FL 33620, U.S.A.

E-mail address: prosen@usf.edu