

Article

Merging Discrete Morse Vector Fields: A Case of Stubborn Geometric Parallelization

Douglas Lenseth and Boris Goldfarb * 

Department of Mathematics and Statistics, State University of New York, Albany, NY 12222, USA; dlenseth@albany.edu

* Correspondence: bgoldfarb@albany.edu

Abstract: We address the basic question in discrete Morse theory of combining discrete gradient fields that are partially defined on subsets of the given complex. This is a well-posed question when the discrete gradient field V is generated using a fixed algorithm which has a local nature. One example is `ProcessLowerStars`, a widely used algorithm for computing persistent homology associated to a grey-scale image in 2D or 3D. While the algorithm for V may be inherently local, being computed within stars of vertices and so embarrassingly parallelizable, in practical use, it is natural to want to distribute the computation over patches P_i , apply the chosen algorithm to compute the fields V_i associated to each patch, and then assemble the ambient field V from these. Simply merging the fields from the patches, even when that makes sense, gives a wrong answer. We develop both very general merging procedures and leaner versions designed for specific, easy-to-arrange covering patterns.

Keywords: discrete Morse theory; gradient vector field; distributed computation

MSC: primary 57Q70; 68Q85; secondary 68U05; 68W10; 68W15



Citation: Lenseth, D.; Goldfarb, B. Merging Discrete Morse Vector Fields: A Case of Stubborn Geometric Parallelization. *Algorithms* **2021**, *14*, 360. <https://doi.org/10.3390/a14120360>

Academic Editors: Michele Squizzato and Gopal Pandurangan

Received: 28 October 2021
Accepted: 9 December 2021
Published: 11 December 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Discrete Morse theory [1] is a fairly new and powerful tool, created by Robin Forman in the 1990s, that has many applications in many fields. One important application of discrete Morse theory is streamlining homology calculations. It has been a very useful tool in topological data analysis, for example, in computing persistent homology.

More specifically, there are algorithms that take both 2D and 3D digital grayscale images and create discrete Morse functions from their grayscale function values defined on pixels or voxels. These images give rise to cubical complexes, where the voxels or pixels are the 0-cells of the complex. One algorithm, `ProcessLowerStars` from [2], takes all of the cells in the cubical complex and puts them either singly in the set of critical cells, C , or as a face-coface (of codimension 1) pair in the discrete vector field, V . The set of critical cells and the discrete vector field are the defining features of a discrete Morse function, and in turn the discrete Morse complex. In discrete Morse theory, the discrete vector field determines the maps in the simplified chain complex of the Morse complex that is created from the critical cells of the discrete Morse function.

Given a general algorithm α applied to cells in a regular CW complex K , which takes data from k -neighborhoods of cells in K and uses those data to either pair cells together in a list, which is to become the discrete vector field V , or places the cells singly in the list of critical cells C . The basic question we address is, how do we merge the vector fields in the elementary situation where $K = U \cup W$, and U, W are two subcomplexes to which α is applied individually to obtain the discrete vector fields $V(U)$ and $V(W)$? The goal is to reconstruct the correct vector field $V(K)$ that we would obtain on all of K . It turns out this is not straightforward in that $V(K) \neq V(U) \cup V(W)$, and so α , is not embarrassingly parallelizable. Being able to perform such a computation will involve looking at where

mistakes occur in naively merging the vector fields together as in the right-hand side of the formula (or, in general, related vector fields) and “crossing out” those mistakes.

What do we mean by an algorithm being embarrassingly parallelizable? One can take an example with operators in calculus to illustrate the difference between embarrassingly and stubbornly parallelizable algorithms. Differentiation is embarrassingly parallelizable with respect to addition and subtraction, i.e., $(f(x) \pm g(x))' = f'(x) \pm g'(x)$. It is stubbornly parallelizable with respect to multiplication and division, i.e., $(f(x)g(x))' = f'(x)g(x) + f(x)g'(x)$, which is a non-obvious replacement of the straightforward but incorrect formula $(f(x)g(x))' = f'(x)g'(x)$.

We will present the main parallelization theorem, Theorem 2, that will allow us to perform this merging for a general uniformly k -local algorithm α . Here, k -local refers to the nature of the algorithm in that it makes the classification decision about a cell or a pair by processing only the information from k -neighborhoods of cells in a regular CW complex. This turns out to be the common type of algorithm in practice. The authors do not know of any α that builds discrete Morse fields or functions and does not have this property. We then make this merging process more efficient in Corollary 1 by noting some geometric properties of the vector fields we get, which leads to a merging that involves applying α on fewer subcomplexes, i.e., a more efficient distributed computation. We then look at the situation where the subcomplexes U and W are in antithetic position, satisfying $(U \cap W)[k] = U[k] \cap W[k]$, where $k \geq 1$ and $U[k]$ is the k -neighborhood of the subcomplex U . This special relative positioning of the patches allows to further improve the efficiency by allowing to use smaller auxiliary patches. This is spelled out in Corollary 2.

In the context of α being `ProcessLowerStars`, it turns out that the rigid structure of the cubical complexes that are formed from the images gives a very efficient formula, related to but not exactly matching the formula for the more general case. We present a formula, in the 2D image case, that allows us to take lists of discrete vector fields directly from two patches and merge them together to obtain the discrete vector field of the union of the two patches, under the condition that the two patches must have an overlap that contains at least one 2-cell. This is done in Section 4 and is an illustration of how the general formulas can be further tweaked to improve the efficiency of the procedure.

Why do we want to parallelize this `ProcessLowerStars` algorithm? For one, if one has a particularly large image with a lot of pixels or voxels, being able to break the image up and apply the algorithm on each piece will reduce the time cost compared to applying the algorithm on the whole image. It also may be that an image may come to you in pieces via online machine learning. The methods of this paper allow to process each piece as it arrives and either gradually build the discrete field or put the whole field together as a final quick step.

In an application to the explicit assembly strategy for the `ProcessLowerStars` algorithm on a 2D image, we assume that our computer has a hypothetical limited constraint (i.e., restricted to processing, at most, a certain number of pixels). In this situation, we have to break up the image so that the number of pixels in each patch is less than the constraint of our computer. These patches are rectangular and allow us to index our patches both in the horizontal and vertical directions. We apply the `ProcessLowerStars` algorithm to each patch (and all intersections, in both the horizontal and vertical direction, of adjacent patches). We create an algorithm that merges the discrete vector field in each of the patches to obtain the discrete vector field of the whole image, using the formula on two patches. In order to do this, our algorithm merges starting in the top left corner of the image and moves along each horizontal strip and their intersections. Once the discrete vector field of every horizontal strip and their intersections is found, the algorithm then moves vertically from top down to merge the discrete vector fields of the horizontal strips two at a time, again making use of the formula that will be provided.

Finally we look at another specific special geometric situation that involves simplicial trees in Section 5. This is done with two purposes. It is known that finite products of trees contain coarse images of the most common structures in geometry of finite asymptotic

dimension. A sufficient partition scheme in a product of trees, therefore, provides a generic template to approach most common regular cellular complexes. On the other hand, we find that a tree is an example that distinguishes the advantages of each of the two main parallelization results of the paper.

We finish with a collection of remarks on applications of the results, the literature, and directions for future work in Sections 6 and 7.

2. Discrete Gradient Vector Fields

All complexes in this paper are regular CW complexes in the sense that characteristic maps of all cells are homeomorphisms onto the closures of cells in the complex. We refer, for the general background in discrete Morse theory on regular CW complexes, to [3].

Notation 1. Suppose K is a regular complex, and suppose f is a discrete Morse function on K . Then there are well-defined sets of critical cells $C(K, f)$ and the discrete vector field that can be thought of as a set of pairs of regular cells $V(K, f)$ or $V(K)$. When we use the notation $(\sigma \leq \tau)$ for a pair in $V(K)$, we mean that σ is a codimension 1 face of τ .

Definition 1. We introduce the following relation among cells in K . Two cells σ_1 and σ_2 are adjacent if the closures of the two cells in K have a point in common. The star of a cell σ is the smallest cellular subcomplex of K that contains all cells adjacent to σ . Given a collection of cells C of K , let the star of C be the smallest subcomplex that contains the stars of all of the cells in C . If K' is a subcomplex, we also call its star the 1-neighborhood of K' denoted $K'[1]$. The star of a star is the 2-neighborhood, etc. We use the notation $K'[n]$ for the n -neighborhood of K' . The k -border of a subcomplex L of K consists of those cells in L whose k -neighborhoods are not contained entirely in L .

We assume that the lists $C(K, f)$ and $V(K, f)$ are results of applying an algorithm $\alpha: K \mapsto V$ of the following nature.

Definition 2. Given a face-coface pair of cells $(\sigma \leq \tau)$, the algorithm decides whether the pair $(\sigma \leq \tau)$ is placed in the list $V(K)$ based on data specific to α coming from some k -neighborhood $\sigma[k]$, for some uniformly fixed number $k > 0$. The result of the algorithm is a discrete vector field $V_\alpha(K)$ which can be therefore realized as $V(K, f)$ for some discrete Morse function f . We say that α with this property is uniformly local or, more precisely, uniformly k -local.

In fact, in all applications that generate $V(K)$ that the authors see in the literature, it is always done using a uniformly 1-local algorithm. This is not surprising, of course, as traditional smooth vector fields are generated through differentiation procedures, which are inherently local.

Example 1. An example of such an algorithm is *ProcessLowerStars*, which produces $V(K)$, starting with an arbitrary positive bounded function (also known as a grayscale function) on cells (pixels or voxels) of a 2D or 3D pixel grid K . Since the algorithm qualifies $(\sigma \leq \tau)$ based on the values of a function on vertices in $\sigma[1]$, this algorithm is uniformly k -local for $k = 1$.

To be more specific, to make these images into a regular CW complex, we translate them to cubical complexes by making the pixels or voxels correspond to 0-cells of the cubical complex, K . This gives us a positive bounded function, called g , on a cubical complex: more specifically, a cubulated plane or space. To apply the algorithm, it is ideal to have the values of g on all of the 0-cells be unique. If they are not, one can make small changes (i.e., a linear ramp) to g to ensure they are unique. Once there are unique values of g , the algorithm inspects all of the cells in the lower star of a 0-cell, x . The lower star of x , $L(x)$, contains all cells $\sigma \in K$ in the cubical complex such that $g(x) = \max_{y \in \sigma} g(y)$. To give an ordering to higher dimensional cells, a new function, G , is introduced, defined as follows: if σ contains the vertices $\{x, y_1, \dots, y_n\}$, then

$G(\sigma) = \{g(x), g(y_{i_1}), \dots, g(y_{i_n})\}$, where $g(x) > g(y_{i_1}) > \dots > g(y_{i_n})$. This allows us to impose the lexicographical ordering on G when performing this algorithm.

The algorithm, itself, then takes all of the cells of $L(x)$ and either possibly pair a cell, τ , with a codimensional 1 face, σ , and place them ($\sigma \leq \tau$) in V , or take a cell, σ , and place it singly in C in the following way:

- If $L(x) = \{x\}$, then $x \in C$. Otherwise, take the minimal (with respect to G) 1-cell, τ , and $(x \leq \tau) \in V$. All other 1-cells are added to a queue called PQzero since they have no remaining unpaired faces in $L(x)$. All cofaces of codimension 1 of τ in $L(x)$ that have 1 unpaired face in $L(x)$ are added to a different queue called PQone.
- The algorithm then takes the minimal cell (with respect to G) that is in PQone and either moves it to PQzero, if it has no unpaired faces remaining in $L(x)$. If it still has an unpaired face, it gets paired with that face and is put into V . Then we look at all cofaces of codimension 1 of both cells that were just paired and put into V . If any of these cofaces have exactly one unpaired face, they are added to the queue PQone.
- Then, if PQzero is not empty, it takes the minimal cell (with respect to G), called δ , and places it singly in C . Then all cofaces of codimension 1 of δ are inspected. If any these cofaces have exactly one unpaired face, it is placed in PQone.

This keeps going until both PQzero and PQone are empty. Note that in particular, $L(x) = x[1]$, using the notation defined above. *ProcessLowerStars*, in fact, only looks at $\sigma[1]$, where σ is a 0-cell.

Example 2. We present a simple example in Figure 1 with $V_\alpha(K) \neq V_\alpha(U) \cup V_\alpha(W)$, where α is the *ProcessLowerStars* algorithm from [2] and U and W are subcomplexes of K such that $K = U \cup W$. Let K be the following 1-dimensional cubical complex with the function values, f , on the 0-cells.



Figure 1. A 1D example of a failed merging of discrete vector fields.

First, we apply α to all of K . This gives $V_\alpha(K) = \{\{3, 31\}, \{4, 42\}\}$. Next, we break up K into the following subcomplexes: $U = \{1, 3, 4, 31, 43\}$ and $W = \{2, 3, 4, 43, 42\}$. When we apply α to U , we obtain $V_\alpha(U) = \{\{3, 31\}, \{4, 43\}\}$. When we apply α to W , we get $V_\alpha(W) = \{\{4, 42\}\}$. Clearly, $V_\alpha(K) \neq V_\alpha(U) \cup V_\alpha(W)$, so simply merging the lists from the subcomplexes does not give us the correct discrete vector field for all of K . In this particular instance, we obtain an extra pairing, the $\{4, 43\}$ from U .

The failure of naive merging is made clearer using the following definition and explanation. Let $\sigma_C[k]$ be the k -star of σ viewed as a cell within the subcomplex C of K . This is intrinsic to the C construction and so may differ from $\sigma[k]$, which is the same as $\sigma_K[k]$, for some combinations of σ and k . Since the algorithm α takes as input the data from $\sigma_C[k]$, there are cases when the discrepancy leads to different constructions of the vector fields as Example 2 shows. This also points to some cases when this discrepancy does not happen.

Proposition 1. Suppose K is a disjoint union of subcomplexes C_i , for i from some index set I . In other words, a connected subset of K is contained entirely within one and only one C_i . Then $V_\alpha(K) = \bigcup_{i \in I} V_\alpha(C_i)$.

Proof. For every pair ($\sigma \leq \tau$), its inclusion in $V_\alpha(C_i)$ is decided by α from the data within $\sigma_{C_i}[k]$. Since in our case $\sigma_{C_i}[k] = \sigma_K[k]$, that decision is precisely the same within C_i as within K . □

This argument makes it clear that failure of the simple merge is caused by possible discrepancies between $\sigma_C[k]$ and $\sigma_K[k]$ for cells σ in the k -border of C (see Definition 1).

3. The Parallelization Theorems

Suppose a complex K is assigned a gradient vector field V using an algorithm α . Suppose further that U and W are two subcomplexes of K that form a covering of K . In particular, there are vector fields $V_\alpha(K)$, $V_\alpha(U[i])$, and $V_\alpha(W[i])$ for various values of i .

Theorem 1. *If α is a uniformly k -local algorithm for some $k \geq 1$, then*

$$V_\alpha(K) = (V_\alpha(U[k]) \cup V_\alpha(W[k])) \setminus (V_\alpha(U[k]) \cup V_\alpha(W[k]) \setminus V_\alpha(U[2k+1] \cap W[2k+1])).$$

In slightly different terms,

$$V_\alpha(K) = (V_\alpha(U[k]) \cup V_\alpha(W[k])) \setminus (V_\alpha(U[k]) \setminus V_\alpha(U[2k+1])) \setminus (V_\alpha(W[k]) \setminus V_\alpha(W[2k+1])).$$

The informal idea is this. It is elementary to check that $V_\alpha(K)$ is contained in $V_\alpha(U[k]) \cup V_\alpha(W[k])$ using that the algorithm α is uniformly k -local. Let us refer to the difference as “mistakes” in building $V_\alpha(K)$ from the union. We see that for the same reason, $V_\alpha(U[k]) \setminus V_\alpha(U[2k+1])$ can be viewed as all the mistakes made near the border of $U[k]$. This set is interpreted as pairs of cells in $V_\alpha(U[k])$ that are not in $V_\alpha(K)$. Similarly, $V_\alpha(W[k]) \setminus V_\alpha(W[2k+1])$ is the set of mistakes made near the border of $W[k]$. Since the cells near the borders of the subcomplexes are far enough away in the adjacency relation chains, they do not contain common adjacent cells. The local nature of the algorithm, again, guarantees that no other mistakes are made away from the borders of $U[k]$ and $W[k]$.

We now give a formal proof.

Proof. If the pair $(\sigma \leq \tau)$ is contained in $V_\alpha(K)$, it is placed in this list according to the algorithm α . According to the assumption, this placement is the outcome of inspection of data relevant to α in $\sigma[k]$. If σ and τ are cells of U , then $\sigma[k]$ is contained in $U[k]$. Since α applied to $U[k]$ would use the exact same data as $\sigma[k] = \sigma_{U[k]}[k]$, this proves that $(\sigma \leq \tau)$ is in $V_\alpha(U[k])$. Repeating the argument for W in place of U , this gives the inclusion $V_\alpha(K) \subset V_\alpha(U[k]) \cup V_\alpha(W[k])$.

We now examine $(\sigma \leq \tau)$ in the difference

$$M = V_\alpha(U[k]) \cup V_\alpha(W[k]) \setminus V_\alpha(K).$$

Observe that whenever the cell σ is outside of $W[k]$, it is contained in U . This makes $\sigma_K[k] = \sigma_{U[k]}[k]$, so the outcome of the decision α makes on the inclusion of $(\sigma \leq \tau)$ in, respectively, $V_\alpha(K)$ and $V_\alpha(U[k])$ is the same. This shows that such $(\sigma \leq \tau)$ is disqualified from M , as it is excluded as part of $V_\alpha(K)$. Our conclusion is that whenever $(\sigma \leq \tau)$ is in M , σ must be in $W[k]$. A symmetric argument proves that it also must also be in $U[k]$.

The last paragraph proves that our pair of cells σ, τ is contained in $U[k+1] \cap W[k+1]$. Therefore, α uses the same data for placing $(\sigma \leq \tau)$ in either $V_\alpha(U[2k+1] \cap W[2k+1])$ or $V_\alpha(K)$. This allows us to rewrite

$$M = V_\alpha(U[k]) \cup V_\alpha(W[k]) \setminus V_\alpha(U[2k+1] \cap W[2k+1]).$$

This gives the first formula.

Let $M_U = V_\alpha(U[k]) \setminus V_\alpha(K)$ and $M_W = V_\alpha(W[k]) \setminus V_\alpha(K)$, so that $M = M_U \cup M_W$. The same argument as above shows that

$$M_U = V_\alpha(U[k]) \setminus V_\alpha(U[2k+1]),$$

and similarly

$$M_W = V_\alpha(W[k]) \setminus V_\alpha(W[2k+1]).$$

Now

$$V_\alpha(K) = V_\alpha(U[k]) \cup V_\alpha(W[k]) \setminus (M_U \cup M_W)$$

gives the second formula. \square

Remark 1. Our interest in this theorem is that it allows to express the global list $V_\alpha(K)$ in terms of generally smaller partial lists $V_\alpha(U[k])$, $V_\alpha(W[k])$, $V_\alpha(U[2k+1])$, and $V_\alpha(W[2k+1])$.

The computation of these smaller lists can be further simplified by the observation that the difference between $V_\alpha(U[k])$ and $V_\alpha(U[2k+1])$ is contained in $V_\alpha(U[2k+1] \cap W[2k+1])$.

Corollary 1. If α is a uniformly k -local algorithm for some $k \geq 1$, then

$$V_\alpha(K) = (V_\alpha(U[k]) \cup V_\alpha(W[k])) \setminus (V_\alpha(U[k] \cap W[k])) \setminus (V_\alpha(U[2k+1] \cap W[2k+1])).$$

Proof. The formula from Theorem 2 is rewritten in terms of the intersection $U[2k+1] \cap W[2k+1]$. Notice that the restriction to $V_\alpha(U[k] \cap W[k])$ is justified because the k -borders of U and W and thus also M are contained in $U[k] \cap W[k]$. \square

We next implement the formula as Algorithm 1 in this simple case of a two-set covering.

Algorithm 1 Two-set parallelization from Corollary 1.

Inputs:

- $k \geq 1$
- α uniformly k -local algorithm
- K regular CW complex
- D data attached to cells in K , specific for use in α
- $U(1), U(2)$ a covering of K by subcomplexes given as subsets of K

Outputs:

- $V = V_\alpha(K, D)$ discrete vector field
- $C = C_\alpha(K, D)$ critical cells

for $i = 1, 2$ **do**

- **build** $U(i)[k]$
- **build** $U(i)[2k+1]$
- **build** $D|U(i)[k]$
- **evaluate** $\alpha: U(i)[k], D|U(i)[k] \rightarrow A(i, k)$

preprocessing step to obtain discrete vector fields in enlargements of patches

end for

build the union $V(1, 2) = A(1, k) \cup A(2, k)$

build the intersection $U(1, 2, k) = U(1)[k] \cap U(2)[k]$

build the intersection $U(1, 2, 2k+1) = U(1)[2k+1] \cap U(2)[2k+1]$

build $D|U(1, 2, k)$

build $D|U(1, 2, 2k+1)$

evaluate $\alpha: U(1, 2, k), D|U(1, 2, k) \rightarrow A(1, 2, k)$

evaluate $\alpha: U(1, 2, 2k+1), D|U(1, 2, 2k+1) \rightarrow A(1, 2, 2k+1)$

build the difference $D = A(1, 2, k) \setminus A(1, 2, 2k+1)$

save the difference $V = V(1, 2) \setminus D$

build N as all cells employed in V

save the difference $C = K \setminus N$

There is a special geometric situation with a more direct identification of intersections of enlargements.

Definition 3. Two subcomplexes U and W of K are said to be in antithetic position if $(U \cap W)[n] = U[n] \cap W[n]$ for all $n \geq 0$.

We can now state the following consequence of Corollary 1.

Corollary 2. *Suppose α is an algorithm that is uniformly k -local for some $k \geq 1$ and suppose U and W are an antithetic pair, then*

$$V_\alpha(K) = (V_\alpha U[k] \cup V_\alpha W[k]) \setminus (V_\alpha(U \cap W)[k] \setminus V_\alpha(U \cap W)[2k + 1]).$$

Proof. The subcomplexes in Corollary 1 now have a description based entirely on the enlargements of covering complexes and their intersections. \square

Remark 2. *In general, building $S[n]$ in K for a subcomplex S and a number $n \geq 1$ requires building $x[n]$ for all x in S and taking the union. So the time complexity of this operation is proportional to the size of S . Therefore, depending on the geometry of K and the choices of U and W , the sizes of the lists the computer needs to work through while building the enlargements can be greatly reduced by working with the intersection $U \cap W$ and one single enlarging procedure. This is particularly true when the size of the intersection $U \cap W$ is significantly smaller than the sizes of U and W . In the case of simplicial trees, which we address later in Section 5, the size of $U \cap W$ can be arranged to be an order of magnitude smaller than either U or W in the Euclidean case, and that order can be made arbitrarily smaller in the case of a tree, depending on the valence of the tree.*

In the rest of the section, we illustrate in the specific context of the algorithm `ProcessLowerStars` [2] that the general theorems can be improved on in situations with specific α and specific geometry.

Suppose K is a cubulated 2D rectangle $I \times J$ and $I = [a, d]$ for $d - a \geq 4$. Suppose b and c are integers nested in $a < b < c < d$. Then we have a decomposition of K as $U \cup W$ with $U = [a, c] \times J$ and $W = [b, d] \times J$.

Theorem 2. *Applying `ProcessLowerStars` as α to the two sets U and W as above allows some savings in the size of processed lists by using a more efficient formula*

$$V_\alpha(K) = (V_\alpha(U) \cup V_\alpha(W)) \setminus V_\alpha(U \cap W) \setminus V_\alpha(U \cap W)[1].$$

Notice that this theorem does not follow from any of the general theorems before because no enlargement of the patches is required for the containment $V_\alpha(K) \subset V_\alpha(U) \cup V_\alpha(W)$.

Proof. The proof is based on a case-by-case analysis of `ProcessLowerStars` applied to stars of vertices near $[b, c] \times J$. As before, we are merging two lists $V_\alpha(U)$ and $V_\alpha(W)$. We know that in each complex, there are pairs that involve the vertices in $b \times J$ that may be included in non-critical pairs in error. In a complementary fashion, the same can be said about the critical cells in $C_\alpha(U)$ and $C_\alpha(W)$. We know the errors can happen in this specific case of α because of a 2D counterexample similar to Example 2. In much the same way, the errors happen at this boundary because that is where incomplete information about their stars in K is used within U and W . Let us call the pairs included in error “mistakes”. We try to identify the mistakes in the merged list $V_\alpha(U) \cup V_\alpha(W)$. Are they guaranteed to be inside $V_\alpha(U \cap W)$? This turns out to be true in this case but is not expected for general uniformly local α or more general geometric decompositions. Next, suppose a pair in $V_\alpha(U) \cup V_\alpha(W)$ is a mistake, then we know that it is not in $V_\alpha((U \cap W)[1])$ because now both cells are in a star entirely contained in $V_\alpha((U \cap W)[1])$. This guarantees that the difference $V_\alpha(U \cap W) \setminus V_\alpha((U \cap W)[1])$ consists precisely of the mistakes in $V_\alpha(U) \cup V_\alpha(W)$ and completes the proof. \square

4. Distributed `ProcessLowerStars` Algorithm on a 2D Digital Image

We apply Theorem 2 to parallelize an algorithm of type α on grayscale 2D images, in particular, the algorithm `ProcessLowerStars` from [2].

These digital images are modeled by cubical complexes with a grid-like structure, which makes it a little easier to work with than other cubical complexes, as we can index

our patches with horizontal and vertical components, which will shortly be outlined more clearly. Our formula for merging the vector field together only makes use of two patches. In some cases, in particular, with a picture with many pixels and a computer program having a constraint on the number of pixels it can process with the algorithm, it is very likely that one would end up with more than two patches. The grid-like structure that these cubical complexes possess will allow us to apply our formula to two pieces at a time. First, the image is split up into patches, which partition the cubical complex. How many patches we end up with is determined by the number of pixels in the image and the constraint we have on the number of pixels our computer program can process. After taking the star of each of the patches in both the vertical and horizontal directions separately, we are left with new patches where there is overlap, both vertically and horizontally, between adjacent patches. This allows us to apply our formula.

After the partitioning into patches and taking the stars in both directions of all the patches, we apply the `ProcessLowerStars` algorithm to each patch and all of the overlaps of the adjacent patches, including the overlaps of the overlaps. The algorithm gives the discrete vector field in each of the patches and all of the overlaps of adjacent patches. We then start applying our formula, starting in the upper left hand side of our complex, with the first two starred patches. This gives the discrete vector field of the union of these two starred patches. We then proceed by applying our formula to this new bigger patch with the adjacent starred patch to its right. We continue until we reach the end of this horizontal strip. We then move down vertically to the next row of starred patches and go through the same procedure. We do this for every horizontal strip of starred patches and obtain the discrete vector field for each horizontal strip.

We then move to the overlaps of the horizontal strips and use the same procedure as the horizontal strips themselves, starting from the left hand side, applying our formula as we move to the right end of the strip. We eventually end up with the discrete vector field of the overlaps of the horizontal strips. We need the discrete vector field of the overlaps of the horizontal strips in the next part of our algorithm when we start applying our formula vertically.

We now have the discrete vector field of all of the horizontal strips and their overlaps, so we can now proceed applying our formula in the vertical direction, starting with the top most horizontal strip and applying our formula with the strip directly below it, making use of the discrete vector field of the overlaps that we computed in the previous steps. We continue downward until we have the discrete vector field of one big patch and the horizontal strip below it, along with the discrete vector field of their overlap. We apply our formula to these last two patches and end up with the discrete vector field of the whole image.

We make this precise with a pseudo-code for the distributed `ProcessLowerStars` algorithm. We assume that the constraint of our computer program in applying the `ProcessLowerStars` algorithm on a digital image is N pixels. Assume that a digital image has D pixels where $D > N$. We break the image up into n disjoint patches such that $n > D/N$, where patch $P_{i,j}^\#$ is a subcomplex and corresponds to the patch in the i th column and j th row for all $i = 1, \dots, m$ and $j = 1, \dots, \ell$ such that $m \cdot \ell = n$.

Let $P_{i,j}^\#[a,b]$ be the subcomplex that is the a -neighborhood of $P_{i,j}^\#$ in the horizontal direction together with the b -neighborhood of $P_{i,j}^\#$ in the vertical direction. We take our disjoint decomposition $P_{i,j}^\#$ and enlarge each patch by 1-neighborhood in each direction and call it $P_{i,j}$. This gives us $P_{i,j} = P_{i,j}^\#[1,1]$, where each patch now overlaps and the overlaps contain at least one 2-cell. Let the following hold:

1. $P_{i,j}^* = \bigcup_{k=1}^i P_{k,j}$ (the union of patches moving across a horizontal strip);

2. $(P_{i,j} \cap P_{i,j+1})^* = \bigcup_{k=1}^i (P_{k,j} \cap P_{k,j+1})$ (the union of patches moving horizontally across the intersection of patches that are vertically adjacent to each other, i.e., moving horizontally along the intersection of strips);
3. $(P_{i,j}[0,1] \cap P_{i,j+1}[0,1])^* = \bigcup_{k=1}^i (P_{k,j}[0,1] \cap P_{k,j+1}[0,1])$ (the union of patches moving horizontally across the intersection of 1-neighborhoods in the vertical direction of the patches that are vertically adjacent to each other).

Figure 2 is used to illustrate this notation. The figure is intended to show a part of the decomposition of the 2D plane centered on the middle patch denoted $P_{i,j}$. The orange, green, and blue overlapping rows are the unions of such patches across the values of the horizontal parameter i . One useful observation is that there are only double overlaps of the rows so that there is always at least 2δ clearance between non-adjacent rows, such as the orange and the blue, for a fixed value of δ . This guarantees that δ -enlargements of this covering of the plane by rows have the same nerve as the rows themselves. The nerve is isomorphic to a triangulation of the real line. The figure also shows the set $P_{i,j}^*$, which is shaded red, and the union of intersections of patches $(P_{i,j} \cap P_{i,j+1})^*$. It is clear that the finer covering of the plane by individual patches has the same property as above: its δ -enlargement has the same nerve as the covering itself.

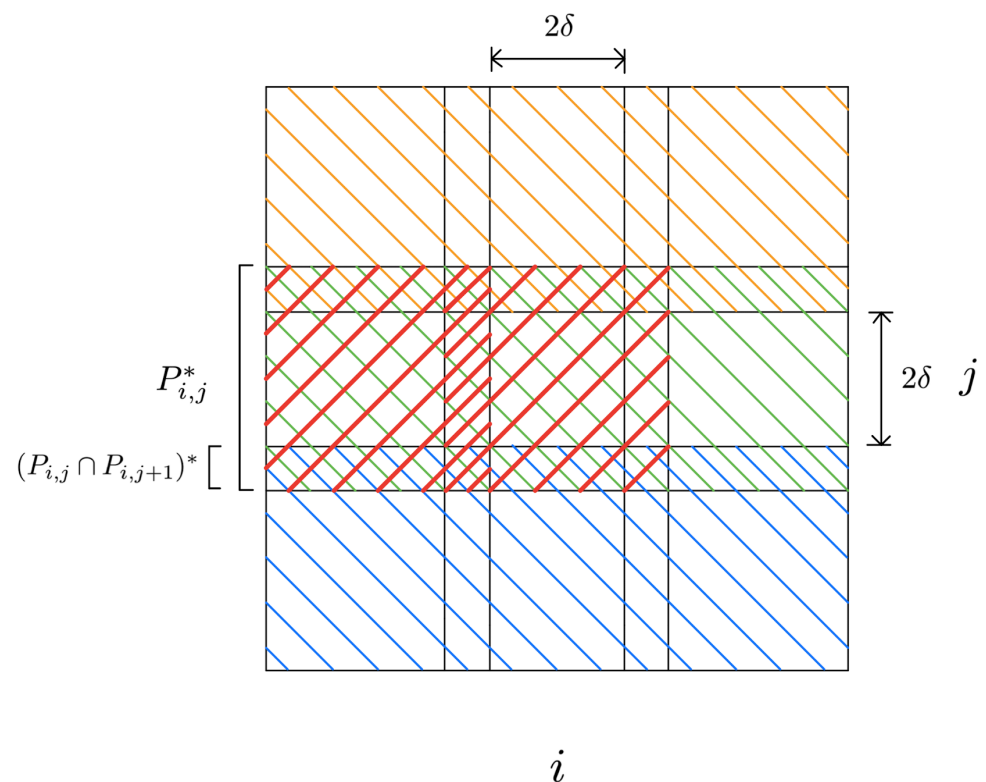


Figure 2. An illustration of the planar grid-like decomposition.

5. Generalization to a Hierarchical Tree-like Decomposition

Our main motivation for this work is the parallelization of the specific algorithm ProcessLowerStars on 2D images that we performed in Sections 3 and 4. This algorithm is proven to work in cubical cellular complexes based on standard cubical grids in 2D and 3D [2]. However, our theorems in Section 3 have only general local constraints on the type of the algorithm and on the geometry of the regular cellular complex. In this section, we want to leverage some geometric material from [4,5] to construct the required antithetic

coverings for grids in nD for all n and, more generally, any subcomplex of a product of finite locally finite trees.

It is known from the result of Dranishnikov [6,7] that all metric spaces which satisfy a very weak and natural geometric condition called finite asymptotic dimension (FAD) can be coarsely embedded in a finite product of locally finite trees with uniform distortion. This allows us to give a useful antithetic covering of any cellular complex, which can be given as a FAD metric with a universal bound on the size of all cells.

As we observed before in Remark 2, the importance of the case of a tree or a product of trees is also an illustration of how crucial the improvements can be in passage from using the most general Theorem 2 to using Corollary 2.

We saw in the previous section a worked-out example of the use of antithetic decompositions in the case of a cubical grid in 2D. Just as in that example, it is most natural to decompose a multi-parameter geometry according to projections to subsets in one chosen parameter. There should result an inductively defined decomposition of the entire cubical complex. The general kind of parameter for our purposes is tree based, with partial order, generalizing from the totally ordered real line.

A simplicial tree is a simplicial complex that is connected and has no cycles. Recall also that a nerve of a collection of subsets of a given set is the simplicial complex with vertices which are the subsets and simplices corresponding to non-empty intersections of families of subsets.

Definition 4. Suppose \mathcal{U} is the covering of a metric space. We denote by $\mathcal{U}[k]$ the covering by k -enlargements of members of \mathcal{U} for a number $k \geq 0$. We say that a covering \mathcal{U} is a tree-like decomposition with margin k if $\mathcal{U}[k]$ is a simplicial tree. Suppose each of the covering sets with the subspace metric also has a tree-like decomposition with margin k . Then we say that the resulting covering by smaller sets is a hierarchical tree-like decomposition with margin k and depth 2. Inductively, for a natural number D , one defines a hierarchical tree-like decomposition with margin k and depth D . We refer to the sets that appear in such hierarchical decomposition and are not unions of other sets as primary sets.

The most useful hierarchical tree-like decomposition of depth D can be obtained for any subset of the product of D simplicial trees. The simplest case of this type of decomposition is when the trees are obtained as triangulations of the real line, which generalizes 2D decompositions as in Algorithm 2 to higher dimensions.

Definition 5. Let T be a simplicial tree where each edge is given length 1 with the global metric induced as a path metric. We fix a vertex v_0 . Given another vertex v in T , we define the “jet” subset $J(v) = \{t \in T \mid v \in [v_0, t]\}$. Let $B(v, r)$ stand for the open metric ball of radius r centered at v and $S(v, r)$ stand for its boundary sphere. We also define the subsets $J(v, l) = J(v) \cap B(v, l)$ for a positive number l , and the differences $J(v; l_1, l_2) = J(v, l_2) \setminus B(v, l_1) \setminus S(v, l_1)$, for $l_2 > l_1 > 0$.

Given a number r greater than 1, consider the collection of open subsets of T consisting of the ball $B(v_0, 2r)$ and the differences $J(v; r - 1, 3r)$ where the vertices v vary over $S(v_0, (2n - 1)r)$ for arbitrary natural numbers n . It is easy to see the following properties.

1. This collection of subsets \mathcal{V} is a covering of T . Its nerve is itself, in general, a forest of trees where the vertices can be indexed by v_0 and the vertices $v \in S(v_0, (2n - 1)r)$. The edges are the pairs (v, v') where $v' \in J(v, 2(n + 1)r)$.
2. The diameter of each set in the covering is bounded by $6r$.
3. The covering has a margin of at least r in the sense that the r -enlargement is again a tree that is isomorphic to the one in (1).

Clearly, intersecting any subset X of T with the produced covering gives a covering of X with exactly the same three properties, except possibly a forest of trees instead of a single tree. This gives a hierarchical tree-like decomposition of X with margin r and depth 1.

Algorithm 2 Distributed ProcessLowerStars on a 2D digital image.**Inputs:**

- D digital image pixels
- g grayscale values on pixels

Outputs:

- C critical cells
- V discrete vector field

Constraint:

- Memory = N Pixels

for all i and j do

Apply ProcessLowerStars algorithm to each of the following to obtain:

- $P_{i,j} \rightarrow V_{i,j}$
 - $P_{i,j} \cap P_{i+1,j} \rightarrow V_{i,j}^{i+1}$
 - $P_{i,j} \cap P_{i,j+1} \rightarrow V_{i,j}^{j+1}$
 - $(P_{i,j} \cap P_{i,j+1}) \cap (P_{i+1,j} \cap P_{i+1,j+1}) \rightarrow V_{i,j}^{\cap}$
 - $P_{i,j}^*[1,0] \cap P_{i+1,j}[1,0] \rightarrow V_{i,j}^{i+1}[1,0]$
 - $(P_{i,j} \cap P_{i,j+1})^*[1,0] \cap (P_{i+1,j} \cap P_{i+1,j+1})[1,0] \rightarrow V_{i,j}^{\cap}[1,0]$
 - $P_{i,j}[0,1] \cap P_{i,j+1}[0,1] \rightarrow V_{i,j}^{j+1}[0,1]$
 - $(P_{i,j}[0,1] \cap P_{i,j+1}[0,1]) \cap (P_{i+1,j}[0,1] \cap P_{i+1,j+1}[0,1]) \rightarrow V_{i,j}^{\cap}[0,1]$
 - $(P_{i,j}[0,1] \cap P_{i,j+1}[0,1])^*[1,0] \cap (P_{i+1,j}[0,1] \cap P_{i+1,j+1}[0,1])[1,0] \rightarrow V_{i,j}^{\cap}[1,1]$
- # Preprocessing Step to obtain discrete vector field in all patches and their overlaps

end for**for $j = 1, \dots, \ell$ do****for $i = 1, \dots, m - 1$ do**

Update: $V'_{i+1,j} = (V'_{i,j} \cup V_{i+1,j}) \setminus (V_{i,j}^{i+1} \setminus V_{i,j}^{i+1}[1,0])$
Moving horizontally along strips

end for**end for****for $j = 1, \dots, \ell$ do****for $i = 1, \dots, m - 1$ do**

Update: $(V'_{i+1,j})' = ((V_{i,j}^{j+1})' \cup V_{i+1,j}^{j+1}) \setminus (V_{i,j}^{\cap} \setminus V_{i,j}^{\cap}[1,0])$
Moving horizontally along intersection of strips

end for**end for****for $j = 1, \dots, \ell$ do****for $i = 1, \dots, m - 1$ do**

Update: $(V'_{i+1,j}[0,1])' = ((V_{i,j}^{j+1}[0,1])' \cup V_{i+1,j}^{j+1}[0,1]) \setminus (V_{i,j}^{\cap}[1,0] \setminus V_{i,j}^{\cap}[1,1])$
Moving horizontally along intersection of strips enlarged by 1 vertically

end for**end for****for $j = 1, \dots, \ell - 1$ do**

Update: $V'_{j+1} = (V'_{m,j} \cup V'_{m,j+1}) \setminus ((V'_{m,j})' \setminus (V'_{m,j}[0,1])')$
Moving vertically down strips

end for

Figure 3 is used to illustrate this notation. In this figure, v_0 is the root of the tree. Notice that the vertices labeled a_1, a_2, a_3 are all at distance 3 from v_0 . The blue, the orange, and the red subsets can be described as the jet subsets $J(a_i, 1, 4)$ for $1 \leq i \leq 3$. They overlap with the green subset which itself can be described as the jet subset $J(b_1, 1, 4)$. Notice that the

bi-colored edges represent the overlaps between the different colored subsets. This should illustrate the pattern of jet-generated covering sets and their overlaps in the whole tree. One more feature that can be seen in this figure is that with this particular set of choices of the vertices and the bounds l_1, l_2 , there are only double overlaps. There are no overlaps between jet subsets of higher multiplicities. In other words, the nerve of such a covering is one dimensional. Now it can be easily seen in the picture that the cellular enlargement of all subsets by 1 unit still has a one-dimensional nerve. The corresponding enlargements are simply the jet subsets $J(a_i, 0, 5)$, for $1 \leq i \leq 3$, and $J(b_1, 0, 5)$. This illustrates how this procedure with the specific choices we make produces a family of patches with margin 1.




Figure 3. An illustration of jet subsets $J(v; l_1, l_2)$ and margins in a tree.

For a product Π of D trees, one starts by performing this construction in each of the factors, then takes the product of the covering families to create a covering of Π . This is a hierarchical tree-like decomposition with margin r and depth D . Again, intersecting any subset X of the product with the covering subsets is a hierarchical tree-like decomposition of X with margin r and depth D .

Remark 3. Recall that a real tree in geometry is a geodesic metric space where every triangle is a tripod. Divergence behavior of geodesics in a real tree allows to generalize the above constructions in simplicial trees verbatim to real trees and their products.

In the rest of this section, we show how to use a hierarchical tree-like decomposition similar to Algorithm 2.

Suppose our complex K is a subcomplex of Π with the product cubical structure. In each factor, we assume the simplicial tree structure. Distributing the computation, one simplifies the computation by applying an algorithm to slices which are reduced in size compared to the total complex.

In our case, this process is performed inductively using the covering sets \mathcal{V} in each tree coordinate. This guarantees that all slices used in the computation form coverings

with a margin of at least r . We remind the reader that because of the assumptions, both \mathcal{V} and $\mathcal{V}[r]$ can be viewed themselves as vertices of forests of trees.

If π_i is the projection onto the i -th coordinate, we have the slices $S_{i,V} = \pi_i^{-1}(V)$ for all V in \mathcal{V}_i . Since π_i is a 1-Lipschitz function, the slices form a covering of K with a margin of at least r . We think of \mathcal{V}_i as analogues of the intervals in Algorithm 2, so the following terminology is natural.

Notation 2. Given a vector x with i -th coordinate a set V_i from \mathcal{V}_i , there are “hypercubes”

$$C_x = S_{1,V_1} \cap S_{2,V_2} \cap \dots \cap S_{D,V_D}.$$

These hypercubes can serve as cells in a grid-like structure in Π . Note that the empty set \emptyset can be made a valid value for V_* . Let x be a vector as above with the property that if \emptyset appears as a value, then all subsequent coordinates must be \emptyset . We use $d(x)$ to denote the highest index for which the value is non-empty. Now the subsets of indices

$$X(s) = \{x \in \prod \mathcal{V}_i \mid D - d(x) = s\}$$

and the corresponding coverings $\mathcal{C}(s)$ by C_x for $x \in X(s)$ for increasing values of s form the covering sets that generalize the rectangles and strips from Algorithm 2. This is precisely the inductive structure that was leveraged in the 2D plane in the case of each tree T_i being a simplicial line covered by the intervals.

6. Remarks on Applications

There is now a large library of applications of discrete gradient fields through its use to simplify computations in topological data analysis but also direct applications to concrete problems. Our treatment of its parallelization in this paper is very general. To our knowledge, all useful gradient fields that appear in the literature are uniformly local and so the methods of this paper can be applied to them.

Just to give an example of gradient fields used for motion planning in robotics, we mention a couple of recent papers of the second author. The problem in motion planning is to create an algorithm for a path a robot would take through the so-called free space of all allowable configurations, the configurations of the robot that avoid all obstructions in the physical environment. The free space is really a parameter space of feasible robot positions. One approach to modeling it is to sample the free space and build the corresponding simplicial approximation to it through some Vietoris–Rips complex built on the sample as vertices [8]. There are several ways to produce a discrete Morse function on a simplicial complex, which restricts to given values on the vertices, or is essentially equivalent to a gradient field, some being very well known, such as [9]. All of them are uniformly local algorithms. One way this can be used for practical motion planning is as follows. Convex polyhedral obstructions usually produce convex polyhedral exclusion zones complementary to the free space. One may choose a density estimator for the sample in free space, with values at every sample point, which is then possible to extend to a Morse function. It is argued in [10] that critical marginal points can act as “lighthouses” in planning a small finite collection of paths that provably contains an optimal path of smallest length which can then be easily extracted.

There are two kinds of applications of our formula in this setting. The construction of the vector field can be distributed by considering an arbitrary or antithetic covering of the free space. This allows one to construct the vector field that identifies the lighthouses in patches. There is also the option of processing the data “as needed”, for example by exploring an adjacent patch and the next lighthouse only when the robot approaches its border.

7. Discussion

There are plenty of comments in the literature that point out that by its nature, `ProcessLowerStars` is embarrassingly parallelizable, cf. Robins et al. [2] and, for another example, Section 2 of the work of Guylassy et al. [11]. In the same way, other uniformly local algorithms α are embarrassingly parallelizable. This perspective of parallelization is not that useful unless α is very expensive to compute by itself. A more urgent need is the type of distributed computation that builds partial vector fields on patches and combines them together. Since discrete vector fields are analogues of smooth vector fields, it is not surprising that they are also very locally defined and are not expensive to compute in each locality. These observations justify our distributed computation in this paper as a useful perspective on parallelization in this context.

Related to the point above, we wonder if there are useful discrete vector fields that are not uniformly local or even those that are uniformly local for $k \geq 2$ but not for $k = 1$. The authors are not aware of any in the literature. Just to mention another example of a broadly used algorithm, it is easy to see that `ExtractRaw` of King et al. [9], which is used to generate a discrete vector field from values of a Morse function on 0-cells is uniformly local for $k = 1$. If we are correct, then our method in the paper applies to all known discrete vector fields.

Our future plans include extending Theorem 2 and Algorithm 2 to the 3D case, where `ProcessLowerStars` was proven to work [2]. We will also address similar custom efficiency improvements of general formulas for another algorithm `MorseReduce` due to Mischaikow and Nanda [12]. This is a versatile algorithm using discrete Morse theory as a preprocessing tool for persistent homology computations in topological data analysis. The striking feature of `MorseReduce` is that it is dimension independent and can be applied to very general regular cellular complexes.

Our results apply to infinite complexes and infinite coverings of complexes. This comment might seem to have no practical implications, but we can use it to point out that our parallelization theorems allow a dynamic approach to processing the data. As in the robotics application described in the preceding section, one may not want to process all available patches at the same time but rather proceed one patch at a time depending on the need of a dynamic process, such as planning a path. In this case, there is a need to incorporate the new patch into the pre-existing framework. There might be infinitely many possible patches in the agnostic planning process. The theorems can be used to extend the definition of the discrete vector field to each successive patch as many times as needed.

Finally, we would like to contrast our theorems with comparable parallelization algorithms [13,14]. We are grateful to the referee who pointed out these algorithms to us. Both papers deal with discrete models approximated by smooth gradient vector fields and geometrically parallelize the computation of the associated Morse–Smale complexes. The goals of these authors are very much akin to ours. In fact, we refer the reader to additional motivation in the Related Work sections in both papers. The main distinguishing feature of our theorems from Section 3 is their generality. They apply to any uniformly local algorithm on any regular cellular complex of any dimension, while the results of the referenced papers are more specific, geared toward the computation of Morse–Smale complexes associated to gradient flows on lower dimensional manifolds. We do not say that these are special cases of the general theorems. They are certainly leaner and more efficient algorithms for that specific task, much like our Theorem 2 is not a special case of Theorem 2 and its corollaries.

Author Contributions: Conceptualization, D.L. and B.G.; methodology, D.L. and B.G.; formal analysis, D.L. and B.G.; investigation, D.L. and B.G.; writing—original draft preparation, D.L. and B.G.; writing—review and editing, D.L. and B.G.; visualization, D.L. and B.G.; supervision, D.L. and B.G.; project administration, D.L. and B.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Forman, R. *A User's Guide to Discrete Morse Theory*; Séminaire Lotharingien de Combinatoire, Institute für Mathematik, Universität Wien: Wien, Austria 2002; Volume 48, 35p.
2. Robins, V.; Wood, P.J.; Sheppard, A. Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. *IEEE Trans. Pattern Anal. Mach. Intell.* **2011**, *33*, 1646–1658. [[CrossRef](#)] [[PubMed](#)]
3. Knudson, K.P. *Morse Theory: Smooth and Discrete*; World Scientific: Singapore, 2015.
4. Carlsson, G.; Goldfarb, B. The integral K -theoretic Novikov conjecture for groups with finite asymptotic dimension. *Invent. Math.* **2004**, *157*, 405–418. [[CrossRef](#)]
5. Goldfarb, B. Singular persistent homology with geometrically parallelizable computation. *Topol. Proc.* **2020**, *55*, 273–294.
6. Dranishnikov, A. On hypersphericity of manifolds with finite asymptotic dimension. *Trans. Am. Math. Soc.* **2003**, *355*, 155–167. [[CrossRef](#)]
7. Kasprowski, D. Coarse embeddings into products of trees. *arXiv* **2021**, arXiv:1810.13361.
8. Upadhyay, A.; Wang, W.; Ekenna, C. Approximating cfree space topology by constructing Vietoris-Rips complex. In Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2019), Macau, China, 4–8 November 2019; pp. 2517–2523.
9. King, H.; Knudson, K.; Mramor, N. Generating discrete Morse functions from point data. *Exp. Math.* **2005**, *14*, 435–444. [[CrossRef](#)]
10. Upadhyay, A.; Goldfarb, B.; Ekenna, C. A topological approach to finding coarsely diverse paths. In Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2021), Prague, Czech Republic, 27 September–1 October 2021.
11. Gyulassy, A.; Bremer, P.-T.; Pascucci, V. Computing Morse-Smale Complexes with Accurate Geometry. *IEEE Trans. Vis. Comput. Gr.* **2012**, *18*, 2014–2022. [[CrossRef](#)] [[PubMed](#)]
12. Mischaikow, K.; Nanda, V. Morse Theory for Filtrations and Efficient Computation of Persistent Homology. *Discr. Comput. Geom.* **2013**, *50*, 330–353. [[CrossRef](#)]
13. Gyulassy, A.; Bremer, P.-T.; Hatmann, B.; Pascucci, V. A Practical Approach to Morse-Smale Complex Computation: Scalability and Generality. *IEEE Trans. Vis. Comput. Gr.* **2008**, *6*, 1619–1626. [[CrossRef](#)] [[PubMed](#)]
14. Shivashankar, N.; Senthilnathan, M.; Natarajan, V. Parallel Computation of 2D Morse-Smale Complexes. *IEEE Trans. Vis. Comput. Gr.* **2012**, *18*, 1757–1770. [[CrossRef](#)] [[PubMed](#)]