

Alexander Duality for Functions: the Persistent Behavior of Land and Water and Shore

Herbert Edelsbrunner*
IST Austria
Klosterneuburg, Austria
edels@ist.ac.at

Michael Kerber
IST Austria
Klosterneuburg, Austria
mkerber@ist.ac.at

ABSTRACT

This note contributes to the point calculus of persistent homology by extending Alexander duality from spaces to real-valued functions. Given a perfect Morse function $f : \mathbb{S}^{n+1} \rightarrow [0, 1]$ and a decomposition $\mathbb{S}^{n+1} = \mathbb{U} \cup \mathbb{V}$ into two $(n + 1)$ -manifolds with common boundary \mathbb{M} , we prove elementary relationships between the persistence diagrams of f restricted to \mathbb{U} , to \mathbb{V} , and to \mathbb{M} .

Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations, Computations on discrete structures*

Keywords

Algebraic topology, homology, Alexander duality, Mayer-Vietoris sequences, persistent homology, point calculus.

1. INTRODUCTION

Persistent homology is a recent extension of the classical theory of homology; see e.g. [6]. Given a real-valued function on a topological space, it measures the importance of a homology class by monitoring when the class appears and when it disappears in the increasing sequence of sublevel sets. A technical requirement is that the function be *tame*, which means it has only finitely many homological critical values, and each sublevel set has finite rank homology groups. Pairing up the births and deaths, and drawing each pair of values as a *dot* (a point in the plane), we get a multiset which we refer to as the *persistence diagram* of the function. It is a combinatorial summary of the homological

*... and Departments of Computer Science and of Mathematics, Duke University, Durham, North Carolina; DeLaunay Lab of Discrete and Computational Geometry, Yaroslavl' State University, Russia; Geomagic, Research Triangle Park, North Carolina.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'12, June 17–20, 2012, Chapel Hill, North Carolina, USA.

Copyright 2012 ACM 978-1-4503-1299-8/12/06 ...\$10.00.

information contained in the sequence of sublevel sets. If we substitute reduced for standard homology groups, we get a slightly modified *reduced persistence diagram*.

As between homology groups, we can observe relationships between persistence diagrams. A prime example is Poincaré duality, which says that the p -th and the $(n - p)$ -th homology groups of an orientable n -manifold are isomorphic. More precisely, this is true if homology is defined for field coefficients, which is what we assume throughout this paper. The extension to functions says that the diagram is symmetric with respect to reflection across the vertical axis; see [3]. Here, we change the homological dimension of a dot from p to $n - p$ whenever we reflect it across the axis. This paper contributes new relationships by extending Alexander duality from spaces to functions. To state our results, we assume a perfect Morse function, $f : \mathbb{S}^{n+1} \rightarrow [0, 1]$, which for the sphere has no critical points other than a minimum and a maximum, and a decomposition of the $(n + 1)$ -dimensional sphere into two $(n + 1)$ -manifolds with boundary, $\mathbb{S}^{n+1} = \mathbb{U} \cup \mathbb{V}$, whose intersection is the common n -manifold boundary $\mathbb{M} = \mathbb{U} \cap \mathbb{V}$. Our first result says that the reduced persistence diagrams of f restricted to \mathbb{U} and to \mathbb{V} are reflections of each other. We call this the Land and Water Theorem. Our second result relates land with shore. Ignoring some modifications, it says that the persistence diagram of f restricted to \mathbb{M} is the disjoint union of the diagram of f restricted to \mathbb{U} and of its reflection. The modifications become unnecessary if we assume that the minimum and maximum of f both belong to a common component of \mathbb{V} . We call this the Euclidean Shore Theorem.

In the example that justifies the title of this paper, and the names of our theorems, we let \mathbb{U} be the planet Earth, not including the water and the air. To a coarse approximation, \mathbb{U} is homeomorphic to a 3-ball, sitting inside the Universe, which we model as a 3-sphere, \mathbb{S} . The function we consider is the negative gravitational potential of the Earth, which is defined on the entire Universe. The sea is then a sublevel set of this function restricted to \mathbb{V} , which is the closure of $\mathbb{S} - \mathbb{U}$. With these definitions, our results relate the persistence diagram of the gravitational potential restricted to the Earth with the shape of the sea as its water level rises. Also, the Euclidean Shore Theorem applies, unmodified, expanding the relationship to include the sea floor, which is swept out by the shoreline as the water level rises.

Besides developing the mathematical theory of persistent homology, there are pragmatic reasons for our interest in the extension of Alexander duality to functions. Persistence has fast algorithms, so that the bulk of the work is often

in the construction of the space and the function for which we compute persistence. A point in case is the analysis of the biological process of cell segregation started in [7]. Modeling the process as a subset of space-time, the function of interest is time which, after compactifying space-time to \mathbb{S}^4 , has no critical points other than a minimum and a maximum. The subset \mathbb{U} of \mathbb{S}^4 is a union of cells times time, whose boundary is a 3-manifold. We can represent \mathbb{U} by a 1-parameter family of alpha complexes, whose disjoint union has the same homotopy type; see e.g. [6, Chapter III]. However, the boundary of that disjoint union is not necessarily a 3-manifold. Using our Euclidean Shore Theorem, we can compute the persistence diagram of the function on the 3-manifold without ever constructing the 3-manifold.

Outline. Sections 2 and 3 introduce background on homology and persistent homology. Sections 4 and 5 present our two results. Section 6 concludes the paper.

2. HOMOLOGY

Starting with a brief introduction of classical homology groups, we present the relevant background on Alexander duality and Mayer-Vietoris sequences. More comprehensive discussions of these topics can be found in textbooks of algebraic topology, such as [8, 10].

Background. The p -dimensional homology of a topological space, \mathbb{X} , is a mathematical language to define, count, and reason about the p -dimensional connectivity of \mathbb{X} . There are different but essentially equivalent theories depending on the choices one makes in the representation of the space, the selection of cycles, and the meaning of addition. For our purpose, the most elementary of these theories will suffice: a simplicial complex, K , that triangulates \mathbb{X} , formal sums of p -simplices with zero boundary as p -cycles, and adding with coefficients in a field, F . Most algorithms on homology assume this model, in particular the ones developed within persistent homology. In this model, we call a formal sum of p -simplices a p -chain, a p -cycle if its boundary is empty, and a p -boundary if it is the boundary of a $(p+1)$ -chain. The p -boundaries form a subgroup of the p -cycles, which form a subgroup of the p -chains: $B_p \subseteq Z_p \subseteq C_p$. The p -th homology group is the quotient of the p -cycles over the p -boundaries: $H_p = Z_p/B_p$. Its elements are sets of p -cycles that differ from each other by p -boundaries. More fully, we denote the p -th homology group by $H_p(K, F)$, or by $H_p(\mathbb{X}, F)$ to emphasize that the group is independent of the simplicial complex we choose to triangulate the space. However, we will fix an arbitrary field F and drop it from the notation. For field coefficients, $H_p(\mathbb{X})$ is necessarily a vector space, which is fully described by its rank, $\beta_p(\mathbb{X}) = \text{rank } H_p(\mathbb{X})$ such that $H_p(\mathbb{X}) \simeq F^{\beta_p(\mathbb{X})}$, called the p -th Betti number of \mathbb{X} . Finally, we will often drop the dimension from the notation and write $H(\mathbb{X}) = \bigoplus_{p \in \mathbb{Z}} H_p(\mathbb{X})$ for the direct sum.

Besides standard homology, we will frequently use *reduced homology groups*, $\tilde{H}_p(\mathbb{X})$, which are isomorphic to the non-reduced groups except possibly for dimensions $p = 0, -1$. To explain the difference, we note that $\beta_0(\mathbb{X})$ counts the components of \mathbb{X} . In contrast, $\tilde{\beta}_0(\mathbb{X}) = \text{rank } \tilde{H}_0(\mathbb{X})$ counts the gaps between components or, equivalently, the edges that are needed to merge the components into one. Hence, $\tilde{\beta}_0(\mathbb{X}) = \beta_0(\mathbb{X}) - 1$, except when \mathbb{X} is empty, in which case $\tilde{\beta}_0(\mathbb{X}) = \beta_0(\mathbb{X}) = 0$. To distinguish this case from a single component, we have $\tilde{\beta}_{-1}(\mathbb{X})$ equal to 1 if $\mathbb{X} = \emptyset$, and

equal to 0, otherwise. Furthermore, we use relative homology, which is defined for pairs of spaces, $\mathbb{Y} \subseteq \mathbb{X}$. Taking a pair relaxes the requirement of a chain to be called a cycle, namely whenever its boundary is contained in \mathbb{Y} , which includes the case when the boundary is empty. We write $H_p(\mathbb{X}, \mathbb{Y})$ for the p -th relative homology group of the pair, and $\beta_p(\mathbb{X}, \mathbb{Y}) = \text{rank } H_p(\mathbb{X}, \mathbb{Y})$ for the p -th relative Betti number. As before, we will suppress the dimension from the notation by introducing $H(\mathbb{X}, \mathbb{Y}) = \bigoplus_{p \in \mathbb{Z}} H_p(\mathbb{X}, \mathbb{Y})$.

As examples, consider the $(n+1)$ -dimensional sphere, $\mathbb{S} = \mathbb{S}^{n+1}$, and a closed hemisphere, $\mathbb{H} \subseteq \mathbb{S}$, which is a ball of dimension $n+1$. In standard homology, we have $\beta_0(\mathbb{S}) = \beta_0(\mathbb{H}) = \beta_{n+1}(\mathbb{S}) = 1$ while all other Betti numbers are zero. In reduced homology, we have $\tilde{\beta}_{n+1}(\mathbb{S}) = 1$ while all other reduced Betti numbers are zero. In particular, $\tilde{\beta}_p(\mathbb{H}) = 0$ for all p . In relative homology, we have $\beta_{n+1}(\mathbb{S}, \mathbb{H}) = 1$ while all other relative Betti numbers are zero. In particular, $\beta_0(\mathbb{S}, \mathbb{H}) = 0$, which may be confusing at first but makes sense because every point on the sphere can be connected by a path to a point in the hemisphere and is thus a 0-boundary.

Alexander duality. Recall that a *perfect Morse function* is one whose number of critical points equals the sum of Betti numbers of the space. For a sphere, this number is 2: a minimum and a maximum. Throughout the remainder of this section, we assume a perfect Morse function $f : \mathbb{S} \rightarrow [0, 1]$, whose values at the minimum and the maximum are 0 and 1. We also assume two $(n+1)$ -manifolds with boundary, \mathbb{U} and \mathbb{V} , whose union is \mathbb{S} , and whose intersection is the common boundary, \mathbb{M} . For technical reasons, we require that \mathbb{U} , \mathbb{V} , and \mathbb{M} are embedded in the sphere in a way that is compatible with taking the simplicial homology and with applying the Mayer-Vietoris sequence and Alexander duality. Specifically, we assume that \mathbb{U} and \mathbb{V} are triangulable closed cofibrations. In other words, there is a deformation retraction from an open neighborhood of \mathbb{U} to \mathbb{U} , and similarly for \mathbb{V} . Furthermore, we assume that the restriction of f to \mathbb{M} is tame, which implies that its restrictions to \mathbb{U} and \mathbb{V} are also tame. For each $t \in \mathbb{R}$, we write $\mathbb{S}_t = f^{-1}[0, t]$ for the *sublevel set* of f , and $\mathbb{U}_t = \mathbb{U} \cap \mathbb{S}_t$, $\mathbb{V}_t = \mathbb{V} \cap \mathbb{S}_t$, $\mathbb{M}_t = \mathbb{M} \cap \mathbb{S}_t$ for the sublevel sets of the restrictions of f to \mathbb{U} , \mathbb{V} , \mathbb{M} . Similarly, we write $\mathbb{S}^t = f^{-1}[t, 1]$ for the *superlevel set*, and $\mathbb{U}^t = \mathbb{U} \cap \mathbb{S}^t$, $\mathbb{V}^t = \mathbb{V} \cap \mathbb{S}^t$, $\mathbb{M}^t = \mathbb{M} \cap \mathbb{S}^t$ for the superlevel sets of the restrictions. Since we will apply Mayer-Vietoris and Alexander duality to sub- and superlevel sets, we make the same technical assumptions as formulated above for \mathbb{U} , \mathbb{V} , and \mathbb{M} . For later reference, we call a pair of $(n+1)$ -manifolds with boundary that satisfy those conditions a *complementary decomposition* of \mathbb{S} , noting that the two $(n+1)$ -manifolds are indeed complementary except for the shared boundary.

The basic version of Alexander duality is a statement about a complementary decomposition of the sphere; see [10, page 424]. More specifically, it states that $H_q(\mathbb{U})$ and $H_p(\mathbb{V})$ are isomorphic, where $q = n - p$, except for $p = 0$ and $q = 0$ when the 0-dimensional group has an extra generator. This implies

$$\beta_0(\mathbb{V}) = \beta_n(\mathbb{U}) + 1, \quad (1)$$

$$\beta_p(\mathbb{V}) = \beta_q(\mathbb{U}), \quad (2)$$

$$\beta_n(\mathbb{V}) = \beta_0(\mathbb{U}) - 1, \quad (3)$$

for $1 \leq p \leq n - 1$. We will also need the version that

deals with a complementary decomposition of the $(n + 1)$ -dimensional ball, \mathbb{B}^{n+1} ; see [10, page 426]. Let $t \in (0, 1)$ denote a regular value of $f|_{\mathbb{M}}$. Note first that \mathbb{S}_t is homeomorphic to \mathbb{B}^{n+1} . By excision, the homology groups of \mathbb{U}_t relative to $\mathbb{U}_t \cap f^{-1}(t)$ are isomorphic to those of \mathbb{U} relative to \mathbb{U}^t . Alexander duality states that $\mathbf{H}_q(\mathbb{U}, \mathbb{U}^t)$ and $\mathbf{H}_p(\mathbb{V}_t)$ are isomorphic, where $q + p = n$, as before, except for $p = 0$, when \mathbb{V}_t has an extra component. This implies

$$\beta_0(\mathbb{V}_t) = \beta_n(\mathbb{U}, \mathbb{U}^t) + 1, \quad (4)$$

$$\beta_p(\mathbb{V}_t) = \beta_q(\mathbb{U}, \mathbb{U}^t), \quad (5)$$

for $1 \leq p \leq n$.

Mayer-Vietoris. We can connect the homology groups of \mathbb{U} and \mathbb{V} with those of \mathbb{M} and \mathbb{S} using the Mayer-Vietoris sequence of the decomposition. This sequence is *exact*, meaning the image of every map equals the kernel of the next map. Counting the trivial homology groups, the sequence is infinitely long, with three terms per dimension:

$$\dots \rightarrow \mathbf{H}_{p+1}(\mathbb{S}) \rightarrow \mathbf{H}_p(\mathbb{M}) \rightarrow \mathbf{H}_p(\mathbb{U}) \oplus \mathbf{H}_p(\mathbb{V}) \rightarrow \mathbf{H}_p(\mathbb{S}) \rightarrow \dots,$$

where the maps between homology groups of the same dimension are induced by the inclusions. The only non-trivial homology groups of \mathbb{S} are in dimensions 0 and $n + 1$, with ranks $\beta_0(\mathbb{S}) = \beta_{n+1}(\mathbb{S}) = 1$. It follows that for $1 \leq p \leq n - 1$, the groups defined by \mathbb{M} , \mathbb{U} , \mathbb{V} are surrounded by trivial groups. This implies that the groups of \mathbb{M} and the direct sums of the groups of \mathbb{U} and \mathbb{V} are isomorphic. For $p = 0$ and $p = n$, the non-trivial homology groups of \mathbb{S} prevent this isomorphism, and we get

$$\beta_0(\mathbb{M}) = \beta_0(\mathbb{U}) + \beta_0(\mathbb{V}) - 1, \quad (6)$$

$$\beta_p(\mathbb{M}) = \beta_p(\mathbb{U}) + \beta_p(\mathbb{V}), \quad (7)$$

$$\beta_n(\mathbb{M}) = \beta_n(\mathbb{U}) + \beta_n(\mathbb{V}) + 1, \quad (8)$$

for $1 \leq p \leq n - 1$. We also have a Mayer-Vietoris sequence for the spaces \mathbb{S}_t , \mathbb{U}_t , \mathbb{V}_t and \mathbb{M}_t :

$$\dots \rightarrow \mathbf{H}_{p+1}(\mathbb{S}_t) \rightarrow \mathbf{H}_p(\mathbb{M}_t) \rightarrow \mathbf{H}_p(\mathbb{U}_t) \oplus \mathbf{H}_p(\mathbb{V}_t) \rightarrow \mathbf{H}_p(\mathbb{S}_t) \rightarrow \dots$$

For $0 \leq t < 1$, \mathbb{S}_t is a point or a closed $(n + 1)$ -dimensional ball, and its only non-trivial homology group is of dimension 0, with rank $\beta_0(\mathbb{S}_t) = 1$. It follows that for $1 \leq p \leq n$, the groups defined by \mathbb{M}_t , \mathbb{U}_t , \mathbb{V}_t are surrounded by trivial groups, which implies that the groups of \mathbb{M}_t and the direct sums of the groups of \mathbb{U}_t and \mathbb{V}_t are isomorphic. We get

$$\beta_0(\mathbb{M}_t) = \beta_0(\mathbb{U}_t) + \beta_0(\mathbb{V}_t) - 1 \quad (9)$$

$$\beta_p(\mathbb{M}_t) = \beta_p(\mathbb{U}_t) + \beta_p(\mathbb{V}_t), \quad (10)$$

for $1 \leq p \leq n$. We may also consider the relative homology groups, again connected by a Mayer-Vietoris sequence:

$$\dots \rightarrow \mathbf{H}_{p+1}(\mathbb{S}, \mathbb{S}^t) \rightarrow \mathbf{H}_p(\mathbb{M}, \mathbb{M}^t)$$

$$\rightarrow \mathbf{H}_p(\mathbb{U}, \mathbb{U}^t) \oplus \mathbf{H}_p(\mathbb{V}, \mathbb{V}^t) \rightarrow \mathbf{H}_p(\mathbb{S}, \mathbb{S}^t) \rightarrow \dots$$

For $0 < t \leq 1$, the only non-trivial homology group of the pair $(\mathbb{S}, \mathbb{S}^t)$ is in dimension $n + 1$, which implies

$$\beta_p(\mathbb{M}, \mathbb{M}^t) = \beta_p(\mathbb{U}, \mathbb{U}^t) + \beta_p(\mathbb{V}, \mathbb{V}^t), \quad (11)$$

$$\beta_n(\mathbb{M}, \mathbb{M}^t) = \beta_n(\mathbb{U}, \mathbb{U}^t) + \beta_n(\mathbb{V}, \mathbb{V}^t) + 1, \quad (12)$$

for $0 \leq p \leq n - 1$.

Separating manifold. Combining Alexander duality with Mayer-Vietoris, we get relations between the homology of

the sublevel sets of \mathbb{M} and \mathbb{U} . More specifically, we get

$$\beta_p(\mathbb{M}) = \beta_p(\mathbb{U}) + \beta_q(\mathbb{U}), \quad (13)$$

$$\beta_p(\mathbb{M}_t) = \beta_p(\mathbb{U}_t) + \beta_q(\mathbb{U}, \mathbb{U}^t), \quad (14)$$

$$\beta_p(\mathbb{M}, \mathbb{M}^t) = \beta_p(\mathbb{U}, \mathbb{U}^t) + \beta_q(\mathbb{U}_t), \quad (15)$$

for all $0 \leq p \leq n$. Here, we combine (6), (7), (8) with (1), (2), (3) to get (13). Similarly, we combine (9), (10) with (4), (5) to get (14). Finally, we exploit the symmetry between \mathbb{U} and \mathbb{V} and combine (11), (12) with (4), (5) to get (15). Note that (14) and (15) imply $\beta_p(\mathbb{M}_t) = \beta_q(\mathbb{M}, \mathbb{M}^t)$, which is a consequence of Lefschetz duality; see [10].

An example. We illustrate the above relationships with an example. Let \mathbb{M} be a 2-dimensional torus in \mathbb{S}^3 . Accordingly, \mathbb{U} and \mathbb{V} are the two solid tori that decompose the 3-sphere and intersect in \mathbb{M} . This is sketched in Figure 1, where we assume that \mathbb{U} is the part of space surrounded by \mathbb{M} , while \mathbb{V} is the space surrounding \mathbb{M} . The only non-zero Betti numbers of the solid torus are $\beta_0(\mathbb{U}) = \beta_1(\mathbb{U}) = 1$. We get

$$\beta_0(\mathbb{M}) = \beta_0(\mathbb{U}) + \beta_2(\mathbb{U}) = 1,$$

$$\beta_1(\mathbb{M}) = \beta_1(\mathbb{U}) + \beta_1(\mathbb{U}) = 2,$$

$$\beta_2(\mathbb{M}) = \beta_2(\mathbb{U}) + \beta_0(\mathbb{U}) = 1$$

from (13). These are the correct Betti numbers of the 2-dimensional torus. Next, choose t so that the sublevel set of \mathbb{U} is half the solid torus. Its only non-zero Betti number is




Figure 1: The level set defined by t splits the solid torus into two halves.

$\beta_0(\mathbb{U}_t) = 1$, and the only non-zero relative Betti number of the pair is $\beta_1(\mathbb{U}, \mathbb{U}^t) = 1$. We thus get

$$\beta_0(\mathbb{M}_t) = \beta_0(\mathbb{U}_t) + \beta_2(\mathbb{U}, \mathbb{U}^t) = 1,$$

$$\beta_1(\mathbb{M}_t) = \beta_1(\mathbb{U}_t) + \beta_1(\mathbb{U}, \mathbb{U}^t) = 1,$$

$$\beta_2(\mathbb{M}_t) = \beta_2(\mathbb{U}_t) + \beta_0(\mathbb{U}, \mathbb{U}^t) = 0$$

from (14). These are the correct Betti numbers of the sublevel set of \mathbb{M} . Finally, we get

$$\beta_0(\mathbb{M}, \mathbb{M}^t) = \beta_2(\mathbb{U}_t) + \beta_0(\mathbb{U}, \mathbb{U}^t) = 0,$$

$$\beta_1(\mathbb{M}, \mathbb{M}^t) = \beta_1(\mathbb{U}_t) + \beta_1(\mathbb{U}, \mathbb{U}^t) = 1,$$

$$\beta_2(\mathbb{M}, \mathbb{M}^t) = \beta_0(\mathbb{U}_t) + \beta_2(\mathbb{U}, \mathbb{U}^t) = 1$$

from (15). These are the correct relative Betti numbers of the pair $(\mathbb{M}, \mathbb{M}^t)$. Note again that the vector of Betti numbers of \mathbb{M}_t is the reverse of that of $(\mathbb{M}, \mathbb{M}^t)$.




Figure 2: The persistence diagram of the function restricted to the solid torus (dots drawn as circles), and to the complementary solid torus (dots drawn as squares), as sketched on the left. The white dots belong to the reduced diagram, the black dots belong to the standard diagram, and the shaded dots belong to both.

3. PERSISTENCE

Starting with a brief introduction of persistent homology, we review its combinatorial expression as a point calculus; see [6] and [1, 2] for more comprehensive discussions.

Background. We take the step from homology to persistent homology by replacing a space with the sequence of (closed and open) sublevel sets of a function on the space. To explain this, let \mathbb{X} be compact and $g : \mathbb{X} \rightarrow [0, 1]$ tame. As before, we write $\mathbb{X}_t = g^{-1}[0, t]$ for the sublevel set and $\mathbb{X}^t = g^{-1}[t, 1]$ for the superlevel set defined by t . Since g is tame, it has finitely many homological critical values, s_0 to s_m . We assume w.l.o.g. that $s_0 = 0$ and $s_m = 1$, else, we add them to the list. We interleave the s_i with homological regular values t_i , such that $s_0 < t_0 < s_1 < \dots < t_{m-1} < s_m$. Taking the homology of closed and open sublevel sets at the regular values, we get the *persistent module*

$$\mathbb{X}_0 \rightarrow \dots \rightarrow \mathbb{X}_{m-1} \rightarrow \mathbb{X}_m \rightarrow \mathbb{X}_{m+1} \rightarrow \dots \rightarrow \mathbb{X}_{2m},$$

where \mathbb{X}_i is $\mathbb{H}(\mathbb{X}_{t_i})$, if $0 \leq i < m$, $\mathbb{H}(\mathbb{X})$, if $i = m$, and $\mathbb{H}(\mathbb{X}, \mathbb{X}^{t_{2m-i}})$, if $m < i \leq 2m$. For notational convenience, we add the trivial groups, $\mathbb{X}_{-1} = 0$ and $\mathbb{X}_{2m+1} = 0$, at the beginning and end. The maps connecting the homology groups are induced by the inclusions $\mathbb{X}_s \subseteq \mathbb{X}_t$ and $\mathbb{X}^t \subseteq \mathbb{X}^s$, whenever $s \leq t$. The maps compose and we write $\mathbf{g}_{i,j} : \mathbb{X}_i \rightarrow \mathbb{X}_j$. Reading the module from left to right, we see homology classes appear and disappear. To understand these events, we say a class $\alpha \in \mathbb{X}_i$ is *born* at \mathbb{X}_i if it does not belong to the image of $\mathbf{g}_{i-1,i}$. If furthermore $\mathbf{g}_{i,j}(\alpha)$ belongs to the image of $\mathbf{g}_{i-1,j}$ but $\mathbf{g}_{i,j-1}(\alpha)$ does not belong to the image of $\mathbf{g}_{i-1,j-1}$, then we say α *dies entering* \mathbb{X}_j . Since the module starts and ends with trivial groups, every homology class has well-defined birth and death values. Every event is associated with the immediately preceding homological critical value, namely with s_i if the event is at \mathbb{X}_i , and with s_{m-i} if the event is at \mathbb{X}_{m+i+1} , for $0 \leq i \leq m-1$. Events at \mathbb{X}_m are associated with $s_m = 1$. We represent a class by a dot in the plane whose two coordinates mark its birth and its death. More precisely, the coordinates signal the increase and decrease of Betti numbers, and the dot represents an entire coset of classes that are born and die with α . All dots have coordinates in $[0, 1]$, by construction.

Note that we use each s_i twice, once during the first pass

in which the sublevel set grows from \emptyset to \mathbb{X} , and then again during the second pass in which \mathbb{X} minus the superlevel set shrinks back from \mathbb{X} to \emptyset . When we collect the dots, we separate the passes by drawing each coordinate axis twice. The result is a multiset of dots, which we refer to as the *persistence diagram* of the function, and denote as $\text{Dgm}(g)$, or $\text{Dgm}_p(g)$ if we restrict ourselves to the dots representing p -dimensional classes; see Figure 2, where we label each dot with the dimension of the classes it represents. We further distinguish four regions within the diagram: the *ordinary*, the *horizontal*, the *vertical*, and the *relative subdiagrams*, denoted as $\text{Ord}(g)$, $\text{Hor}(g)$, $\text{Vcl}(g)$, and $\text{Rel}(g)$. For example, a dot belongs to the ordinary subdiagram if its birth and death both happen during the first pass; see again Figure 2.

Reduced persistence diagrams. It is sometimes convenient to use reduced instead of standard homology groups in the filtration. There are small differences caused by the (-1) -dimensional class, which exists if the space is empty. As a first step, we introduce reduced versions of the relative homology groups, which are isomorphic to the standard relative groups, for all p . To define them, let ω be a new (dummy) vertex, write $\omega \cdot \mathbb{X}^t$ for the cone of ω over \mathbb{X}^t , and let $\mathbb{X}_\omega^t = \mathbb{X} \cup \omega \cdot \mathbb{X}^t$ be the result of gluing the cone along \mathbb{X}^t to \mathbb{X} . Then the *reduced relative homology group* of the pair is $\tilde{\mathbb{H}}_p(\mathbb{X}, \mathbb{X}^t) = \tilde{\mathbb{H}}_p(\mathbb{X}_\omega^t)$. For example, if $\mathbb{X}^t = \emptyset$, then ω forms a separate component, so that the reduced 0-th Betti number is equal to the number of components of \mathbb{X} , just as $\beta_0(\mathbb{X}, \mathbb{X}^t)$. The *reduced persistence diagram* is now defined as before, but for the filtration of reduced homology groups. To describe this, we suppress the dimension and write $\tilde{\mathbb{X}}_i$ for $\tilde{\mathbb{H}}(\mathbb{X}_{t_i})$, if $0 \leq i < m$, for $\tilde{\mathbb{H}}(\mathbb{X})$, if $i = m$, and for $\tilde{\mathbb{H}}(\mathbb{X}_\omega^{t_{2m-i}})$, if $m < i \leq 2m$. The resulting sequence of reduced homology groups, from $\tilde{\mathbb{X}}_0$ to $\tilde{\mathbb{X}}_{2m}$, is connected from left to right by maps induced by inclusion. Finally, we define $\tilde{\text{Dgm}}(g)$ by matching up the births and the deaths and by drawing each pair as a dot in the plane. Similar to before, we write $\tilde{\text{Dgm}}_p(g)$ when we restrict ourselves to homology classes of dimension p .

For dimension $p \geq 1$, the reduced diagrams are the same as the standard ones, simply because the groups are the same. More formally, the persistence diagrams are the same because the two persistence modules form a commutative diagram with vertical isomorphisms:

$$\begin{array}{ccccccc} 0 & \rightarrow & \mathbb{X}_0^p & \rightarrow & \mathbb{X}_1^p & \rightarrow & \dots \rightarrow \mathbb{X}_{2m}^p \rightarrow 0 \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & \downarrow \\ 0 & \rightarrow & \tilde{\mathbb{X}}_0^p & \rightarrow & \tilde{\mathbb{X}}_1^p & \rightarrow & \dots \rightarrow \tilde{\mathbb{X}}_{2m}^p \rightarrow 0; \end{array}$$

see the Persistence Equivalence Theorem [6, page 159]. This is not true for $p = -1$, where the standard diagram is empty, while the reduced diagram contains a single dot marking the transition from an empty to a non-empty space. To describe the difference for $p = 0$, we call a dot (u, w) in the 0-th standard persistence diagram *extreme* if no other dot has a first coordinate smaller than u and a second coordinate larger than w . Here, we assume for simplicity that no two dots in the 0-th diagram share the same first coordinate or the same second coordinate.

We note that only dots in the horizontal subdiagram can be extreme. The non-extreme dots also belong to the reduced diagram of g , while the extreme dots exchange their coordinates to form new dots in the reduced diagram; see Figure 2. The way the coordinates are exchanged will be important later, so we describe the details now.




Figure 3: The cascade combines the coordinates of the four black dots with 0 and 1 to form the five white dots.

Cascades. We let $\ell + 1$ be the number of extreme dots and, for a reason that will become clear shortly, denote their coordinates with different indices as (u_k, w_{k+1}) , for $0 \leq k \leq \ell$, ordering them such that $u_0 < u_1 < \dots < u_\ell$. By definition of extreme, the corresponding second coordinates satisfy $w_1 < w_2 < \dots < w_{\ell+1}$. Observe that u_0 is the globally minimum value and $w_{\ell+1}$ is the globally maximum value. Let us now construct the reduced diagram, focusing on dimension 0. We start the upward pass with the birth of the (-1) -dimensional homology class at 0, which dies at u_0 , so we have $(0, u_0)$ in $\tilde{\text{Dgm}}_{-1}(g)$. The minima of all other connected components of \mathbb{X} mark the births of 0-dimensional classes. They can be interpreted as gaps between components. We start the downward pass with the birth of a 0-dimensional class at 1, which dies at $w_{\ell+1}$, so we have $(1, w_{\ell+1})$ in $\tilde{\text{Dgm}}_0(g)$. The maxima of all other connected components mark the deaths of 0-dimensional classes. If w is the second coordinate of a non-extreme dot, then it marks the death of the gap that opened up when we passed the corresponding first coordinate, u , during the upward pass. Hence, (u, w) belongs to $\tilde{\text{Dgm}}_0(g)$, and we note that this dot also belongs to the standard diagram. However, if $w = w_k$ is the second coordinate of an extreme dot, then it marks the death of the gap that has opened up when we passed the minimum of the next component with extreme dot during the upward pass. This minimum value is u_k , so (u_k, w_k) belongs to $\tilde{\text{Dgm}}_0(g)$, for $1 \leq k \leq \ell$.

In summary, we see that contiguous extreme dots exchange one pair of coordinates, using 0 on the left and 1 on the right to complete the process. More precisely, the reduced diagram is the same as the standard one, except for substituting the dots (17) for the dots (16):

$$(u_0, w_1), (u_1, w_2), \dots, (u_{\ell-1}, w_\ell), (u_\ell, w_{\ell+1}), \quad (16)$$

$$(0, u_0), (u_1, w_1), (u_2, w_2), \dots, (u_\ell, w_\ell), (1, w_{\ell+1}). \quad (17)$$

For later use, we call this substitution a *cascade* and write $\tilde{\text{Dgm}}(g) = \text{Dgm}(g)^C$; see Figure 3. Note that all dots belong to the horizontal and vertical subdiagrams of dimension 0, except for $(1, w_{\ell+1})$, which belongs to the relative subdiagram of dimension 0, and $(0, u_0)$, which belongs to the ordinary subdiagram of dimension -1 . We will often have $u_0 = 0$ and $w_{\ell+1} = 1$, in which case the first and the last dots in (17) lie on the baseline and can be ignored.

Point calculus. As described in [1, 2], the persistence diagram can be harvested for a wealth of homological information. This includes the ranks of the homology groups of the

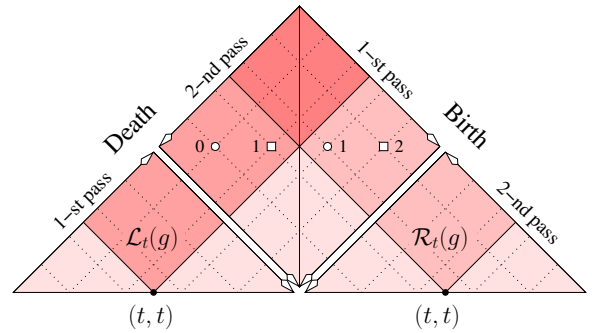


Figure 4: The persistence diagram of the function $g = f|_{\mathbb{M}}$, as defined in Figure 1; the circle dots also belong to the diagram of $f|_{\mathbb{U}}$, but the square dots do not. The rectangle on the left, $\mathcal{L}_t(g)$, represents the homology of \mathbb{M}_t , while the rectangle on the right, $\mathcal{R}_t(g)$, represents the homology of $(\mathbb{M}, \mathbb{M}^t)$; compare with the Betti numbers computed for the example at the end of Section 2.

sets \mathbb{X}_t and of the pairs $(\mathbb{X}, \mathbb{X}^t)$. To explain this, we write $\mathcal{L}_t^p(g)$ for the multiset of dots in the rectangle with lower corner (t, t) on the base edge of $\text{Ord}_p(g)$. Similarly, we write $\mathcal{R}_t^p(g)$ for the multiset of dots if the lower corner, (t, t) , lies on the base edge of $\text{Rel}_p(g)$; see Figure 4. Often, we drop the dimension from the notation and write $\mathcal{L}_t(g)$ for the disjoint union of the multisets $\mathcal{L}_t^p(g)$, over all p , and similar for $\mathcal{R}_t(g)$. We can read the Betti numbers of \mathbb{X}_t by counting the dots in $\mathcal{L}_t(g)$, and we can read the relative Betti numbers of $(\mathbb{X}, \mathbb{X}^t)$ by counting the dots in $\mathcal{R}_t(g)$. Writing $\tilde{\mathcal{L}}_t^p(g)$ and $\tilde{\mathcal{R}}_t^p(g)$ for the corresponding multisets in the reduced diagrams, we note that the same relations hold between the reduced Betti numbers and the reduced persistence diagrams.

Betti numbers and persistence diagrams. To motivate the first result of this paper, we rewrite (1) to (5) for reduced Betti numbers:

$$\tilde{\beta}_p(\mathbb{V}_t) = \beta_q(\mathbb{U}, \mathbb{U}^t), \quad (18)$$

for all $0 \leq p \leq n$ and all regular values t . Observe also that $\mathbb{V}_t = \emptyset$ iff $\mathbb{U} - \mathbb{U}^t = \mathbb{S} - \mathbb{S}^t$, which implies that (18) also holds for $p = -1$. In words, for every dimension p and every regular value t , the rectangles $\tilde{\mathcal{L}}_t^p(f|_{\mathbb{V}})$ and $\tilde{\mathcal{R}}_t^p(f|_{\mathbb{U}})$ have equally many dots. This 1-parameter family of relations is satisfied if the reduced diagrams of $f|_{\mathbb{V}}$ and $f|_{\mathbb{U}}$ are reflections of each other. This is our first result, stated as the Land and Water Theorem in Section 4. We see an illustration in Figure 2, which shows the reduced and non-reduced diagrams of the function f in Figure 1 restricted to the solid torus, \mathbb{U} , and to the complementary solid torus, \mathbb{V} . Removing the black dots, we are left with the two reduced diagrams, which are indeed reflections of each other.

To motivate our second result, we consider the relations (13) to (15). They say that for every regular value t , the rectangle $\mathcal{L}_t^p(f|_{\mathbb{M}})$ has the same number of dots as $\mathcal{R}_t^q(f|_{\mathbb{U}})$ and $\mathcal{L}_t^p(f|_{\mathbb{U}})$ together. Similarly, $\mathcal{R}_t^q(f|_{\mathbb{M}})$ has the same number of dots as $\mathcal{L}_t^q(f|_{\mathbb{U}})$ and $\mathcal{R}_t^p(f|_{\mathbb{U}})$ together. This 1-parameter family of relations is satisfied if the persistence diagram of $f|_{\mathbb{M}}$ is the disjoint union of the diagram of $f|_{\mathbb{U}}$ and of its reflection. This is our second result, stated as the Euclidean Shore Theorem in Section 5. It is illustrated in Figure 4,

which shows the persistence diagram of the function $f|_{\mathbb{M}}$ in Figure 1. Comparing Figures 2 and 4, we see that the two circle dots in Figure 4 also belong to the diagram of $f|_{\mathbb{U}}$, whereas the two square dots are reflections of the circles.

While (1) to (5) and (13) to (15) motivate our two results, these relations are not sufficient to prove them. Indeed, (13) to (15) hold in general, but the Euclidean Shore Theorem requires that the minimum and maximum of the perfect Morse function belong to a common component of \mathbb{V} . If this condition is violated, then we can have coordinate exchanges among the dots that contradict the Euclidean Shore Theorem without affecting the relations (13) to (15). An example of this phenomenon can be seen in Figure 5, which will be discussed in Section 5.

4. LAND AND WATER

In this section, we present our first result, which extends Alexander duality from spaces to functions.

Compatible pairings. A *pairing* between two finite-dimensional vector spaces X and Y over the field F is a bilinear map $\langle \cdot, \cdot \rangle : X \times Y \rightarrow F$. The pairing is *non-degenerate* if every $\xi \in X$ has at least one $\eta \in Y$ with $\langle \xi, \eta \rangle = 1$, and every $\eta \in Y$ has at least one $\xi \in X$ with $\langle \xi, \eta \rangle = 1$. A non-degenerate pairing implies an isomorphism between the vector spaces, by choosing bases of both spaces and reducing the induced matrix to diagonal form.

Now consider two filtrations, each consisting of $m + 1$ finite-dimensional vector spaces over F , and non-degenerate pairings connecting the filtrations contravariantly:

$$\begin{array}{ccccccc} X_0 & \rightarrow & X_1 & \rightarrow & \dots & \rightarrow & X_m \\ \times & & \times & & & & \times \\ Y_m & \leftarrow & Y_{m-1} & \leftarrow & \dots & \leftarrow & Y_0. \end{array}$$

Write $g_i^j : X_i \rightarrow X_j$ and $h_i^j : Y_i \rightarrow Y_j$ for the maps upstairs and downstairs. We call the pairings *compatible* with these maps if $\langle \xi, \eta \rangle = \langle \xi', \eta' \rangle$ for every $\xi' \in X_i$ and $\eta \in Y_{m-j}$, where $\xi = g_i^j(\xi')$ and $\eta' = h_{m-j}^{m-i}(\eta)$. The reason for considering compatible pairings is the following result.

CONTRAVARIANT PE THEOREM. *Let X_0 to X_m and Y_0 to Y_m be two filtrations contravariantly connected by non-degenerate and compatible pairings. Then there is a bijection between the groups upstairs and downstairs such that a class is born at X_i and dies entering X_j iff the corresponding class is born at Y_{m-j+1} and dies entering Y_{m-i+1} .*

PROOF. Let X_i^* be the dual vector space of X_i , that is, the space of homomorphisms from X_i to F , and similarly let Y_{m-j}^* be the dual vector space of Y_{m-j} . We consider the linear maps

$$\phi : X_i \rightarrow Y_{m-j}^* \quad \text{and} \quad \psi : Y_{m-j} \rightarrow X_i^*$$

defined by mapping x to $\phi(x)$, which sends $y \in Y_{m-j}$ to $\langle g_i^j(x), y \rangle$, and by mapping y to $\psi(y)$, which sends $x \in X_i$ to $\langle x, h_{m-j}^{m-i}(y) \rangle$. Since the pairings are compatible, we have

$$\phi(x)(y) = \langle g_i^j(x), y \rangle = \langle x, h_{m-j}^{m-i}(y) \rangle = \psi(y)(x).$$

Therefore, ϕ and ψ are adjoint, which means that the associated matrices are transposes of each other. It follows that ϕ and ψ have the same rank. Note that $\phi = \phi' \circ g_i^j$, where $\phi' : X_j \rightarrow Y_{m-j}^*$ is defined by mapping x' to $\phi'(x')$, which sends $y \in Y_{m-j}$ to $\langle x', y \rangle$. Since the pairing is non-degenerate, ϕ' is an isomorphism. Hence, $\text{rank } g_i^j = \text{rank } \phi$,

and by a similar argument, $\text{rank } h_{m-j}^{m-i} = \text{rank } \psi$. This now implies

$$\text{rank } g_i^j = \text{rank } h_{m-j}^{m-i}. \quad (19)$$

These ranks are the persistent Betti numbers, which characterize the births and deaths in the two persistence modules. Specifically, the number of classes born at X_i and dying entering X_j is

$$\text{rank } g_i^{j-1} - \text{rank } g_{i-1}^{j-1} - \text{rank } g_i^j + \text{rank } g_{i-1}^j.$$

Similarly, the number of classes born at Y_{m-j+1} and dying entering Y_{m-i+1} is

$$\text{rank } h_{m-j+1}^{m-i} - \text{rank } h_{m-j}^{m-i} - \text{rank } h_{m-j+1}^{m-i+1} + \text{rank } h_{m-j}^{m-i+1}.$$

But these two sums are equal by (19). \square

Alexander duality for functions. Let $f : \mathbb{S} \rightarrow [0, 1]$ be a perfect Morse function on the $(n + 1)$ -dimensional sphere, and $\mathbb{S} = \mathbb{U} \cup \mathbb{V}$ a complementary decomposition into two $(n + 1)$ -manifolds with shared boundary $\mathbb{M} = \mathbb{U} \cap \mathbb{V}$. Assuming $f|_{\mathbb{M}}$ is tame, we have only finitely many homological critical values, including 0 and 1. Letting these critical values be s_0 to s_m , we interleave them with homological regular values to get $s_0 < t_0 < s_1 < \dots < t_{m-1} < s_m$. We are interested in the filtrations defined by the restrictions of f to \mathbb{U} and to \mathbb{V} . To describe them, we write \tilde{U}_i^p for the p -th reduced homology group of \mathbb{U}_{t_i} , if $0 \leq i < m$, of \mathbb{U} , if $i = m$, and of $(\mathbb{U}, \mathbb{U}^{t_{2m-i}})$, if $m < i \leq 2m$. In the same way, we define \tilde{V}_i^q for $0 \leq i \leq 2m$ with respect to \mathbb{V} . Using Alexander duality to form isomorphisms, we get two contravariantly connected persistence modules:

$$\begin{array}{ccccccc} \tilde{U}_0^p & \rightarrow & \dots & \rightarrow & \tilde{U}_m^p & \rightarrow & \dots & \rightarrow & \tilde{U}_{2m}^p \\ \downarrow & & & & \downarrow & & & & \downarrow \\ \tilde{V}_{2m}^q & \leftarrow & \dots & \leftarrow & \tilde{V}_m^q & \leftarrow & \dots & \leftarrow & \tilde{V}_0^q, \end{array}$$

where $p + q = n$. At the end of this section, we will prove that there exists a sequences of compatible pairings between \tilde{U}_i^p and \tilde{V}_{2m-i}^q , for $0 \leq i \leq 2m$. Using these pairings, we get our first result, which we formulate using the superscript ‘ T ’ for the operation that reflects a dot across the vertical axis of a persistence diagram and, at the same time, changes its dimension from p to $q = n - p$.

LAND AND WATER THEOREM. *Let \mathbb{U} and \mathbb{V} be a complementary decomposition of $\mathbb{S} = \mathbb{S}^{n+1}$, and let $f : \mathbb{S} \rightarrow [0, 1]$ be a perfect Morse function whose restriction to the n -manifold $\mathbb{M} = \mathbb{U} \cap \mathbb{V}$ is tame. Then $\tilde{D}\text{gm}(f|_{\mathbb{V}}) = \tilde{D}\text{gm}(f|_{\mathbb{U}})^T$.*

PROOF. We apply the Contravariant PE Theorem to the persistence modules \tilde{U}_i^p and \tilde{V}_i^q : If a class α is born at \tilde{U}_i^p and dies entering \tilde{U}_j^p then its corresponding class is born at \tilde{V}_{2m-j+1}^q and dies entering \tilde{V}_{2m-i+1}^q . If $1 \leq i < j \leq m$, then the class α is represented by (s_i, s_j) in the p -th ordinary subdiagram of $f|_{\mathbb{U}}$. The corresponding class is represented by (s_j, s_i) in the q -th relative subdiagram of $f|_{\mathbb{V}}$. The reflection maps the first and second coordinate-axes of the ordinary subdiagram to the second and first coordinate-axes of the relative subdiagram. It follows that it maps (s_i, s_j) in the former to (s_j, s_i) in the latter subdiagram. Other cases are similar, and we conclude that $\tilde{D}\text{gm}_p(f|_{\mathbb{U}})$ and $\tilde{D}\text{gm}_q(f|_{\mathbb{V}})$ are reflections of each other. Writing this more succinctly gives the claimed relationship between the reduced persistence diagrams of f restricted to the two complementary subsets of the sphere. \square

Alexander pairing. We fill the gap in the proof of the Land and Water Theorem by establishing compatible pairings between the groups upstairs and downstairs. For the sake of simplicity, we restrict ourselves to the case $F = \mathbb{Z}_2$. While the argument for general fields is similar, it requires oriented simplices, which is a technical formalism we prefer to avoid. We begin with the pairing implicit in Lefschetz duality for manifolds with possibly non-empty boundary:

$$\langle \cdot, \cdot \rangle_L : \tilde{H}_p(\mathbb{U}_t) \times \tilde{H}_{q+1}(\mathbb{U}_t, \partial\mathbb{U}_t) \rightarrow \mathbb{Z}_2,$$

which is defined by mapping two classes to the parity of the number of intersections between representing cycles. We get such a pairing for every regular value t , and these pairings are compatible with the horizontal maps induced by inclusion of the sublevel sets; see [3]. Next, we reduce the dimension of the second factor using a mapping, φ , which we compose from four simpler mappings:

$$\begin{aligned} \tilde{H}_{q+1}(\mathbb{U}_t, \partial\mathbb{U}_t) &\xrightarrow{\varphi_1} \tilde{H}_{q+1}(\mathbb{S}, \mathbb{S}^t \cup \mathbb{V}_t) \\ &\xrightarrow{\varphi_2} \tilde{H}_q(\mathbb{S}^t \cup \mathbb{V}_t) \\ &\xrightarrow{\varphi_3} \tilde{H}_q(\mathbb{S}^t \cup \mathbb{V}_t, \mathbb{S}^t) \\ &\xrightarrow{\varphi_4} \tilde{H}_q(\mathbb{V}, \mathbb{V}^t). \end{aligned}$$

The first mapping, φ_1 , is an isomorphism defined by excision. Next, φ_2 , is the connecting homomorphism of the exact sequence of the pair $(\mathbb{S}, \mathbb{S}^t \cup \mathbb{V}_t)$. Since the q -th reduced homology of \mathbb{S} is trivial, for all $q \neq n+1$, φ_2 is an isomorphism for all $0 \leq q < n$. It is a surjection for $q = n$. The third mapping, φ_3 , is induced by inclusion. It occurs in the exact sequence of the pair $(\mathbb{S}^t \cup \mathbb{V}_t, \mathbb{S}^t)$. For $0 < t \leq 1$, the reduced homology groups of \mathbb{S}^t are all trivial, implying that φ_3 is an isomorphism. For $t = 0$, φ_3 is the trivial isomorphism. Finally, φ_4 is again an isomorphism defined by excision. In summary, we get $\varphi : \tilde{H}_{q+1}(\mathbb{U}_t, \partial\mathbb{U}_t) \rightarrow \tilde{H}_q(\mathbb{V}, \mathbb{V}^t)$, which is an isomorphism for $0 \leq q < n$ and a surjection for $q = n$. More specifically, each class $\beta \in \tilde{H}_n(\mathbb{V}, \mathbb{V}^t)$ has two preimages. Indeed, an n -cycle representing the corresponding class in $\tilde{H}_n(\mathbb{S}^t \cup \mathbb{V}_t)$ partitions the components of \mathbb{U}_t into two subsets, and each subset generates an $(n+1)$ -dimensional relative homology class that maps to β .

Note that the roles of \mathbb{U} and \mathbb{V} can be interchanged to get a mapping from $\tilde{H}_{q+1}(\mathbb{V}_t, \partial\mathbb{V}_t)$ to $\tilde{H}_q(\mathbb{U}, \mathbb{U}^t)$. With this, we are ready to construct the *Alexander pairing*:

$$\langle \cdot, \cdot \rangle_A : \tilde{H}_p(\mathbb{U}_t) \times \tilde{H}_q(\mathbb{V}, \mathbb{V}^t) \rightarrow \mathbb{Z}_2,$$

defined by $\langle \alpha, \beta \rangle_A = \langle \alpha, \beta' \rangle_L$, where β' is a preimage of β under φ . This fixes the pairing of \tilde{U}_i^p and \tilde{V}_{2m-i}^q for $0 \leq i < m$. Similarly, we define $\langle \cdot, \cdot \rangle_A$ for $\tilde{H}_p(\mathbb{U}, \mathbb{U}^t) \times \tilde{H}_q(\mathbb{V}_t)$ which fixes the pairing for $m < i \leq 2m$. The pairing is clearly well-defined for $0 < p < n$, where φ is an isomorphism. The remaining two cases are symmetric, and we consider $p = 0$ and $\beta \in \tilde{H}_n(\mathbb{V}, \mathbb{V}^t)$. As noted before, β has two preimages, β' and β'' , which are generated by complementary subsets of \mathbb{U}_t . Every $\alpha \in \tilde{H}_0(\mathbb{U}_t)$ is represented by an even number of points, which the complementary subsets partition into two even or two odd subsets. In either case, we have $\langle \alpha, \beta' \rangle_L = \langle \alpha, \beta'' \rangle_L$, so the pairing is well-defined in all cases. Moreover, the compatibility of the Alexander pairings within the left and right halves follows from the compatibility of the Lefschetz pairings.

It remains to define the pairing in the middle of the sequence, for $\tilde{U}_m^p = \tilde{H}_p(\mathbb{U})$ and $\tilde{V}_m^q = \tilde{H}_q(\mathbb{V})$. To this end, let

α be a reduced p -cycle in \mathbb{U} , β a reduced q -cycle in \mathbb{V} , and β' a $(q+1)$ -chain whose boundary is β . We define $\langle \alpha, \beta \rangle_A$ by counting the intersections between α and β' . Note that β' is the preimage of β under the previously defined mapping, $\varphi : \tilde{H}_{q+1}(\mathbb{U}_t, \partial\mathbb{U}_t) \rightarrow \tilde{H}_q(\mathbb{V}, \mathbb{V}^t)$, for any regular value t . Therefore, the pairing in the middle is compatible with the pairings in left half, again exploiting the compatibility of the Lefschetz pairing. Alternatively, we can define the pairing by taking a $(p+1)$ -chain α' with boundary α and counting its intersections with β . In this case, α' is the preimage of α under the map $\varphi : \tilde{H}_{p+1}(\mathbb{V}_t, \partial\mathbb{V}_t) \rightarrow \tilde{H}_p(\mathbb{U}, \mathbb{U}^t)$, and compatibility with the pairings in the right half follows. Indeed, both definitions are equivalent, as already observed by Lefschetz [9]. We give a simple proof for the case $F = \mathbb{Z}_2$:

BRIDGE LEMMA. *Let α and β be non-intersecting reduced cycles on the $(n+1)$ -sphere whose dimensions add up to n . Let α' and β' be chains whose boundaries are α and β , respectively. Then $\langle \alpha', \beta \rangle_L = \langle \alpha, \beta' \rangle_L$.*

PROOF. We can assume that α and β do not intersect. The intersection of α' and β' is a 1-chain. Its boundary consists of an even number of points, and is the disjoint union of intersections of α' with β and of α with β' . This implies that both types of intersections occur with the same parity. \square

5. SHORE

This section presents our second result, which extends the combination of Alexander duality and Mayer-Vietoris sequences from spaces to functions.

Mayer-Vietoris sequence of filtrations. Assuming that $f|_{\mathbb{M}}$ is tame, we write $s_0 < t_0 < s_1 < \dots < t_{m-1} < s_m$ for the interleaved sequence of homological critical and regular values, as before. The main tool in this section is the diagram obtained by connecting the filtrations of f and its restrictions with Mayer-Vietoris sequences. We describe this using shorthand notation for the groups. Consistent with earlier notation, we write S_i^p for the p -th homology group of \mathbb{S}_{t_i} , for $0 \leq i < m$, of \mathbb{S} , for $i = m$, and of $(\mathbb{S}, \mathbb{S}^{t_{2m-i}})$, for $m < i \leq 2m$. Similarly, we write M_i^p for the p -th homology group of \mathbb{M}_{t_i} , of \mathbb{M} , and of $(\mathbb{M}, \mathbb{M}^{t_{2m-i}})$. Finally, we write D_i^p for the direct sum of the p -th homology groups of \mathbb{U}_{t_i} and \mathbb{V}_{t_i} , of \mathbb{U} and \mathbb{V} , and of $(\mathbb{U}, \mathbb{U}^{t_{2m-i}})$ and $(\mathbb{V}, \mathbb{V}^{t_{2m-i}})$. Drawing the filtrations from left to right and the Mayer-Vietoris sequences from top to bottom, we get

$$\begin{array}{ccccccc} \downarrow & & \downarrow & & \downarrow & & \downarrow \\ S_0^{p+1} & \rightarrow & \dots & \rightarrow & S_m^{p+1} & \rightarrow & \dots & \rightarrow & S_{2m}^{p+1} \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ M_0^p & \rightarrow & \dots & \rightarrow & M_m^p & \rightarrow & \dots & \rightarrow & M_{2m}^p \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ D_0^p & \rightarrow & \dots & \rightarrow & D_m^p & \rightarrow & \dots & \rightarrow & D_{2m}^p \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ S_0^p & \rightarrow & \dots & \rightarrow & S_m^p & \rightarrow & \dots & \rightarrow & S_{2m}^p \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \end{array}$$

All squares commute, which is particularly easy to see for the squares that connect groups of the same dimension, whose maps are all induced by inclusion. For $1 \leq p \leq n$, all groups S_i^p are trivial. By exactness of the Mayer-Vietoris sequences, this implies that the maps $m_i^p : M_i^p \rightarrow D_i^p$ are isomorphisms, for $1 \leq p \leq n-1$. The persistence diagram of $f|_{\mathbb{M}}$ is therefore the disjoint union of the persistence diagrams of $f|_{\mathbb{U}}$ and $f|_{\mathbb{V}}$. More generally, we claim:

GENERAL SHORE THEOREM. *Let n be a positive integer, let \mathbb{U} and \mathbb{V} be a complementary decomposition of $\mathbb{S} = \mathbb{S}^{n+1}$, and let $f : \mathbb{S} \rightarrow [0, 1]$ be a perfect Morse function whose restriction to the n -manifold $\mathbb{M} = \mathbb{U} \cap \mathbb{V}$ is tame. Then*

$$\text{Dgm}_0(f|_{\mathbb{M}}) = [\text{Dgm}_0(f|_{\mathbb{U}}) \sqcup \text{Dgm}_0(f|_{\mathbb{V}})]^C, \quad (20)$$

$$\text{Dgm}_p(f|_{\mathbb{M}}) = \text{Dgm}_p(f|_{\mathbb{U}}) \sqcup \text{Dgm}_p(f|_{\mathbb{V}}), \quad (21)$$

$$\text{Dgm}_n(f|_{\mathbb{M}}) = [\text{Dgm}_0(f|_{\mathbb{U}}) \sqcup \text{Dgm}_0(f|_{\mathbb{V}})]^{CT}, \quad (22)$$

for $1 \leq p \leq n-1$, where C stands for applying the cascade and T stands for reflecting the diagram.

We note that the assumption of n being positive is necessary since the formulas do not hold for $n = 0$. Concerning the proof of the theorem, note that (21) is clear, and that (22) follows from (20) using the duality of persistence diagrams from [3]. We will need to study the impact of the non-trivial groups \mathbb{S}_i^0 to prove (20).

Latitudinal manifolds. Call the minimum of f the *south-pole* and the maximum the *north-pole* of the sphere. Let \mathbb{M}' be a component of \mathbb{M} , and note that it is an n -manifold that decomposes \mathbb{S} into two complementary subsets. If it separates the two poles, we refer to \mathbb{M}' as a *latitudinal n -manifold*. Assuming neither pole lies on \mathbb{M} , we order the latitudinal n -manifolds from south to north as $\mathbb{M}_1, \mathbb{M}_2, \dots, \mathbb{M}_\ell$. Letting u_k and w_k be the minimum and maximum values of f restricted to \mathbb{M}_k , we get $u_1 < u_2 < \dots < u_\ell$ as well as $w_1 < w_2 < \dots < w_\ell$.

For each component \mathbb{M}' of \mathbb{M} , there are *neighboring components* \mathbb{U}' of \mathbb{U} and \mathbb{V}' of \mathbb{V} defined such that $\mathbb{U}' \cap \mathbb{V}' = \mathbb{M}'$. A single component of \mathbb{U} or \mathbb{V} can be neighbor to an arbitrary number of n -manifolds, but not to more than two latitudinal n -manifolds. Specifically, there are components $\mathbb{S}_0, \mathbb{S}_1, \dots, \mathbb{S}_\ell$ of \mathbb{U} and \mathbb{V} such that $\mathbb{M}_k = \mathbb{S}_{k-1} \cap \mathbb{S}_k$ for each $1 \leq k \leq \ell$. We refer to the \mathbb{S}_k as *latitudinal components*. For example, in Figure 5, we have $\ell = 2$ latitudinal 1-manifolds and $\ell + 1 = 3$ latitudinal components. Setting $u_0 = 0$ and $w_{\ell+1} = 1$, the minimum and maximum values of f restricted to \mathbb{S}_k are u_k and w_{k+1} , for $0 \leq k \leq \ell$. Assuming \mathbb{S}_k belongs to \mathbb{U} , it gives rise to zero or more dots in the ordinary subdiagram of $\text{Dgm}_0(f|_{\mathbb{U}})$ and to exactly one dot, (u_k, w_{k+1}) , in the horizontal subdiagram. We say the dot in the horizontal subdiagram *represents* the 0-dimensional homology class defined by \mathbb{S}_k . Note that its coordinates are indexed consistently with the notation used in the introduction of the cascade. There is indeed a connection, namely the dots in the 0-th horizontal subdiagram representing latitudinal components are the extreme ones in the multiset:

LATITUDINAL COMPONENT LEMMA. *The latitudinal components of \mathbb{U} and \mathbb{V} correspond bijectively to the extreme dots in the disjoint union of $\text{Dgm}_0(f|_{\mathbb{U}})$ and $\text{Dgm}_0(f|_{\mathbb{V}})$.*

We omit the proof, which is not difficult. The statement includes the case in which there is no latitudinal n -manifold so that \mathbb{S}_0 is the only latitudinal component. It contains both poles and is therefore represented by $(0, 1)$, which is the only extreme dot in the disjoint union of 0-th diagrams.

Proof of (20). We compare $\text{Dgm}_0(f|_{\mathbb{M}})$ with the disjoint union of $\text{Dgm}_0(f|_{\mathbb{U}})$ and $\text{Dgm}_0(f|_{\mathbb{V}})$, noting that their dots all belong to the ordinary and horizontal subdiagrams. Consider first a dot (a, b) in the ordinary subdiagram $\text{Ord}_0(f|_{\mathbb{M}})$. It represents a component in the sublevel set that is born at a minimum $x \in \mathbb{M}$, with $f(x) = a$, and that dies at a

saddle point $y \in \mathbb{M}$, with $f(y) = b$. Let $\mathbb{M}' \subseteq \mathbb{M}$ be the connected n -manifold that contains x and y , and recall that $\mathbb{M}' = \mathbb{U}' \cap \mathbb{V}'$. Assume that x is a local minimum of $f|_{\mathbb{U}}$. In the sequence of sublevel sets of $f|_{\mathbb{U}}$, we see the birth of a component at $f(x) = a$ and its death at $f(y) = b$. It follows that (a, b) is also a dot in $\text{Ord}_0(f|_{\mathbb{U}})$ and therefore of the disjoint union of the diagrams of $f|_{\mathbb{U}}$ and $f|_{\mathbb{V}}$. The argument can be reversed, which implies that the 0-th ordinary subdiagrams are the same.

Consider second a dot (a, b) in the horizontal subdiagram $\text{Hor}_0(f|_{\mathbb{M}})$. It represents a connected n -manifold $\mathbb{M}' \subseteq \mathbb{M}$, which splits \mathbb{S} into two subsets. If \mathbb{M}' is non-latitudinal, then one subset contains both poles and the other contains neither. The latter subset contains a (non-latitudinal) neighboring component, which is represented by (a, b) in $\text{Hor}_0(f|_{\mathbb{U}})$ or in $\text{Hor}_0(f|_{\mathbb{V}})$. Again, the argument can be reversed. If on the other hand, \mathbb{M}' is latitudinal, then $(a, b) = (u_i, w_i)$, for some i , where we write u_1, u_2, \dots, u_ℓ for the minimum values and w_1, w_2, \dots, w_ℓ for the maximum values of the latitudinal n -manifolds, as before. More generally, we get the dots (u_i, w_i) in $\text{Hor}_0(f|_{\mathbb{M}})$, for $1 \leq i \leq \ell$. Adding $u_0 = 0$ and $w_{\ell+1} = 1$, we get the dots (u_i, w_{i+1}) in $\text{Hor}_0(f|_{\mathbb{U}}) \sqcup \text{Hor}_0(f|_{\mathbb{V}})$, for $1 \leq i \leq \ell$. After adding $(0, 0)$ and $(1, 1)$, which are both irrelevant, we get precisely the dots specified in (16) and in (17). This implies that the two diagrams are related to each other by a cascade, which completes the proof of (20).

The General Shore Theorem can also be proven with algebraic arguments: Di Fabio and Landi [4, 5] consider decompositions $\mathbb{X} = \mathbb{A} \cup \mathbb{B}$, for a general topological space \mathbb{X} , and relate the ordinary, relative, and extended (horizontal plus vertical) persistent Betti numbers of a function restricted to these sets with each other. These relationships depend on the ranks of maps between certain absolute and relative homology groups; see Theorems 3.1, 3.3, and 3.4 in [5] for details. Applied to the case of a sphere, these maps are trivial except for the extended case in dimensions 0 and n . Investigating this case more carefully leads to the same result as (20).




Figure 5: Left: the height function on the 2-sphere decomposed into an annulus and a pair of disks. Right: the reduced persistence diagram of f restricted to the annulus (shaded circular dots), to the pair of disks (white circular dots), and to the two circles (squares).

Euclidean case. Considering that (13) to (15) give elegant relations between the Betti numbers of the shore and the land, it is perhaps surprising that we need cascades to formulate a similar result for persistence diagrams. Indeed, (13) to (15) suggest that the persistence diagrams of $f|_{\mathbb{M}}$ be the disjoint union of the persistence diagram of $f|_{\mathbb{U}}$ and

of its reflection. The example in Figure 5 shows that this simple relation does not hold in general. Indeed, we have the dots (a, c) , (b, d) , (c, a) , (d, b) in the standard diagram of $f|_{\mathbb{M}}$, while the standard diagram of $f|_{\mathbb{U}}$ contains (a, d) and (c, b) . This clearly violates the suggested relation. However, there is a natural setting in which the relation does hold. To state this result, we let \mathbb{A} be a compact $(n + 1)$ -manifold with boundary in \mathbb{R}^{n+1} . To be compatible with the rest of this paper, we assume \mathbb{A} is triangulable and there exists an open neighborhood and a deformation retraction of that neighborhood to \mathbb{A} .

EUCLIDEAN SHORE THEOREM. *Let n be a positive integer, let $\mathbb{A} \subset \mathbb{R}^{n+1}$ be a compact $(n + 1)$ -manifold with boundary satisfying the technical conditions as described, and suppose that $e : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ has no homological critical values and its restriction to $\partial\mathbb{A}$ is tame. Then*

$$\text{Dgm}(e|_{\partial\mathbb{A}}) = \text{Dgm}(e|_{\mathbb{A}}) \sqcup \text{Dgm}(e|_{\mathbb{A}})^T.$$

PROOF. We can extend e to a perfect Morse function $f : \mathbb{S}^{n+1} \rightarrow \mathbb{R}$ and \mathbb{A} to a subset \mathbb{U} of $\mathbb{S} = \mathbb{S}^{n+1}$ such that the persistence diagrams of e restricted to \mathbb{A} and to $\partial\mathbb{A}$ are the same as those of f restricted to \mathbb{U} and $\mathbb{M} = \partial\mathbb{U}$. It thus suffices to show that the persistence diagram of $f|_{\mathbb{M}}$ is the disjoint union of the diagram of $f|_{\mathbb{U}}$ and of its reflection. For $0 < p < n$, this follows from (21), from

$$\tilde{\text{Dgm}}_p(f|_{\mathbb{V}}) = \tilde{\text{Dgm}}_{n-p}(f|_{\mathbb{U}})^T,$$

as stated in the Land and Water Theorem, and the fact that the reduced diagrams are equal to the non-reduced ones. For $p = 0$, we start with (20) and note that there is only one extreme dot in $\text{Dgm}_0(f|_{\mathbb{U}}) \sqcup \text{Dgm}_0(f|_{\mathbb{V}})$, namely the one representing the component of \mathbb{V} that contains both the minimum and the maximum of f . It follows that the cascade leaves $\text{Dgm}_0(f|_{\mathbb{U}})$ unchanged while it turns $\text{Dgm}_0(f|_{\mathbb{V}})$ into $\tilde{\text{Dgm}}_0(f|_{\mathbb{V}})$. Using the Land and Water Theorem again, we obtain:

$$\begin{aligned} \text{Dgm}_0(f|_{\mathbb{M}}) &= [\text{Dgm}_0(f|_{\mathbb{U}}) \sqcup \text{Dgm}_0(f|_{\mathbb{V}})]^C \\ &= \text{Dgm}_0(f|_{\mathbb{U}}) \sqcup \tilde{\text{Dgm}}_0(f|_{\mathbb{V}}) \\ &= \text{Dgm}_0(f|_{\mathbb{U}}) \sqcup \tilde{\text{Dgm}}_n(f|_{\mathbb{U}})^T \\ &= \text{Dgm}_0(f|_{\mathbb{U}}) \sqcup \text{Dgm}_n(f|_{\mathbb{U}})^T. \end{aligned}$$

For $p = n$, we exploit Poincaré duality for manifolds, which implies $\text{Dgm}_n(f|_{\mathbb{M}}) = \text{Dgm}_0(f|_{\mathbb{M}})^T$; see [3]. \square

The 4-dimensional version of the theorem, for $n = 3$, brings us back full circle to the motivation for this work, namely the computation of the persistence diagram of the space-time shape formed by a moving collection of biological cells [7]. Modeling space-time as \mathbb{R}^4 , the shape is compact, and we consider the time function restricted to that shape. This data satisfies the assumptions of the Euclidean Shore Theorem, so we can infer the persistence diagram of the function on the boundary from the diagram of the function on the solid 4-dimensional shape.

6. DISCUSSION

The main contributions of this paper are two extensions of Alexander duality from spaces to functions. The first extension is direct and relates the persistence diagrams of a perfect Morse function restricted to two $(n + 1)$ -manifolds

with shared boundary, forming a complementary decomposition of the $(n + 1)$ -dimensional sphere with each other. The second extension relates the persistence diagram of the function restricted to the n -manifold boundary with the diagram of the function restricted to one $(n + 1)$ -manifold. A key tool in its proof is the persistence module of Mayer-Vietoris sequences (or the Mayer-Vietoris sequence of persistence modules). This suggests the study of more general modules of exact sequences. Besides the hope to develop a general purpose device that pervades persistent homology in the same way exact sequences pervade homology, we motivate the study with a few concrete questions.

- Can our results be extended to the more general setting of a decomposition $\mathbb{X} = \mathbb{A} \cup \mathbb{B}$ of a topological space \mathbb{X} , as considered by Di Fabio and Landi [4, 5].
- As pointed out to us by Yuriy Mileyko, an alternative approach to generalizing our results for functions can be based on the generalization of Alexander duality given in [8, Proposition 3.46]. Combining this statement of duality with the Mayer-Vietoris sequence has the potential to generalize our two Shore Theorems.
- Recall that our Euclidean Shore Theorem relates the persistence diagram of a function f on an $(n + 1)$ -manifold with boundary $\mathbb{A} \subseteq \mathbb{R}^{n+1}$ with that of its restriction to the boundary. Here, \mathbb{A} is assumed to be compact. As suggested to us by Amit Patel, we can consider other inclusions $\mathbb{M} \subseteq \mathbb{A}$ and ask how much information about the diagram of f restricted to \mathbb{M} we can infer from the diagram of f restricted to \mathbb{A} . Similarly, how much can we say about \mathbb{M} by comparing the persistence diagrams of f restricted to the two spaces?

Acknowledgments

The authors thank an anonymous referee for suggesting the simplified proof of the Contravariant PE Theorem given in this paper. They also thank Frederick Cohen, Yuriy Mileyko and Amit Patel for helpful discussions. This research is partially supported by the National Science Foundation (NSF) under grant DBI-0820624, the European Science Foundation under the Research Networking Programme, and the Russian Government Project 11.G34.31.0053.

7. REFERENCES

- [1] P. BENDICH, S. CABELLO AND H. EDELSBRUNNER. A point calculus for interlevel set homology. *Pattern Rec. Lett.*, to appear.
- [2] G. CARLSSON, V. DE SILVA AND D. MOROZOV. Zigzag persistent homology and real-valued functions. In “Proc. 25th Ann. Sympos. Comput. Geom., 2009”, 247–256.
- [3] D. COHEN-STEINER, H. EDELSBRUNNER AND J. HARER. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.* **9** (2009), 79–103.
- [4] B. DI FABIO AND C. LANDI. A Mayer-Vietoris formula for persistent homology with an application to shape recognition in the presence of occlusions. *Found. Comput. Math.* **11** (2011), 499–527.
- [5] B. DI FABIO AND C. LANDI. Persistent homology and partial similarity of shapes. *Pattern Recognition Letters*, to appear.

- [6] H. EDELSBRUNNER AND J. L. HARER. *Computational Topology. An Introduction*. Amer. Math. Soc., Providence, Rhode Island, 2010.
- [7] H. EDELSBRUNNER, C.-P. HEISENBERG, M. KERBER AND G. KRENS. The medusa of spatial sorting. Manuscript, IST Austria, Klosterneuburg, Austria, 2012.
- [8] A. HATCHER. *Algebraic Topology*. Cambridge Univ. Press, Cambridge, England, 2002.
- [9] S. LEFSCHETZ. Intersections and transformations of complexes and manifolds. *Amer. Math. Soc.* **28** (1926), 1–49.
- [10] J. R. MUNKRES. *Elements of Algebraic Topology*. Perseus, Cambridge, Massachusetts, 1984.