Homology, Homotopy and Applications, vol. 23(2), 2021, pp.121–140

MAGNITUDE HOMOLOGY, DIAGONALITY, AND MEDIAN SPACES

RÉMI BOTTINELLI AND TOM KAISER

(*communicated by J. Daniel Christensen*)

Abstract

We verify that the Künneth and Mayer–Vietoris formulae for magnitude homology of graphs, proven by Hepworth and Willerton, generalise naturally to the metric setting. Similarly, we extend the notion of diagonality of graphs to metric spaces, and verify its stability under products, retracts, and filtrations. As an application, we show that median spaces are diagonal; in particular any Menger convex median space has vanishing magnitude homology.

1. Introduction

These notes have two somewhat independent purposes, both related to the notion of *magnitude homology*—a categorification of *magnitude*—in short defined as follows. Let X be a metric space. Given two points x, y of X, the interval $[x, y] \subseteq X$ is the set of points $z \in X$ satisfying $d(x, z) + d(z, y) = d(x, y)$. The *magnitude homology* of X is the homology of its *magnitude complex*, which, at height k, is spanned freely by $(k + 1)$-tuples of consecutively distinct points in X. Its boundary operator is the alternating sum of the maps ∂_i, each defined as sending the tuple $\langle x_0, \dots, x_i, \dots, x_k\rangle$ to $\langle x_0, \dots, \hat{x}_i, \dots, x_k\rangle$ in case $x_i \in [x_{i-1}, x_{i+1}]$, and to 0 otherwise.

Defining the *length* $l(\mathbf{x})$ of a $(k + 1)$-tuple $\mathbf{x} := \langle x_0, \dots, x_k\rangle$ as the sum $l(\mathbf{x}) := \sum_{i=0}^{k-1} d(x_i, x_{i+1})$, we see that the boundary operator preserves length. It follows that the magnitude homology groups are graded (in $\mathbb{R}_{\geqslant 0}$) by the length.

1.1. Magnitude homology of median spaces

Our first aim is to analyse the magnitude homology of median metric spaces. Recall that the (metric) space X is called *median* if, for any three distinct points $x, y, z \in X$, the intersection $[x, y] \cap [y, z] \cap [z, x]$ is a singleton. Median spaces are relatively common: examples include trees (more generally $\mathbb{R}$-trees), any product of median spaces with the l^1 metric, and skeleta of CAT(0) cube complexes. For more examples and information on median spaces, see [5]. We call a $(k + 1)$-tuple $\langle x_0, \dots, x_k\rangle$ *saturated* if each interval $[x_i, x_{i+1}]$ contains only x_i and x_{i+1}, and, extending the notion of

The first author was supported by the Swiss National Science Foundation project no. PP00P2-144681/1.
Received April 27, 2020, revised May 26, 2020, August 16, 2020; published on April 21, 2021.
2010 Mathematics Subject Classification: 55N35.
Key words and phrases: magnitude, metric space.
Article available at http://dx.doi.org/10.4310/HHA.2021.v23.n2.a7

diagonality defined in [8], say that X is *diagonal* if all its magnitude homology groups are spanned by linear combinations of saturated tuples.

We obtain the following:

Proposition 6.3. *Median metric spaces are diagonal.*

In Section 3, we verify that median *graphs* are diagonal. The key tool here is the characterisation of median graphs as retracts of hypercubes, due to Bandelt [2]. Then, in Section 5, we formally define diagonality and verify that it is stable under some constructions on metric spaces (e.g. retracts, products, and filtrations). Finally, Section 6 is devoted to proving diagonality of median spaces, by what amounts to an approximation of a median space by finite median graphs. The argument uses the stability of diagonality mentioned above, along with an equivalence between finite median spaces and finite median graphs, due to Avann [1].

1.2. Künneth and Mayer–Vietoris formulae

Section 4 is dedicated to the second part of these notes: verifying that the Künneth and Mayer–Vietoris formulae for graph magnitude homology, proven in [8], generalise to the metric setting.

Recall that the l^1 *product* of two metric spaces X, Y has the Cartesian product $X \times Y$ as underlying set, and its distance map is given by

$$d_{X\times Y}((x,y),(x',y')) := d_X(x,x') + d_Y(y,y').$$

The "metrised" variant of the Künneth theorem in [8] has the form:

Proposition 4.3 (Künneth theorem—metric case). *If X, Y are metric spaces and $X \times Y$ is their l^1 product, then there exists a natural "cross-product" morphism*

$$\mathrm{MH}_*(X) \otimes \mathrm{MH}_*(Y) \xrightarrow{\square} \mathrm{MH}_*(X \times Y)$$
$$[f] \otimes [g] \mapsto [f \square g],$$

which fits into a natural short exact sequence

$$0 \longrightarrow \mathrm{MH}_*(X) \otimes \mathrm{MH}_*(Y) \xrightarrow{\square} \mathrm{MH}_*(X \times Y) \longrightarrow \mathrm{Tor}(\mathrm{MH}_{*-1}(X), \mathrm{MH}_*(Y)) \longrightarrow 0.$$

The proof found in [8] translates rather easily to the metric setting, and is done in Section 4.1.

The Mayer–Vietoris formula in [8] uses so-called *projecting decompositions* of graphs. We will use the following metric analogue of a projecting decomposition. A *gated decomposition* of the metric space X is a pair of subspaces Y, Z satisfying $X = Y \cup Z$ and such that:

- $Y \cap Z$ is *gated* relative to Z: given any $z \in Z$, there is a unique element $y_z \in Y \cap Z$, called a *gate*, such that for any $y \in Y \cap Z$, y_z is in $[z, y]$.
- $Y \cap Z$ lies "between" Y and Z: for any $z \in Z$ and $y \in Y$, there exists some $w \in Y \cap Z$ such that w is in $[y, z]$.

The notion of gated subsets in metric spaces is fairly well-studied; see e.g. [7].

Our "metrised" variant of the Mayer–Vietoris theorem of [8] has the form:

Theorem 4.14 (Mayer–Vietoris—metric case)**.** *If $X = Y \cup Z$ is a gated decomposition of X and $W = Y \cap Z$, then the inclusions*

$$j_Y : W \to Y, \quad j_Z : W \to Z, \quad i_Y : Y \to X, \quad i_Z : Z \to X$$

induce a short exact sequence

$$0 \longrightarrow \mathrm{MH}_*(W) \xrightarrow{\langle (j_Y)_*, -(j_Z)_* \rangle} \mathrm{MH}_*(Y) \oplus \mathrm{MH}_*(Z) \xrightarrow{(i_Y)_* \oplus (i_Z)_*} \mathrm{MH}_*(X) \longrightarrow 0.$$

Its proof requires a bit more care than the Künneth theorem, since the discreteness of graphs allows for some simplifying assumptions. Unlike in [**8**], no naturality property for the Mayer–Vietoris formula is given here, though we do expect a naturality statement similar to the one in [**8**] to hold in the metric case, assuming the right setup.

1.3. A remark on betweenness

In a metric space X, we say that a point z is *between* two given points x and y if $z \in [x, y]$, that is, if $d(x, y) = d(x, z) + d(z, y)$. When one discards the grading, magnitude homology groups only depend on the "betweenness" relation. Thus, many arguments can be worked out without appealing to either notion of length or distance, instead relying only on betweenness. We strove to make this reliance as apparent as possible, while de-emphasising the length grading. The advantage of this approach appears, for instance, in the proof of Proposition 6.3.

Acknowledgments

We thank Victor Chepoi for sharing Proposition 6.1 with us, and a referee for their thorough and insightful report.

2. Background

In this section, we provide necessary definitions and settle on notation.

2.1. Metric spaces

Let (X, d) be a metric space. Finite sequences of points in X are written using angle brackets: $\mathbf{x} = \langle x_0, \ldots, x_k \rangle$ and identified with maps $[k] \to X$, where $[k] := \{0, \ldots, k\}$. If $x_i \neq x_{i+1}$ for all $0 \leqslant i < k$, we call such a sequence a *k-path*; the set of k-paths in X is denoted by $P_k(X)$, and the set of all paths by $P(X)$. The *length* $l(\mathbf{x})$ of a k-path $\mathbf{x} = \langle x_0, \ldots x_k \rangle$ is defined as the sum

$$l(\mathbf{x}) := \sum_{i=0}^{k-1} d(x_i, x_{i+1}).$$

Given two points $x, y \in X$, we say that a third point $z \in X$ lies *between* them if

$$d(x, y) = d(x, z) + d(z, y).$$

In other words, z turns the triangle inequality into an equality. If furthermore $z \neq x$ and $z \neq y$, we say that z lies *strictly between* x and y. We write $[x, y]$ for the set of points between x and y and $]x, y[$ for those strictly between. We call $[x, y]$ and $]x, y[$

intervals for obvious reasons. A k-path $\mathbf{x}$ is *saturated* if each strict interval $]x_i, x_{i+1}[$ is empty. A metric space is *Menger convex* if no strict interval between distinct points is empty.

A map $f\colon X \to Y$ between metric spaces is *non-expanding* (or 1-*Lipschitz*) if for all $x, x' \in X$, we have

$$d(fx, fx') \leqslant d(x, x').$$

We let Met denote the category with objects metric spaces and morphisms given by 1-Lipschitz maps.

A subset A of a metric space X is *convex* if for all $a, b \in A$, the interval $[a, b]$ in X is contained in A; in other words, any point between two points of A is also in A. Note that this definition is stronger than the one found in [**8**] for graphs.

If X is a set and A is a directed set, a *filtration* of X is a family $(U_\alpha)_{\alpha\in A}$ of subsets of X satisfying $\bigcup_\alpha U_\alpha = X$ and such that for any $\alpha \leqslant \beta$, we have $U_\alpha \subseteq U_\beta$.

If Y is a subspace of X, a *retraction* of X onto Y is a 1-Lipschitz map $f\colon X \to Y$ satisfying $f|Y = \mathrm{Id}_Y$.

Finally, if X and Y are metric spaces, we will always endow the set $X \times Y$ with the l^1 metric, that is:

$$d_{X\times Y}((x, y), (x', y')) := d_X(x, x') + d_Y(y, y').$$

Note that the l^1 product of metric space is *not* the categorical product in Met. Instead, under the interpretation of metric spaces as categories enriched over $[0, \infty]$ (so-called *Lawvere metric spaces*), the l^1 product becomes the *tensor* product in that category (see [**12**, Section 1.4]).

2.2. Graphs

Recall that a *graph* G is defined as a pair (V, E), where V is an arbitrary set and E is an antireflexive, symmetric binary relation on V (in other words, a set of ordered pairs of elements of V such that $(u, v) \in E$ implies both $(v, u) \in E$ and $u \neq v$). An element of V is called a *vertex*, and an element of E an *edge*. The *induced distance* on G, $d_G\colon V \times V \to \mathbb{R}_{\geqslant 0} \cup \{\infty\}$, assigns to a pair (u, v) of vertices the least number of edges needed to connect u to v, that is, the least k such that there exist vertices $w_0, \ldots, w_k$ satisfying $w_0 = u, w_k = v$, and $(w_i, w_{i+1}) \in E$ for each $i \in \{0, \ldots, k-1\}$. If no such k exists for some pair (u, v), their distance is set to ∞ and the graph is said to be *disconnected*. For simplicity's sake, we will assume all graphs to be connected. This assumption makes the pair (V, d_G) into a metric space; note that d_G takes values in $\mathbb{N}$, thus has discrete range.

In the case of graphs, the condition for 1-Lipschitz maps can be restated in terms of edges as follows: If $G = (V, E)$ and $G' = (V', E')$ are graphs, then $f\colon V \to V'$ is 1-Lipschitz if and only if, for all $u, v \in V$, $(u, v) \in E$ implies $(fu, fv) \in E'$ or $fu = fv$.

From the data (V, d_G), the set E of edges of G can easily be recovered: it is the set of pairs (u, v) satisfying $d_G(u, v) = 1$. With this fact in mind, we will henceforth view graphs as special cases of metric spaces (given by pairs (V, d_G)) rather than combinatorial objects (given by pairs (V, E)).

A subset U of vertices of a graph $G = (V, E)$ can be understood to define a metric space in two ways. On the one hand, one can look at the graph $G(U) := (U, E \cap U^2)$, and consider the induced distance $d_{G(U)}$; on the other hand, one can simply restrict

d_G to $U \times U$. According to our interpretation of graphs as "special" metric spaces, we will only use the second definition.

In the case of graphs, it is easily seen that the graph-theoretic Cartesian product agrees with our definition of the l^1 product of metric spaces. Hence, from now on, we will use the term "Cartesian product" with this definition in mind, and not that of a *categorical* direct product.

Since our definitions differ from the ones in [**8**], some care has to be taken in the "translation" process. The results stated in [**8**] tend to assume a more combinatorial setting, while the generalisations we present here are stated from a metric viewpoint.

2.3. Magnitude homology

Magnitude homology is a $\mathbb{R}_{\geqslant 0}$-graded homology introduced in [**8**] for graphs, and in [**13**] for arbitrary metric spaces, among other structures. It is a categorification of magnitude, itself an invariant (in the form of a power series) of finite metric spaces introduced in [**12**]. More precisely, in the case of a finite graph X, the magnitude $\#X$ of X is recovered from its magnitude homology $(\mathrm{MH}_k^l(X))_{k\in\mathbb{N}}^{l\geqslant 0}$ via the formula (borrowed from [**8**, p. 32]):

$$\#X(q) = \sum_{k,l} (-1)^k \operatorname{rank}(\mathrm{MH}_k^l(X)) \cdot q^l,$$

that is, as a "weighted" Euler characteristic. We will focus strictly on magnitude *homology*, but the interested reader can find an up-to-date bibliography of magnitude and magnitude homology maintained by T. Leinster in [**11**].

Let us recall its definition, in the metric case:

Definition 2.1 (Magnitude complex)**.** Let X be a metric space. The *magnitude complex* of X is the chain complex $(\mathrm{MC}_*(X), \partial_*)$ defined by:

$$\mathrm{MC}_k(X) := \mathbb{Z}[P_k(X)],$$

(the free abelian group on the set of all k-paths in X) with boundary map

$$\partial_k := \sum_{i=1}^{k-1} (-1)^i \partial_{k,i} \colon \mathrm{MC}_k(X) \to \mathrm{MC}_{k-1}(X),$$

where $\partial_{k,i} \colon \mathrm{MC}_k(X) \to \mathrm{MC}_{k-1}(X)$ is defined by

$$\partial_{k,i}\langle x_0, \dots, x_i, \dots, x_k\rangle := \begin{cases} \langle x_0, \dots, \widehat{x_i}, \dots, x_k\rangle & \text{if } x_i \text{ is between } x_{i-1} \text{ and } x_{i+1}; \\ 0 & \text{otherwise.} \end{cases}$$

We will write $\mathrm{MZ}_k(X)$ for $\ker \partial_k$ and $\mathrm{MB}_k(X)$ for $\operatorname{im} \partial_{k+1}$ respectively, so that the k-th *magnitude homology* group of X is

$$\mathrm{MH}_k(X) := \mathrm{MZ}_k(X)/\mathrm{MB}_k(X).$$

The magnitude complex enjoys a grading on $\mathbb{R}_{\geqslant 0}$ by letting $\mathrm{MC}_k^l(X)$ be the subgroup of $\mathrm{MC}_k(X)$ spanned by the k-paths of length l.

A 1-Lipschitz map $f \colon X \to Y$ induces a morphism of magnitude complexes by letting:

$$\mathrm{MC}_k(f)\langle x_0,\dots,x_k\rangle := \begin{cases}\langle fx_0,\dots,fx_k\rangle & \text{if } l\langle fx_0,\dots,fx_k\rangle = l\langle x_0,\dots,x_k\rangle;\\ 0 & \text{otherwise.}\end{cases}$$

Thus, $\mathrm{MC}_*(\bullet)$ defines a functor from the category Met of metric spaces and 1-Lipschitz maps to the category of chain complexes over $\mathbb{Z}$ with $\mathbb{R}_{\geqslant 0}$-grading.

3. Median graphs

Definition 3.1 (Median graphs)**.** A graph X is said to be *median* if, for any three pairwise distinct points $x, y, z \in X$, the intersection $[x, y] \cap [y, z] \cap [x, z]$ consists of a single point, written $m(x, y, z)$.

Median graphs have been thoroughly studied. The two prototypical examples of median graphs are trees and hypercubes. Since, in a Cartesian product of graphs, the interval between two pairs of points is exactly the product of the respective intervals in each coordinate, it is easily seen that a Cartesian product of median graphs is still median. Any convex subgraph of a median graph is median as well, thus yielding many other examples. Finally, it has been shown by Chepoi ([**6**, Theorem 6.1]) that median graphs are exactly the 1-skeleta of CAT(0) cube complexes. The cycle graph of length 3 and the graph obtained by adding a diagonal path of length 2 to a square are both examples of *non* median graphs, the former failing at *existence*, and the latter at *uniqueness* of middle points (that is, vertices in the intersection of the intervals given by three distinct vertices).

In a graph X, all paths have integral length, which implies that the groups $\mathrm{MC}_k^l(X)$ vanish when $l \notin \mathbb{N}$. Furthermore, no k-path has length less than k, so the groups $\mathrm{MC}_k^l(X)$ also vanish when $l < k$. Graphically, this means that, when placed on a quadrant of the $\mathbb{Z}^2$ grid, all magnitude homology groups vanish "above the diagonal".

Definition 3.2 (Diagonality ([**8**, Definition 7.1]))**.** A graph X is *diagonal* if, for all $k \neq l \in \mathbb{N}$, the group $\mathrm{MH}_k^l(X)$ vanishes.

In other words, X is diagonal if the magnitude homology groups vanish outside the diagonal.

In this section, and as preparation for Sections 5 and 6, we prove the following:

Proposition 3.3. *Median graphs are diagonal.*

We will shortly present a beautiful characterisation of median graphs, due to Bandelt, which will be used in the proof of Proposition 3.3. Some preliminary definitions are necessary:

Fix a set X, and let Q_X be the graph whose vertices are the *finite* subsets of X, and which has an edge between two subsets if and only if either contains exactly one more element than the other. A graph isomorphic to some Q_X is called a *hypercube* ([**4**]).

If X is a graph, $Y \subseteq X$ a subspace (recall our convention on subspaces of graphs), and $f\colon X \to Y$ a 1-Lipschitz map satisfying $f|_Y = \mathrm{Id}_Y$ *and* such that $d(x, x') = 1$ implies $d(fx, fx') = 1$ for all $x, x' \in X$, then we say that Y is an *edge-preserved* retract of X. Note that edge-preserved retracts are in particular retracts, by definition.

Theorem 3.4 ([2, Theorem 2]). *Median graphs are precisely the edge-preserved retracts of hypercubes.*

We need three simple properties of diagonality, whose proofs we will not linger on, since generalisations will be given in Section 5.

Proposition 3.5 ([8, Proposition 7.3]). *Cartesian products of diagonal graphs are diagonal.*

Proof. By applying the Künneth formula and noting that diagonal graphs have torsion-free homologies. See [8, Proposition 7.3] for the full argument. □

Corollary 3.6. *Finite hypercubes are diagonal.*

Proof. The complete graph on two vertices K_2 is diagonal, and a finite hypercube is a finite Cartesian product of copies of K_2. □

Proposition 3.7. *Retracts of diagonal graphs are diagonal.*

Proof. A retraction $f\colon X \to Y$ has left inverse the inclusion $\iota\colon Y \to X$. Functoriality of MH^l_k implies that $\mathrm{MH}^l_k(f) : \mathrm{MH}^l_k(X) \to \mathrm{MH}^l_k(Y)$ is surjective. Hence, if $\mathrm{MH}^l_k(X)$ vanishes outside the diagonal, so does $\mathrm{MH}^l_k(Y) = 0$. □

Proposition 3.8. *Graphs with filtrations by diagonal graphs are diagonal.*

Proof. Let $(U_\alpha)_\alpha$ be a filtration of X, so that $\mathrm{MZ}^l_k(X) = \bigcup_\alpha \mathrm{MZ}^l_k(U_\alpha)$ and $\mathrm{MB}^l_k(X) = \bigcup_\alpha \mathrm{MB}^l_k(U_\alpha)$. If for all α and $k \neq l$, $\mathrm{MZ}^l_k(U_\alpha) = \mathrm{MB}^l_k(U_\alpha)$, then $\mathrm{MZ}^l_k(X) = \mathrm{MB}^l_k(X)$ and the homology vanishes outside the diagonal. □

We can now proceed with the proof of Proposition 3.3:

Proof of Proposition 3.3. Fix a median graph X and a hypercube Q of which X is a retract. Q has a filtration by finite hypercubes, which are diagonal (Corollary 3.6); hence so is Q (Proposition 3.8), and thus X (Proposition 3.7). □

4. The Künneth and Mayer–Vietoris formulae

In [8], Hepworth and Willerton describe versions of the Künneth, excision, and Mayer–Vietoris formulae for magnitude homology of graphs. In [13], Leinster and Shulman, extending magnitude homology to metric spaces, asked whether those extend to this new setting. The answer is yes, assuming the right reinterpretations. More precisely:

Hepworth and Willerton's statement and proof of the Künneth theorem ([8, Theorem 5.3]) extend verbatim to l^1 products of metric spaces.

Similarly, Hepworth and Willerton's statement and proof of the excision and Mayer–Vietoris formulae ([8, Theorem 6.5]) extend to *gated* decompositions of metric spaces (Definition 4.12) with minimal changes. Those "minimal changes" are a bit trickier than simple generalisations. In particular, the "metric excision formula" that we define is not strictly a generalisation of the graph-theoretic one of [8], since the definitions we use are not themselves generalisations of the ones in [8].

Since our arguments mainly consist in tweaking the original constructions of Hepworth and Willerton, having a copy of [8] at hand will prove useful!

4.1. The Künneth Formula

If X, Y are metric spaces, we endow the Cartesian product $X \times Y$ with the l^1 metric:

$$d((x,y),(x',y')) := d(x,x') + d(y,y').$$

This implies that the intervals satisfy the identity

$$[(x,y),(x',y')] = [x,y] \times [x',y'],$$

which explains the l^1 metric's appearance in this context. Recall also that the l^1 product reduces to the usual Cartesian product in the case of graphs.

Remark 4.1 (Summary of differences). The arguments in [**8**, Section 8] go through verbatim when proving the Künneth formula in the case of metric spaces, since the main ingredient is the "interval structure", which generalises directly from graphs to metric spaces. Our downplaying of length as a grading of magnitude homology simplifies some expressions by virtue of getting rid of some $\bigoplus_l$s and $\bigvee_l$s; this is syntactical. Our arguments do not provide any new insight, but merely confirm that the generalisation holds.

We now retrace [**8**, Section 8] closely, with the metric case in mind.

Definition 4.2 (Interleavings, cross product ([**8**, Definition 5.2]))**.** Fix $n, l \in \mathbb{N}$, let $k = n + l$, and write $[k]$ for the set $\{0, \ldots, k\}$.

A map $\sigma = \langle \sigma_h, \sigma_v \rangle : [n+l] \to [n] \times [l]$ satisfying

- $\sigma(0) = (0,0)$ and $\sigma(n+l) = (n,l)$;
- if $\sigma(i) = (a,b)$, then $\sigma(i+1)$ is either $(a+1,b)$ or $(a,b+1)$,

is called a *staircase path.* Write $\llcorner n, l \urcorner$ for the set of (n,l) staircase paths. A staircase path is just a geodesic from $(0,0)$ to (n,l) in the obvious grid. The *sign* $\operatorname{sgn}\sigma$ of σ is $(-1)^s$, where s is the number of squares "below the staircase", i.e.

$$s = |\{(a,b) \in [n] \times [l] \,:\, a = \sigma_h(i) \Rightarrow b < \sigma_v(i)\}|.$$

If $\mathbf{x}$ is an n-path in X, $\mathbf{y}$ an l-path in Y, and $\sigma \in \llcorner n, l \urcorner$ a staircase path, the *interleaving of* $\mathbf{x}$ *and* $\mathbf{y}$ *along* σ is the k-path $\mathbf{x} \overset{\sigma}{\times} \mathbf{y}$ defined by $\mathbf{x} \overset{\sigma}{\times} \mathbf{y} := (\mathbf{x} \times \mathbf{y}) \circ \sigma$ (where $\mathbf{x} \times \mathbf{y}$ is identified with the map $[n] \times [l] \to X \times Y$).

The *cross product* is the morphism of chain complexes

$$\square \colon \mathrm{MC}_*(X) \otimes \mathrm{MC}_*(Y) \to \mathrm{MC}_*(X \times Y),$$

sending a pure tensor $\mathbf{x} \otimes \mathbf{y}$ to the alternating sum of all its interleavings:

$$\mathbf{x} \otimes \mathbf{y} \overset{\square}{\mapsto} \sum_{\sigma \in \llcorner n, l \urcorner} \operatorname{sgn}(\sigma)(\mathbf{x} \overset{\sigma}{\times} \mathbf{y}).$$

Verifying that $\square$ really defines a morphism mostly involves the same arguments as for the cross product map in simplicial homology: we will refer to those arguments as "generic". The non-generic part appears due to the possible vanishing of a "partial boundary" $\partial_i \mathbf{z}$, which happens when $z_i \notin [z_{i-1}, z_{i+1}]$. Thus, to show that $\square$ indeed commutes with the boundaries, we must understand the betweenness relations in a given interleaving $\mathbf{z} = \mathbf{x} \overset{\sigma}{\times} \mathbf{y}$ in terms of the betweenness relations in the paths $\mathbf{x}$ and $\mathbf{y}$, and of σ.

We visualise a staircase path $\sigma = \langle \sigma_h, \sigma_v \rangle$ as an actual (irregular) staircase on the $[n] \times [l]$ grid, going from bottom-left $(0,0)$ to top-right (n,l), with horizontal coordinate given by $\mathbf{x}$ and vertical by $\mathbf{y}$. Any $0 < m < n+l$ defines a point $\sigma(m)$ on the staircase; exactly one of:

A "corner": which means that its predecessor and successor differ in both coordinates. There are two distinct types of corners, looking like ⌜ and ⌟ respectively. In that case, $(\mathbf{x} \overset{\sigma}{\times} \mathbf{y})_m$ will necessarily be between its neighbors.

A "flat": which means that $\sigma_v(m+1) = \sigma_v(m) = \sigma_v(m-1)$ and $\sigma_h(m) = \sigma_h(m-1)+1 = \sigma_h(m+1)-1$. In that case $(\mathbf{x} \overset{\sigma}{\times} \mathbf{y})_m$ is between its neighbors if and only if $\mathbf{x}_{\sigma_h(m)}$ is between its neighbors, independently of $\mathbf{y}$.

A "wall": which means that $\sigma_h(m+1) = \sigma_h(m) = \sigma_h(m-1)$ and $\sigma_v(m) = \sigma_v(m-1)+1 = \sigma_v(m+1)-1$. In that case $(\mathbf{x} \overset{\sigma}{\times} \mathbf{y})_m$ is between its neighbors if and only if $\mathbf{y}_{\sigma_v(m)}$ is between its neighbors, independently of $\mathbf{x}$.

If σ is a staircase path and m is a corner point for σ, then there exists a unique other staircase path σ' with "dual" corner and opposite sign. More precisely, σ' and σ are equal, except at m, where

$$\sigma'(m) = (\sigma_h(m-1) + \sigma_v(m) - \sigma_v(m-1), \sigma_v(m-1) + \sigma_h(m) - \sigma_h(m-1)).$$

When applying ∂_m, the two interleavings $\mathbf{x} \overset{\sigma}{\times} \mathbf{y}$ and $\mathbf{x} \overset{\sigma'}{\times} \mathbf{y}$ will thus cancel out and the betweenness relations in $\mathbf{x}$ and $\mathbf{y}$ do not matter.

If m is a flat of σ, one can delete the column with coordinate $\sigma_h(m)$ and get a new staircase σ' in $\llcorner n-1, l \urcorner$. The sign of σ' differs from that of σ by $(-1)^{\sigma_v(m)}$, and

$$\partial_m(\mathbf{x} \overset{\sigma}{\times} \mathbf{y}) = \partial_{\sigma_h(m)} \mathbf{x} \overset{\sigma'}{\times} \mathbf{y}.$$

A similar description holds for walls. With these identities, it is a simple matter to adapt the "generic arguments" to the case of the magnitude complex.

Proposition 4.3 (Künneth theorem ([8, Theorem 4.3]))**.** *The cross product map induces a morphism*

$$\mathrm{MH}_*(X) \otimes \mathrm{MH}_*(Y) \xrightarrow{\square} \mathrm{MH}_*(X \times Y)$$
$$[f] \otimes [g] \mapsto [f \square g]$$

which fits into a natural short exact sequence

$$0 \longrightarrow \mathrm{MH}_*(X) \otimes \mathrm{MH}_*(Y) \xrightarrow{\square} \mathrm{MH}_*(X \times Y) \longrightarrow \mathrm{Tor}(\mathrm{MH}_{*-1}(X), \mathrm{MH}_*(Y)) \longrightarrow 0.$$

To prove Proposition 4.3 (at the end of the subsection), we first set the stage and verify a few preliminary results.

Definition 4.4 ([8, Definition 8.1])**.** If X is a metric space, we define the pointed simplicial set $\mathrm{MS}(X)$ with, as k-simplices, the $(k+1)$-tuples of points $\langle x_0, \ldots, x_k \rangle : [k] \to X$ in X, plus basepoint simplices pt_k, along with face and degeneracy maps defined by

$$\mathsf{d}_{k,i}\langle x_0,\dots,x_i,\dots,x_k\rangle := \begin{cases}\langle x_0,\dots,\widehat{x_i},\dots,x_k\rangle & \text{if } x_i\in[x_{i-1},x_{i+1}]\ ,\\ \mathrm{pt}_{k-1} & \text{otherwise,}\end{cases}$$

and

$$\mathsf{s}_{k,i}\langle x_0,\dots,x_i,\dots,x_k\rangle := \langle x_0,\dots,x_i,x_i,\dots,x_k\rangle,$$

and on basepoints:

$$\mathsf{d}_{k,i}\mathrm{pt}_k := \mathrm{pt}_{k-1},$$
$$\mathsf{s}_{k,i}\mathrm{pt}_k := \mathrm{pt}_{k+1}.$$

In [8, Definition 8.1], for a graph G and $l\in\mathbb{N}$, the simplicial set $M_l(G)$ is defined. In our notation, this set corresponds to the sub-simplicial set of $\mathrm{MS}(G)$ obtained by restricting to paths of length l, so that our $\mathrm{MS}(G)$ is equal to their $\bigvee_l M_l(G)$.

Proposition 4.5 ([8, Proposition 8.2])**.** *Let X,Y be metric spaces. The morphism of pointed simplicial sets*

$$\square\colon \mathrm{MS}(X)\wedge\mathrm{MS}(Y)\to\mathrm{MS}(X\times Y)$$
$$[\mathbf{x},\mathbf{y}]\mapsto\langle\mathbf{x},\mathbf{y}\rangle$$

is an isomorphism.

Let us clarify the notation: $\mathbf{x}$ and $\mathbf{y}$ are k-simplices of $\mathrm{MS}(X)$ and $\mathrm{MS}(Y)$ respectively, that is, maps $[k]\to X$ and $[k]\to Y$. Thus, the pair $(\mathbf{x},\mathbf{y})$ is an element of $\mathrm{MS}(X)\times\mathrm{MS}(Y)$, and $[\mathbf{x},\mathbf{y}]$ an element of $\mathrm{MS}(X)\wedge\mathrm{MS}(Y)$. Finally, $\langle\mathbf{x},\mathbf{y}\rangle:[k]\to X\times Y$ is the "product" of the given maps, hence an element of $\mathrm{MS}(X\times Y)$.

Proof. Bijectivity and commutation with degeneracy maps is clear. For face maps, one uses the product identity for intervals in the l^1 product, plus the fact that $\mathsf{d}_{k,i}(\mathbf{x},\mathbf{y})\neq\mathrm{pt}_{k-1}$ if and only if both $\mathsf{d}_{k,i}\mathbf{x}\neq\mathrm{pt}_{k-1}$ and $\mathsf{d}_{k,i}\mathbf{y}\neq\mathrm{pt}_{k-1}$ hold. □

Still following [8], given a pointed simplicial set S, the *normalised reduced* chain complex $\overline{\mathrm{N}}_*(S)$ associated to S is defined by

$$\overline{\mathrm{N}}_k(S) := \mathbb{Z}[\{k\text{-simplices}\}]/\mathbb{Z}[\{\text{degenerate and basepoint simplices}\}],$$

with boundary map induced by

$$\partial_k=\sum_{i=0}^{k}(-1)^i\mathsf{d}_{k,i}.$$

Since a simplex in $\mathrm{MS}(X)$ is degenerate if and only if two consecutive points are equal, the following holds:

Proposition 4.6 ([8, Lemma 8.3])**.** *$\overline{\mathrm{N}}_k(\mathrm{MS}(X))$ and $\mathrm{MC}_k(X)$ are isomorphic chain complexes.*

Proof. $\mathrm{MC}_k(X)$ is generated by the k-paths in X; that is, the $(k+1)$-tuples of consecutively distinct points in X. $\overline{\mathrm{N}}_k(\mathrm{MS}(X))$ is generated by the non-degenerate non-basepoint simplices of $\mathrm{MS}(X)$, which are exactly the k-paths. Thus, the groups are isomorphic. On $\overline{\mathrm{N}}_k(\mathrm{MS}(X))$, the boundary is defined as $\partial_k = \sum_{i=1}^{k-1}(-1)^i \mathsf{d}_{k,i}$. Since $\mathsf{d}_{k,i}$ sends a simplex $\mathbf{x}$ to a basepoint if and only if $x_i \notin [x_{i-1}, x_{i+1}]$, $\mathsf{d}_{k,i}$ sends $\mathbf{x}$ to zero at the level of chain maps, which shows that the boundary maps agree. □

From now on, we will identify $\overline{\mathrm{N}}_k(\mathrm{MS}(X))$ with $\mathrm{MC}_k(X)$.

Remember that given a simplicial set S, there exists, for any k, a natural bijection

$$s(\bullet)\colon S_k \leftrightarrow \mathrm{Mor}(\Delta[k], S),$$

where $\Delta[k]$ is the canonical k-simplex. The (inverse of this) bijection is obtained by sending a morphism $f : \Delta[k] \to S$ to the image through f of the single non-degenerate k-simplex $\mathrm{Id} : [k] \to [k]$ in $\Delta[k]$. Under the interpretation of simplicial sets as presheaves, this bijection is simply the Yoneda lemma.

If $\sigma = \langle \sigma_h, \sigma_v \rangle$ is a (n,l)-staircase, σ defines a morphism of simplicial complexes

$$\sigma_*\colon \Delta[n+l] \to \Delta[n] \times \Delta[l],$$

by sending a face $\phi : [m] \to [n+l]$ of $\Delta[n+l]$ to the pair of faces

$$(\sigma_h \circ \phi : [m] \to [n], \sigma_v \circ \phi : [m] \to [l])$$

in $\Delta[n] \times \Delta[l]$. Finally, if $\mathbf{x}$ and $\mathbf{y}$ are simplices in the simplicial sets S and T respectively, they are naturally associated to morphisms $s(\mathbf{x}) : \Delta[n] \to S, s(\mathbf{y})\colon \Delta[l] \to T$, so that one has a morphism:

$$s(\mathbf{x}) \times s(\mathbf{y}) : \Delta[n] \times \Delta[l] \to S \times T.$$

Given pointed simplicial sets S, T, we now define the *reduced Eilenberg–Zilber map*

$$\begin{aligned} \nabla^{\overline{N}}\colon \overline{\mathrm{N}}_*(S) \otimes \overline{\mathrm{N}}_*(T) &\to \overline{\mathrm{N}}_*(S \wedge T) \\ \mathbf{x} \otimes \mathbf{y} \in \overline{\mathrm{N}}_n(S) \otimes \overline{\mathrm{N}}_l(T) &\mapsto \sum_{\sigma \in \llcorner n,l \urcorner} [s^{-1}((s(\mathbf{x}) \times s(\mathbf{y})) \circ \sigma_*)], \end{aligned}$$

where $[\bullet]\colon S \times T \to S \wedge T$ is the collapsing map.

The following abstract property of $\nabla^{\overline{N}}$ is proven in [8]:

Proposition 4.7 ([8, Proposition 8.4])**.** $\nabla^{\overline{N}}$ *is a quasi-isomorphism.*

Let us now concretely describe the map $\nabla^{\overline{N}}$ in the case at hand. Fix generators $\mathbf{x} \in \overline{\mathrm{N}}_n(\mathrm{MS}(X))$ and $\mathbf{y} \in \overline{\mathrm{N}}_l(\mathrm{MS}(Y))$. When $\mathbf{x}$ is seen as a simplex of $\mathrm{MS}(X)$, we have

$$\mathbf{x}\colon [n] \to X.$$

Through the identification "simplex↔morphism", $\mathbf{x}$ becomes

$$\begin{gathered} s(\mathbf{x})\colon \Delta[n] \to \mathrm{MS}(X) \\ (\phi\colon [m] \to [n]) \in \Delta[n]_m \mapsto (\mathbf{x} \circ \phi : [m] \to X) \in \mathrm{MS}(X)_m. \end{gathered}$$

Thus, the composite

$$(s(\mathbf{x}) \times s(\mathbf{y})) \circ \sigma_*\colon \Delta[n+l] \to \mathrm{MS}(X) \times \mathrm{MS}(Y)$$

is defined as

$$(\phi : [m] \to [n+l]) \in \Delta[n+l]_m \mapsto ((s(\mathbf{x}) \times s(\mathbf{y})) \circ \sigma_*)(\phi) \in \mathrm{MS}(X) \times \mathrm{MS}(Y) = (\mathbf{x} \circ \sigma_h \circ \phi, \mathbf{y} \circ \sigma_v \circ \phi),$$

and passing back from morphisms to simplices (evaluating at $\mathrm{Id} : [n+l] \to [n+l]$), the result is simply

$$(\mathbf{x} \circ \sigma_h \circ \mathrm{Id}, \mathbf{y} \circ \sigma_v \circ \mathrm{Id}) = (\mathbf{x} \circ \sigma_h, \mathbf{y} \circ \sigma_v) \in \mathrm{MS}(X)_n \times \mathrm{MS}(Y)_l.$$

Proposition 4.8 ([8, Proof of Theorem 5.3])**.** *The cross product*

$$\square\colon \mathrm{MC}_*(X) \otimes \mathrm{MC}_*(Y) \to \mathrm{MC}_*(X \times Y)$$

is a quasi-isomorphism.

Proof. The proof is essentially obtained by forgetting the grading in [8, Proof of Theorem 5.3]. More precisely, the map $\square$ is a quasi-isomorphism if and only if it is one at each "level" of the grading, i.e. if, for each $l \geqslant 0$, it restricts to a quasi-isomorphism:

$$\square_l\colon \bigoplus_{l_1+l_2=l} \mathrm{MC}_*^{l_1}(X) \otimes \mathrm{MC}_*^{l_2}(Y) \to \mathrm{MC}_*^{l}(X \times Y),$$

which is the content of [8, Proof of Theorem 5.3]. □

Proof of Proposition 4.3. Applying the algebraic Künneth formula to $\mathrm{MC}_*(X)$ and $\mathrm{MC}_*(Y)$ yields a short exact sequence

$$0 \to \mathrm{MH}_*(X) \otimes \mathrm{MH}_*(Y) \to H_*(\mathrm{MC}_*(X) \otimes \mathrm{MC}_*(Y)) \to \mathrm{Tor}(\mathrm{MH}_{*-1}(X), \mathrm{MH}_*(Y)) \to 0.$$

By Proposition 4.8, the middle term is isomorphic, through $H_*(\square)$, to $\mathrm{MH}_*(X \times Y)$. Naturality follows from naturality in the algebraic Künneth formula and that of the cross product map. □

Note that the "length aware" sequence in [8] can easily be recovered by fixing l in $H_*(\mathrm{MC}_*(X) \otimes \mathrm{MC}_*(Y))$.

4.2. The excision formula

Definition 4.9 (Gated sets)**.** Given a metric space X, a subset A of X is said to be *gated* if for any $x \in X$, there exists some $a_x \in A$ such that a_x is between x and all $a \in A$. The point a_x is called a *gate* between x and A.

Gated sets enjoy the following properties (see [7]—more about gated sets in median spaces can be found in [5]).

Proposition 4.10 ([7, pp. 114,112, 115 respectively])**.** *Let X be a metric space and A a gated subset of X. Then:*

- *A is convex.*
- *For any $x \in X$, there exists a* unique *gate $a_x \in A$ for x.*
- *The map $x \mapsto a_x$ is non-expanding, and is the identity on A.*

From now on, we write $\pi\colon X \to A$ for the map sending x to its gate a_x. Note that by the above, π is a (1-Lipschitz) retraction from X to A; it follows that $\mathrm{MC}_*(\pi)$: $\mathrm{MC}_*(X) \to \mathrm{MC}_*(A)$ is a (well-defined) epimorphism.

Remark 4.11 (Summary of differences). As for the Künneth formula, the proofs of excision and Mayer–Vietoris in [8] essentially generalise without trouble, yet some more care is needed, mainly because of slight differences in the definitions.

Apart from definitional differences, the main obstacle to generalising [8] appears in their [8, Lemma 9.5] and [8, Proof of Theorem 6.6][1] in which, once a length l is fixed, [8] uses the vanishing of the groups $\mathrm{MC}_k^l(X)$ for $k > l$; this does not hold in general for metric spaces.

Finally, unlike [8], we have not verified naturality of the Mayer–Vietoris sequence in the metric case.

To conclude, we will (again) follow [8, Section 9] very closely and highlight where changes are required.

Definition 4.12 (Gated decomposition)**.** Let X be a metric space and Y, Z, W subspaces satisfying $X = Y \cup Z$ and $W = Y \cap Z$. If W is gated with respect to Z, and for each $z \in Z$ and $y \in Y$, the intersection $W \cap [y, z]$ is non-empty, then we say that the triple $(X; Y, Z)$ is a *gated decomposition* of X.

Note that from W being gated in Z and $W \cap [y, z]$ being non-empty, it follows that $\pi(z)$ lies between z and any element of Y.

Following [8], we write $\mathrm{MC}_*(Y, Z)$ for the subcomplex of $\mathrm{MC}_*(X)$ spanned by paths entirely contained in either Y or Z. We can now state the excision theorem:

Theorem 4.13 (excision—metric setting)**.** *If $(X; Y, Z)$ is a gated decomposition of X, then the inclusion*

$$\mathrm{MC}_*(Y, Z) \hookrightarrow \mathrm{MC}_*(X)$$

is a quasi-isomorphism.

The Mayer–Vietoris theorem follows easily from excision.

Theorem 4.14 (Mayer–Vietoris—metric case)**.** *If $(X; Y, Z)$ is a gated decomposition of X and $W = Y \cap Z$, then the inclusions*

$$j_Y\colon W \to Y, \quad j_Z\colon W \to Z, \quad i_Y\colon Y \to X, \quad i_Z\colon Z \to X$$

induce a short exact sequence

$$0 \to \mathrm{MH}_*(W) \xrightarrow{\langle (j_Y)_*, -(j_Z)_* \rangle} \mathrm{MH}_*(Y) \oplus \mathrm{MH}_*(Z) \xrightarrow{(i_Y)_* \oplus (i_Z)_*} \mathrm{MH}_*(X) \to 0.$$

Proof. The proof is obtained from the argument in [8, Proof of Theorem 6.6, assuming Theorem 6.5] by dropping the grading (and noting that since we deal with metric spaces, we do not have to worry about path-connected components as they do). In our case Theorem 4.13 takes on the role of their Theorem 6.5. ☐

[1]Note that [8] contains both a "Proof of Theorem 6.6 assuming Theorem 6.5" (page 50), and a "Proof of Theorem 6.6" (page 52). The latter is actually the proof of Theorem 6.5.

4.2.1. Proof of excision

For the remainder of the section, let us fix a gated decomposition $(X; Y, Z)$ and let $W := Y \cap Z$.

We define, for $a \in Y - Z, b \in Z - Y$ (or vice versa) and $k \geqslant 0$:

$$A_k(a,b) := \mathbb{Z}[\{\langle x_0, \ldots, x_k \rangle \,:\, x_0 = a, x_k = b, x_1, \ldots, x_{k-1} \in W\}] \leqslant \mathrm{MC}_k(X).$$

For $b \in Z - Y$ we define

$$B_k(b) := \mathbb{Z}[\{\langle x_0, \ldots, x_k \rangle \,:\, x_k = b, x_0, \ldots, x_{k-1} \in Y\}] \leqslant \mathrm{MC}_k(X),$$
$$\tilde{B}_k(b) := \mathbb{Z}[\{\langle x_0, \ldots, x_k \rangle \,:\, x_k = b, x_0, \ldots, x_{k-1} \in W\}] \leqslant \mathrm{MC}_k(X),$$

and for $i \in \mathbb{N}$

$$F_k(b; i) := \mathbb{Z}[\{\langle x_0, \ldots, x_k \rangle \,:\, x_k = b, x_0, \ldots, x_{i-1} \in Y, x_i, \ldots, x_{k-1} \in W\}] \leqslant \mathrm{MC}_k(X),$$

and symmetrically for $b \in Y - Z$. Finally for $i \in \mathbb{N}$ set

$$G_k(i) := \mathbb{Z}[\{\langle x_0, \ldots, x_k \rangle \,:\, x_0, \ldots, x_{k-i} \text{ all lie in } Y, \text{ or all lie in } Z\}] \leqslant \mathrm{MC}_k(X).$$

These all define sub-chain complexes of $\mathrm{MC}_*(X)$; the definitions match the ones found in [**8**, Section 9], except that our last $G_*(i)$s correspond to the F_is found in [**8**, Proof of Theorem 6.6, p. 52].

It is clear that $G_*(0) = \mathrm{MC}_*(Y, Z)$, $G_k(l) = \mathrm{MC}_k(X)$ for all $k \leqslant l$, and $G_*(l) \leqslant G_*(l+1)$ for all l. It follows that $\mathrm{MC}_*(X)$ is the direct limit of the system

$$\mathrm{MC}_*(Y, Z) = G_*(0) \leqslant G_*(1) \leqslant \cdots \leqslant G_*(l) \leqslant G_*(l+1) \leqslant \cdots.$$

Thus, to show that the inclusion $\mathrm{MC}_*(Y, Z) \hookrightarrow \mathrm{MC}_*(X)$ is a quasi-isomorphism, it is enough to do so for each inclusion $G_*(l) \hookrightarrow G_*(l+1)$. Indeed, once this is done, the whole system becomes a chain of isomorphisms after passing to homology, and each inclusion of $G_*(l)$ to the limit $\mathrm{MC}_*(X)$ as well. In particular, this is true for the inclusion $\mathrm{MC}_*(Y, Z) \hookrightarrow \mathrm{MC}_*(X)$.

This is essentially the only thing we have to change from the argument of [**8**].

Let l be fixed from now on, and given a chain complex C_*, write $\Sigma^j C_*$ for the shifted chain complex $(\Sigma^j C_*)_k = C_{k-j}$.

Proposition 4.15 ([**8**, Lemma 9.2])**.** *The complex $A_*(a,b)$ is acyclic.*

Proof. The proof of [**8**, Lemma 9.2] applies verbatim, once the grading is dropped.

Recall that gated decompositions are required to satisfy the following: for $y \in Y$ and $z \in Z$, $W \cap [y, z] \neq \emptyset$. This condition is necessary for the case $k = 1$ in the cited proof: if $a \in Y - Z$ and $b \in Z - Y$, it is necessary for a point of W between a and b to exist, so that $\pi(b)$ is well-defined. □

Define the set

$$J_Z(l) := \{\mathbf{x} = \langle x_0, \ldots, x_l \rangle \,:\, x_0, \ldots x_l \in Y,\ x_l \notin Z\},$$

and define $J_Y(l)$ symmetrically.

Proposition 4.16 ([**8**, Lemma 9.5])**.** *For any $b \in Z - Y$, there is an isomorphism*

$$F_*(b, l+1)/F_*(b, l) \cong \bigoplus_{\mathbf{x} \in J_Z(l)} \Sigma^l A_*(x_l, b).$$

In particular, the quotient $F_*(b, l)/F_*(b, 0)$ is acyclic. The same holds for $b \in Y - Z$ with $J_Z(l)$ replaced by $J_Y(l)$.

Proof. Apply the proof of [**8**, Lemma 9.5], or more precisely the part showing that each complex F_i/F_{i-1} (in their notation) is acyclic. □

Proposition 4.17. *For any $b \in Y \Delta Z$, the quotient $B_*(b)/\tilde{B}_*(b)$ is acyclic.*

Proof. Consider the directed system:

$$\tilde{B}_*(b) = F_*(b, 0) \leqslant F_*(b, 1) \leqslant \cdots \leqslant F_*(b, l) \leqslant F_*(b, l+1) \leqslant \cdots .$$

Since we have inclusions $F_*(b, l) \leqslant B_*(b)$ for all l, and for each $k \leqslant l$,

$$F_k(b, l) = B_k(b),$$

it follows that $B_*(b)$ is the direct limit of the above system. Passing to homology, each inclusion $F_*(b, l) \leqslant F_*(b, l+1)$ becomes an isomorphism by Proposition 4.16. It follows then that the inclusion $\tilde{B}_*(b) \leqslant B_*(b)$ also becomes an isomorphism. □

For $k \geqslant l$, define the set

$$K_k(l) := \{\mathbf{x} = \langle x_{k-l}, \ldots, x_k \rangle \,:\, x_{k-l} \in Y \Delta Z\},$$

and if $k < l$, let $K_k(l) := \emptyset$.

Proposition 4.18 ([**8**, Proof of Theorem 6.6 p. 52])**.** *There is an isomorphism of chain complexes*

$$G_k(l+1)/G_k(l) \cong \bigoplus_{\mathbf{x} \in K_k(l)} \left(\Sigma^l B_*(x_{k-l})/\tilde{B}_*(x_{k-l})\right)_k .$$

In particular, the inclusion $G_*(l) \leqslant G_*(l+1)$ is a quasi-isomorphism.

Proof. Apply the part of [**8**, Proof of Theorem 6.6 p. 52] showing that each quotient F_i/F_{i-1} (in their notation) is acyclic. □

We can now prove the excision formula:

Proof of Theorem 4.13. Each inclusion in the directed system

$$\mathrm{MC}_*(Y, Z) = G_*(0) \leqslant \cdots \leqslant G_*(l) \leqslant G_*(l+1) \leqslant \cdots$$

is a quasi-isomorphism by Proposition 4.18, and $\mathrm{MC}_*(X)$ is the direct limit of this system. Thus, the inclusions induce isomorphisms

$$\mathrm{MH}_*(Y, Z) = H(G_*(0)) \cong \cdots \cong H(G_*(l)) \cong H(G_*(l+1)) \cong \cdots$$

and $\mathrm{MH}_*(X)$ (along with the morphisms induced by inclusions into $\mathrm{MH}_*(X)$) is the limit of this diagram. It follows that each inclusion, in particular $\mathrm{MC}_*(Y, Z) \hookrightarrow \mathrm{MC}_*(X)$ is a quasi-isomorphism. □

5. Diagonality

In the first section, we have seen that median graphs are diagonal (in the sense of Hepworth and Willerton). Knowing that median *graphs* are special cases of median *spaces* motivates us to try and find a corresponding description for median spaces. In this section, we introduce the notion of diagonality for metric spaces and verify some of its properties. As hoped, we will see in the next section that median spaces are indeed diagonal.

Recall that a path $\mathbf{x} = \langle x_0, \ldots, x_k \rangle$ is saturated if all strict intervals $]x_i, x_{i+1}[$ are empty. In what follows, for any chain $\sigma := \sum_{\mathbf{x}} \lambda_{\mathbf{x}} \mathbf{x}$, we will write σ_S for the "saturated" part of σ, that is, $\sigma_S := \sum_{\mathbf{x} \text{ saturated}} \lambda_{\mathbf{x}} \mathbf{x}$.

Definition 5.1 (Diagonality)**.** A space X is said to be *diagonal* if $\mathrm{MH}_k(X)$ is generated by linear combinations of saturated paths, for all $k \in \mathbb{N}$.

Let $S_k(X)$ denote the span of saturated paths, as a submodule of $\mathrm{MC}_k(X)$. Recall that the *support* of a chain $\sum_{\mathbf{x} \in P_k(X)} \lambda_{\mathbf{x}} \mathbf{x}$ is by definition the set $\{\mathbf{x} \in P_k(X) : \lambda_{\mathbf{x}} \neq 0\}$.

Proposition 5.2. *The supports of elements in $S_k(X)$ and $\mathrm{MB}_k(X)$ are mutually disjoint. In particular, $S_k(X) \cap \mathrm{MB}_k(X) = \{0\}$.*

Proof. By linearity, it suffices to verify that for any path $\mathbf{x}$, its boundary $\partial \mathbf{x} = \sum_i (-1)^i \partial_i \mathbf{x}$ has no saturated path in its support. This follows by definition of the boundary operator, since $\partial_i \mathbf{x}$ is non-zero exactly when $x_i \in]x_{i-1}, x_{i+1}[$, which implies that $\langle x_0, \ldots, x_{i-1}, x_{i+1}, \ldots x_k \rangle$ is not saturated. □

Thus, diagonality can be restated in different ways:

Proposition 5.3. *The following are equivalent:*

1. *X is diagonal;*
2. *$\mathrm{MZ}_k(X) = (\mathrm{MZ}_k(X) \cap S_k(X)) \oplus \mathrm{MB}_k(X)$;*
3. *Given a cycle $\sigma = \sigma_S + \sigma' \in \mathrm{MZ}_k(X)$, with σ_S the part corresponding to saturated paths and σ' the rest, we must have $\sigma' \in \mathrm{MB}_k(X)$;*
4. *The "inclusion-then-quotient" morphism $\mathrm{MZ}_k(X) \cap S_k(X) \to \mathrm{MH}_k(X)$ is an isomorphism.*

Proof. We show 1⇒2⇒4⇒1and 1⇔3.

1⇒2 By Proposition 5.2, $\mathrm{MZ}_k(X) \cap S_k(X)$ and $\mathrm{MB}_k(X)$ intersect trivially. By diagonality, given $\sigma \in \mathrm{MZ}_k(X)$, there exists $\sigma_s \in \mathrm{MZ}_k(X) \cap S_k(X)$ such that σ and σ_s are equal in homology, i.e. $\sigma - \sigma_s \in \mathrm{MB}_k(X)$. Thus, $\mathrm{MZ}_k(X) \cap S_k(X) + \mathrm{MB}_k(X) = \mathrm{MZ}_k(X)$.

2⇒4 Clear.

4⇒1 Clear.

1⇒3 Fix $\sigma = \sigma_S + \sigma' \in \mathrm{MZ}_k(X)$ as in 3. By diagonality, there exists $\sigma'_S \in \mathrm{MZ}_k(X) \cap S_k(X)$ and $\sigma'' \in \mathrm{MB}_k(X)$ such that $\sigma_S + \sigma' = \sigma'_S + \sigma''$, hence $\sigma_S - \sigma'_S = \sigma'' - \sigma'$. Since the support of $\sigma_S - \sigma'_S$ consists of saturated chains and σ'' has no saturated chains in its support, necessarily $\sigma' = \sigma''$ and $\sigma_S = \sigma'_S$.

3⇒1 Clear. □

In particular, it follows directly from the last item that:

Corollary 5.4. *A diagonal space has torsion-free homology.*

The following proposition supports our choice of terminology:

Proposition 5.5. *A graph is diagonal in the sense of Hepworth and Willerton if and only if it is diagonal in the above sense.*

Proof. In a graph, a k-path is saturated if and only if it has length k. Thus, having vanishing homology outside the diagonal and having homology groups generated by (linear combinations of) saturated paths are equivalent conditions. □

Corollary 5.6. *A diagonal Menger convex space has vanishing homology (except possibly at $k = 0$).*

Proof. A Menger convex space has no saturated path. □

In the next few propositions, we verify stability of diagonality under some usual constructions.

Proposition 5.7. *If $(U_\alpha)_\alpha$ is a filtration of a space X such that each U_α is diagonal, then so is X.*

Proof. Fix $\sigma \in \mathrm{MZ}_k(X)$, and write $\sigma = \sigma_S + \sigma'$. Fix α large enough that U_α contains all points in the support of σ, and, for each non-saturated path $\mathbf{x}$ in the support of σ', also contains a "witness" to non-saturation of the path (that is, a point $p \in]x_i, x_{i+1}[$ for some i). Then, $\sigma \in \mathrm{MZ}_k(U_\alpha)$, σ_S still consists of saturated paths *in* U_α, and σ' of non-saturated paths *in* U_α. Since U_α is diagonal, we conclude that $\sigma' \in \mathrm{MB}_k(U_\alpha) \leqslant \mathrm{MB}_k(X)$. □

Proposition 5.8. *An l^1 product of diagonal spaces is diagonal.*

Proof. By applying the (metric) Künneth formula. Let X, Y be diagonal spaces. For any fixed k, we have a short exact sequence

$$0 \to \bigoplus_{n+l=k} \mathrm{MH}_n(X) \otimes \mathrm{MH}_l(Y) \xrightarrow{\square} \mathrm{MH}_k(X \times Y) \to \mathrm{Tor}(\dots, \dots) \to 0.$$

The torsion part being zero (Corollary 5.4), an isomorphism

$$\bigoplus_{n+l=k} \mathrm{MH}_n(X) \otimes \mathrm{MH}_l(Y) \xrightarrow{\square} \mathrm{MH}_k(X \times Y)$$

remains. Since both $\mathrm{MH}_n(X)$ and $\mathrm{MH}_l(Y)$ are generated by (linear combinations of) saturated paths, the whole domain of the isomorphism is generated by pure tensors of such, and $\mathrm{MH}_k(X \times Y)$ by their images. Noting that if $\mathbf{z}$ is an interleaving of two paths $\mathbf{x}, \mathbf{y}$ as in the definition of the map □, then $\mathbf{z}$ is saturated if and only if both $\mathbf{x}, \mathbf{y}$ are, we conclude that $\mathrm{MH}_k(X \times Y)$ is generated by (linear combinations of) saturated paths. □

Proposition 5.9. *If $(X; Y, Z)$ is a gated decomposition and Y, Z are convex and diagonal, then X is diagonal.*

Proof. By applying a fragment of the (metric) Mayer–Vietoris sequence, we have an epimorphism

$$\mathrm{MH}_*(Y) \oplus \mathrm{MH}_*(Z) \to \mathrm{MH}_*(X),$$

and the images of saturated paths in either Y or Z are still saturated in X by convexity. □

Definition 5.10 (Betweenness preservation, reflection)**.** If X, Y are two metric spaces, and $f\colon X \to Y$ is an injective map, we say that it

- *preserves betweenness* if $y \in [x, z]$ implies $fy \in [fx, fz]$; and
- *reflects betweenness* if $fy \in [fx, fz]$ implies $y \in [x, z]$.

In case f both preserves and reflects betweenness, we say it is a *betweenness embedding.* If it is also surjective, it becomes a *betweenness isomorphism.*

Let us give a few examples of maps that do or do not preserve or reflect betweenness:

- If X is a metric space, $\lambda > 0$ and λX denotes the metric space obtained by rescaling X by λ, then the identity map $X \to \lambda X$ is both preserving and reflecting.
- The composition of preserving (resp. reflecting) maps is still preserving (resp. reflecting). Thus, one sees that being preserving or reflecting is not related to being 1-Lipschitz, since, at least for finite metric spaces, one can always compose a map with a rescaling to be (or not) 1-Lipschitz.
- The projections onto a coordinate $\mathbb{R}^2 \to \mathbb{R}$ are preserving but not reflecting.
- The piecewise-linear map on the unit interval $[0, 1] \to [0, 1]$ defined by

$$x \mapsto \begin{cases} 2x & \text{if } x \leqslant \frac{1}{2} \\ 2(1-x) & \text{if } x \geqslant \frac{1}{2}, \end{cases}$$

  is neither preserving nor reflecting.

Proposition 5.11. *If $f\colon X \to Y$ is a betweenness embedding, then f induces a morphism of chain complexes:*

$$\begin{aligned} f_*\colon \mathrm{MC}_*(X) &\to \mathrm{MC}_*(Y) \\ (x_0, \ldots, x_k) &\mapsto (fx_0, \ldots, fx_k). \end{aligned}$$

If f is bijective, this turns into an isomorphism.

Note that this induced morphism is not the one given by the magnitude functor, since a betweenness embedding may fail to preserve the length of paths.

Proof. Injectivity plus the betweenness preserving and reflecting imply that f commutes with boundaries. □

Proposition 5.12. *If $f\colon X \to Y$ is a betweenness isomorphism between metric spaces and Y is diagonal, then so is X.*

Proof. Since f preserves and reflects betweenness, images of saturated paths are saturated, and vice versa. The same can be said of homological cycles and boundaries. □

Proposition 5.13. *Retracts of diagonal spaces are diagonal.*

Proof. Let $f\colon X \to Y$ be a retraction.

If $\mathbf{y}$ is a path in Y, then $\mathbf{y}$ is saturated in Y if and only if it is saturated in X. Indeed, assume first that $\mathbf{y}$ is not saturated in X, so that there exists some $x \in X$ strictly between y_i and y_{i+1}. Then, since f is non-expanding and fixes y_i and y_{i+1}, it follows that $fx \in Y$ is also strictly between y_i and y_{i+1}, so that $\mathbf{y}$ is not saturated in Y. This shows that non-saturatedness in X implies non-saturatedness in Y, and the converse is obvious.

Consider now a cycle $\sigma = \sigma_S + \sigma' \in \mathrm{MZ}_k(Y)$. Since $Y \subseteq X$, σ is still a cycle in X, and its decomposition into "saturated+non-saturated" in X is still $\sigma = \sigma_S + \sigma'$. Thus, assuming X is diagonal, $\sigma' \in \mathrm{MB}_k(X)$, that is, there exists some $\tau \in \mathrm{MC}_{k+1}(X)$ with $\sigma' = \partial\tau$. Applying $f_*\colon \mathrm{MC}_*(X) \to \mathrm{MC}_*(Y)$ to σ', we get $f_*\sigma' = f_*\partial\tau = \partial f_*\tau$. Since σ' has support in Y, $f_*\sigma' = \sigma'$. Thus $\sigma' = \partial f_*\tau \in \mathrm{MB}_k(Y)$. This shows that Y is diagonal. □

6. Median spaces are diagonal

We will need the following fact due to Avann:

Proposition 6.1 ([**1**]). *If X is a finite median space, then there exists a finite graph $G(X)$ and a betweenness isomorphism $\phi\colon X \to G(X)$. In particular, $G(X)$ is median.*

Another important property of median spaces:

Proposition 6.2. *Any median space has a filtration by finite median subspaces.*

Proof. Any finite subset of a median space X has a finite so-called *median hull*, that is, a smallest median subspace of X containing the set in question (see [**3**, p. 7]). Taking all such finite median hulls yields a filtration by finite median subspaces. □

It is now easy to conclude with the following:

Proposition 6.3. *Median spaces are diagonal.*

Proof. Finite median graphs are diagonal, and by applying Propositions 5.12 and 6.1, so are finite median spaces. □

Corollary 6.4. *Median Menger convex spaces have vanishing homology (except possibly at $k = 0$).*

Menger convexity is already known to be strongly related to vanishing in the magnitude homology:

- In [**13**], it is shown that $\mathrm{MH}_1(X) = 0$ if and only if X is Menger convex, and $\mathrm{MH}_2(X) = 0$ if X is Menger convex and satisfies two more "straightness" conditions (Corollary 4.5 and Theorem 4.21 respectively).
- [**9**, Corollary 4.9] shows that if $\mathrm{MH}_n(X) = 0$ for some $n \geqslant 1$, then X is Menger convex. Conversely, and extending the results of [**13**], [**9**, Corollary 7.3] shows that assuming Menger convexity and the same "straightness" conditions as above, $\mathrm{MH}_n(X) = 0$ for all $n \geqslant 1$ (this result is also obtained in [**10**]).

Those two straightness conditions are *geodeticy* and *no* 4*-cuts*:

- X is geodetic if for any $x, y, z, w \in X$, it follows from $z, w \in [x, y]$ that either $z \in [x, w]$ or $z \in [w, y]$.

- X has no 4-cuts if for any sequence $x \neq y \neq z \neq w \in X$, it follows from $y \in [x, z]$ and $z \in [y, w]$ that $y, z \in [x, w]$.

Note that median spaces do not necessarily satisfy either geodeticy or having no 4-cuts: the four-cycle graph C_4 being a prime example of a graph satisfying neither. Conversely, Euclidean space is geodetic and has no 4-cuts, but is not median. As such, Corollary 6.4 yields a new crop of spaces with vanishing homology.

References

[1] S. P. Avann. Metric ternary distributive semi-lattices. *Proc. Amer. Math. Soc.*, 12:407–414, 1961.

[2] H.-J. Bandelt. Retracts of hypercubes. *J. Graph Theory*, 8(4):501–510, 1984.

[3] H.-J. Bandelt and V. Chepoi. Metric graph theory and geometry: a survey. In *Surveys on discrete and computational geometry*, volume 453 of *Contemp. Math.*, pages 49–86. Amer. Math. Soc., Providence, RI, 2008.

[4] H.-J. Bandelt and H. M. Mulder. Infinite median graphs, (0, 2)-graphs, and hypercubes. *J. Graph Theory*, 7(4):487–497, 1983.

[5] I. Chatterji, C. Druţu, and F. Haglund. Kazhdan and Haagerup properties from the median viewpoint. *Adv. Math.*, 225(2):882–921, 2010.

[6] V. Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.

[7] A. W. M. Dress and R. Scharlau. Gated sets in metric spaces. *Aequationes Math.*, 34(1):112–120, 1987.

[8] R. Hepworth and S. Willerton. Categorifying the magnitude of a graph. *Homology Homotopy Appl.*, 19(2):31–60, 2017.

[9] B. Jubin. On the magnitude homology of metric spaces, 2018, arXiv:1803.05062.

[10] R. Kaneta and M. Yoshinaga. Magnitude homology of metric spaces and order complexes, 2018, arXiv:1803.04247.

[11] T. Leinster. Magnitude: a bibliography. `https://www.maths.ed.ac.uk/~tl/magbib/`. [Accessed October 1, 2020].

[12] T. Leinster. The magnitude of metric spaces. *Doc. Math.*, 18:857–905, 2013.

[13] T. Leinster and M. Shulman. Magnitude homology of enriched categories and metric spaces, 2017. To appear in *Alg. Geom. Topol.* arXiv:1711.00802.

Rémi Bottinelli `remi.bottinelli@unine.ch`

Institut de Mathématiques, Université de Neuchâtel, Rue Emile-Argand 11, CH-2000 Neuchâtel, Switzerland

Tom Kaiser `tomkaiser456@gmail.com`

Institut de Mathématiques, Université de Neuchâtel, Rue Emile-Argand 11, CH-2000 Neuchâtel, Switzerland