CATEGORIFYING THE MAGNITUDE OF A GRAPH

RICHARD HEPWORTH AND SIMON WILLERTON

(*communicated by Bill Murray*)

Abstract

The magnitude of a graph can be thought of as an integer power series associated to a graph; Leinster introduced it using his idea of magnitude of a metric space. Here we introduce a bigraded homology theory for graphs which has the magnitude as its graded Euler characteristic. This is a categorification of the magnitude in the same spirit as Khovanov homology is a categorification of the Jones polynomial. We show how properties of magnitude proved by Leinster categorify to properties such as a Künneth Theorem and a Mayer-Vietoris Theorem. We prove that joins of graphs have their homology supported on the diagonal. Finally, we give various computer calculated examples.

1. Introduction

1.1. Overview

The magnitude of a finite metric space was introduced by Leinster [9] by analogy with his notion of the Euler characteristic of a category [8]. This was found to have connections with topics as varied as intrinsic volumes [13], biodiversity [12], potential theory [16], Minkowski dimension [16] and curvature [20].

This invariant of finite metric spaces can be used to construct an invariant of finite graphs. For G a finite graph and $t > 0$, we equip the set of vertices of G with the shortest path metric on G where each edge is given length t. Leinster [11] showed that as a function of t, the magnitude of this metric space is a rational function in e^{-t}. Writing $q = e^{-t}$, the magnitude can be expanded as a formal power series in q and Leinster proved that this power series has integer coefficients. It is this integer power series that we will take as the magnitude of G, and we will write it as $\#G$. For example, the five-cycle graph has magnitude which starts as follows:

$$\#C_5 = 5 - 10q + 10q^2 - 20q^4 + 40q^5 - 40q^6 - 80q^8 + \cdots.$$

In this paper we will categorify the magnitude of graphs by defining *magnitude homology of graphs*. This is a bigraded homology theory $\mathrm{MH}_{*,*}$. It is functorial with
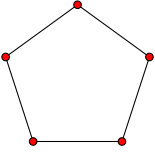
	k	0	1	2	3	4	5	6	7	8	9	10	11
	0	5											
	1		10										
	2			10									
	3			10	10								
	4				30	10							
	5					50	10						
l	6					20	70	10					
	7						80	90	10				
	8							180	110	10			
	9							40	320	130	10		
	10								200	500	150	10	
	11									560	720	170	10

Table 1: The ranks of $\mathrm{MH}_{k,l}(C_5)$, the magnitude homology groups of the pictured five-cycle graph, as computed using Sage.

respect to maps of graphs that send vertices to vertices and preserve or contract edges, and its graded Euler characteristic recovers the magnitude:

$$\#G = \sum_{k,l \geqslant 0} (-1)^k \cdot \mathrm{rank}\big(\mathrm{MH}_{k,l}(G)\big) \cdot q^l = \sum_{l \geqslant 0} \chi\big(\mathrm{MH}_{*,l}(G)\big) \cdot q^l. \qquad (1)$$

Thus our categorification is in exactly the same spirit as as existing categorifications of polynomial invariants in knot theory and graph theory, for example Khovanov's categorification of the Jones polynomial [6], Ozsvath-Szabo's categorification of the Alexander polynomial [17], Helme-Guizon and Rong's categorification of the chromatic polynomial [4], Jasso-Hernandez and Rong's categorification of the Tutte polynomial [5], and Stošić's categorification of the dichromatic polynomial [18].

As an example of magnitude homology, the ranks of the magnitude homology groups of the five-cycle graph are given in Table 1. You can verify that the Euler characteristic of each of the first few rows is the corresponding coefficient in $\#C_5$.

Being a bigraded abelian group rather than just a power series, the magnitude homology has a richer structure than the magnitude. For example, functoriality means that for a given graph its magnitude homology is equipped with an action of its automorphism group. We will see below that various properties of magnitude described by Leinster in [11] are shadows of properties of magnitude homology.

Leinster has given a counting formula [11, Proposition 3.9] for the magnitude. It expresses the coefficient of q^l in $\#G$ as

$$\sum_{k \geqslant 0} (-1)^k \Big| \big\{ (x_0, \dots, x_k) : \ x_i \in V(G), \ x_i \neq x_{i+1}, \ \textstyle\sum_{i=0}^{k-1} d(x_i, x_{i+1}) = l \big\} \Big|.$$

This expression is precisely the alternating sum of the ranks of the magnitude chain groups, and in general these ranks are considerably larger than the ranks of magnitude homology groups. In Table 2, the ranks of the magnitude chain groups for the five-cycle graph are given and this should be compared with Table 1. Again the Euler characteristic of each row gives a coefficient of the magnitude, but the terms grow

	0	1	2	3	4	k 5	6	7	8	9	10	11
0	5											
1		10										
2		10	20									
3			40	40								
4			20	120	80							
5				120	320	160						
6				40	480	800	320					
7					320	1600	1920	640				
8					80	1600	4800	4480	1280			
9						800	6400	13440	10240	2560		
10						160	4800	22400	35840	23040	5120	
11							1920	22400	71680	92160	51200	10240

Table 2: The ranks of $\mathrm{MC}_{k,l}(C_5)$ the magnitude chain groups of the five-cycle graph, as computed using Sage.

exponentially as you move down diagonally. This means that the magnitude homology groups are counting something much subtler than Leinster's formula is.

1.2. Categorifying properties of the magnitude

Many of the properties of the magnitude that were proved by Leinster in [11] can be categorified, meaning that they follow from properties of the magnitude homology upon taking the graded Euler characteristic. The categorified properties are subtler, being properties of the homology rather than its Euler characteristic, and are correspondingly harder to prove. We list the categorifications here.

1.2.1. Disjoint unions

Leinster shows that magnitude is additive with respect to the disjoint union of graphs [11, Lemma 3.5]:

$$\#(G \sqcup H) = \#G + \#H$$

Our categorification of this, Proposition 4.1, is the additivity of the magnitude homology:

$$\mathrm{MH}_{*,*}(G \sqcup H) \cong \mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(H).$$

Taking the graded Euler characteristic of both sides recovers Leinster's formula for $\#(G \sqcup H)$.

1.2.2. Products

Leinster shows that magnitude is multiplicative with respect to the cartesian product $\square$ of graphs [11, Lemma 3.6]:

$$\#(G \square H) = \#G \cdot \#H.$$

The categorification of this is Theorem 5.3, a Künneth Theorem which says that there is a non-naturally split, short exact sequence:

$$0 \to \mathrm{MH}_{*,*}(G) \otimes \mathrm{MH}_{*,*}(H) \to \mathrm{MH}_{*,*}(G \,\square\, H)$$
$$\to \mathrm{Tor}\big(\mathrm{MH}_{*+1,*}(G), \mathrm{MH}_{*,*}(H)\big) \to 0.$$

Taking the graded Euler characteristic of this sequence recovers the multiplicativity of the magnitude. Moreover, if either G or H has torsion-free magnitude homology, then this sequence reduces to an isomorphism $\mathrm{MH}_{*,*}(G) \otimes \mathrm{MH}_{*,*}(H) \cong \mathrm{MH}_{*,*}(G \,\square\, H)$. At the time of writing, despite quite a bit of computation, we don't know whether any graphs have torsion in their magnitude homology.

1.2.3.   Unions

Magnitude can be extended to infinite metric spaces [15] and the Convexity Conjecture [13] gives an explicit formula for the magnitude of compact, convex subsets of $\mathbb{R}^n$. A corollary of the conjecture would be that the magnitude of compact, convex subsets satisfies an inclusion-exclusion formula. Leinster showed that an analogue of this corollary holds for graphs. If $(X; G, H)$ is a projecting decomposition (see Section 6), so that in particular, $X = G \cup H$, then the inclusion-exclusion formula holds for magnitude [11, Theorem 4.9]:

$$\#X = \#G + \#H - \#(G \cap H).$$

Our categorification of this result, Theorem 6.6, is that if $(X; G, H)$ is a projecting decomposition, then there is a naturally split short exact sequence

$$0 \to \mathrm{MH}_{*,*}(G \cap H) \to \mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(H) \to \mathrm{MH}_{*,*}(X) \to 0$$

(which we think of as a form of Mayer-Vietoris sequence) and consequently a natural isomorphism

$$\mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(H) \cong \mathrm{MH}_{*,*}(X) \oplus \mathrm{MH}_{*,*}(G \cap H).$$

Taking the Euler characteristic recovers the inclusion-exclusion formula for magnitude.

1.3.   Diagonality

Leinster [10] noted many examples of graphs which had magnitude with alternating coefficients; these examples included complete graphs, complete bipartite graphs, forests and graphs with up to four vertices. This phenomenon can be explained in terms of magnitude homology. Call a graph G *diagonal* if $\mathrm{MH}_{k,l}(G) = 0$ if $k \neq l$. In this case the formula (1) becomes

$$\#G = \sum_{l \geqslant 0} (-1)^l \cdot \mathrm{rank}\, \mathrm{MH}_{l,l}(G) \cdot q^l,$$

and shows in particular that the coefficients of the magnitude alternate in sign. Recall that the join $G \star H$ of graphs G and H is obtained by adding an edge between every vertex of G and every vertex of H. This is a very drastic operation, for instance the diameter of the resulting join is at most 2. We prove in Theorem 7.5 that any join $G \star H$ of non-empty graphs has diagonal magnitude homology. This tells us immediately that complete graphs and complete bipartite graphs are diagonal. Together with the

Figure 1: The left hand picture shows the magnitude function for a certain 11-point subset of $\mathbb{R}^2$. The right hand picture shows the magnitude function for a certain graph with 11 vertices and 33 edges. The dotted line shows the approximation to the magnitude function using the first seven terms of the expansion in powers of $q = e^{-t}$.

other properties of magnitude homology mentioned above, we recover the alternating magnitude property of all the graphs noted by Leinster, as well as many more.

1.4.   The power series expansion and asymptotics

It is worth commenting here on how the magnitude of graphs fits in with the general theory of magnitude of metric spaces. One nice class of metric spaces, as far as magnitude is concerned, is the class of subsets of Euclidean space (this is a subclass of the class of positive definite metric spaces [15]). For A a finite metric space let tA be A with the metric scaled by a factor of t. If X is a non-empty finite subset of Euclidean space, the *magnitude function* $|tX|$ is defined for all $t > 0$, satisfies $t \geqslant 1$ and is continuous on $(0, \infty)$; see Corollaries 2.4.5 and 2.5.4 of [9] and Corollary 5.5 of [16]. It is not known whether the magnitude function of such a space is increasing or not — see the discussion after Corollary 6.2 of [16] — but certainly all computed examples are increasing. However, it is known [13] that for any finite metric space A and for $t \gg 0$ the magnitude $|tA|$ is defined and increasing in t, with $|tA| \to \mathrm{card}(A)$ as $t \to \infty$. A random and seemingly typical example of a subset of Euclidean space is given in Figure 1.

The metric space obtained from a graph is generally not isometric to a subset of Euclidean space and the magnitude function $|tG|$ will have many singularities, see Figure 1. However, we know from the above that the magnitude will eventually become nice in that the magnitude $|tG|$ is defined and increasing to the number of vertices of G as $t \to \infty$. We defined the magnitude of a graph G to be the formal power series about $q = 0$, which, as $q = e^{-t}$, corresponds to expanding in negative exponentials near $t = \infty$, this avoids the bad behaviour of the magnitude function. Again, see Figure 1.

This perspective can be compared with previous examinations of asymptotics of the magnitude of infinite spaces in [13] and [20]. There polynomial contributions to the aymptotics of the magnitude were shown to come from things like volume, surface area, total scalar curvature and Euler characteristic, which are obtained by integrating local phenomena such as curvature. Here we are looking at exponentially decaying

contributions to the asymptotics and these come from counting global phenomena, such as certain paths of a given length.

1.5. Further directions and open questions

The results described above demonstrate that magnitude homology is natural, nontrivial, and can shed light on properties of the magnitude that are otherwise unexplained. Here are a few questions and avenues for further study.

- There are examples of non-isomorphic graphs with isomorphic magnitude homology, for example any two trees with the same number of vertices. Are there graphs with the same magnitude but different magnitude homology groups?

- Is there a graph whose magnitude homology contains torsion?

- Leinster showed that if two graphs differ by a Whitney twist with adjacent gluing points, then their magnitudes are equal. Do two graphs related by a Whitney twist have isomorphic magnitude homology?

- Prove the magnitude homology of cyclic graphs is as is conjectured in Appendix A.1.

- Computations suggest that the icosahedral graph (i.e. the 1-skeleton of the icosahedron) has diagonal homology. We have not been able to apply any of our techniques for proving that graphs are diagonal in this case, and in particular the graph is not a join. Is the icosahedral graph diagonal, and if so why? More generally, is there a graph-theoretic characterization of diagonal graphs?

- We anticipate that there is a theory of magnitude *cohomology* dual to the homological theory developed in this paper. As with cohomology of spaces, it should be possible to equip this theory with a product structure

- One can define $MH_{*,*}(G)$ as the reduced homology of a sequence of pointed simplicial sets. This is used in Section 8, see in particular Remark 8.6. We have chosen not to emphasise this approach in the present paper, but there may be advantages to doing so in future.

1.6. Organisation of the paper

The paper is organised as follows. Sections 2 and 3 define magnitude homology as a functor and prove that it categorifies the magnitude. Sections 4, 5 and 6 cover the magnitude homology of disjoint unions, cartesian products and unions, describing in detail the categorified properties of the magnitude discussed above. Then section 7 discusses diagonal graphs, in particular the fact that joins are diagonal. Sections 8, 9 and 10 give some lengthy deferred proofs. Finally, Appendix A records and discusses some computer calculations of magnitude homology.

Acknowledgments

2. The definition of magnitude homology

In this section we define the magnitude homology of a graph G, give some very basic examples and properties, and establish the relationship between a graph's magnitude homology and its magnitude. First we state our conventions, which are taken directly from [11]. By a *graph* we mean a finite undirected graph with no loops or multiple edges. The set of vertices of a graph G is denoted $V(G)$ and the set of edges is denoted $E(G)$. If x and y are vertices of a graph G, then the *distance* $d_G(x, y)$ (or simply $d(x, y)$ where it will not cause confusion) is defined to be the length of a shortest edge path from x to y. If x and y lie in different components of G then $d_G(x, y) = \infty$. Thus d_G is a metric on $V(G)$, so long as we allow metrics to take value ∞ on certain pairs.

From this point on we will assume a basic understanding of homological algebra. We recommend section 2.1 of [3] and chapter 12 of [14] as helpful introductions to the subject, and [19] as a standard reference.

Definition 2.1. Let G be a graph. The *length* of a tuple $(x_0, \ldots, x_k)$ of vertices of G is

$$\ell(x_0, \ldots, x_k) = d(x_0, x_1) + \cdots + d(x_{k-1}, x_k).$$

For $i = 0, \ldots, k$ the triangle inequality guarantees that

$$\ell(x_0, \ldots, \widehat{x_i}, \ldots, x_k) \leqslant \ell(x_0, \ldots, x_k) \tag{2}$$

where a 'hat' indicates a term that has been omitted.

Definition 2.2 (The magnitude chain complex)**.** The *magnitude chain complex* $\mathrm{MC}_{*,*}(G)$ of a graph G is the direct sum of chain complexes

$$\bigoplus_{l \geqslant 0} \mathrm{MC}_{*,l}(G)$$

where the chain complex $\mathrm{MC}_{*,l}$ is defined as follows. The group $\mathrm{MC}_{k,l}(G)$ is freely generated by tuples $(x_0, \ldots, x_k)$ of vertices of G satisfying $x_0 \neq x_1 \neq \cdots \neq x_k$ and $\ell(x_0, \ldots, x_k) = l$. The differential

$$\partial \colon \mathrm{MC}_{k,l}(G) \to \mathrm{MC}_{k-1,l}(G)$$

is the alternating sum $\partial = \partial_1 - \partial_2 + \cdots + (-1)^{k-1}\partial_{k-1}$ of the maps defined by

$$\partial_i(x_0, \ldots, x_k) = \begin{cases} (x_0, \ldots, \widehat{x_i}, \ldots, x_k) & \text{if } \ell(x_0, \ldots, \widehat{x_i}, \ldots, x_k) = l, \\ 0 & \text{otherwise.} \end{cases}$$

It is shown in Lemma 2.11 below that $\partial \circ \partial = 0$, so that each $\mathrm{MC}_{*,l}(G)$ is indeed a chain complex.

Remark 2.3. The condition $\ell(x_0, \ldots \widehat{x_i}, \ldots, x_k) = l$ appearing in the definition of the differential can be replaced with the equivalent condition $d(x_{i-1}, x_i) + d(x_i, x_{i+1}) = d(x_{i-1}, x_{i+1})$.

Definition 2.4 (Magnitude homology)**.** The *magnitude homology* $\mathrm{MH}_{*,*}(G)$ of a graph G is the bigraded abelian group defined by

$$\mathrm{MH}_{k,l}(G) = \mathrm{H}_k(\mathrm{MC}_{*,l})$$

for $k, l \geqslant 0$.

Example 2.5 (Complete graphs). Let K_n denote the complete graph on n vertices. Then for $l \geqslant 0$, $\mathrm{MH}_{l,l}(K_n)$ is the free abelian group on $(l+1)$-tuples $(x_0, \ldots, x_l)$ of vertices of K_n satisfying $x_0 \neq \cdots \neq x_l$, and $\mathrm{MH}_{k,l}(K_n) = 0$ if $k \neq l$. This holds because $\mathrm{MC}_{*,*}(K_n)$ admits exactly the same description, as $d(x_i, x_j) = 1$ for $i \neq j$, and in particular its differentials are all zero.

Example 2.6 (Discrete graphs). Let E_n denote the discrete graph on n vertices, meaning that it has no edges. Then $\mathrm{MH}_{0,0}(E_n)$ is the free abelian group on the vertices of E_n, and all other magnitude homology groups of E_n vanish. Again, this follows because the magnitude chain complex admits exactly the same description.

In the two examples above, the magnitude homology was concentrated on the diagonal, by which we mean that $\mathrm{MH}_{k,l}(G) = 0$ for $k \neq l$. In Table 1 we see that according to computer calculations this does not seem to always be the case, and we verify this in the example below.

Example 2.7 (The cyclic graph C_5). Let C_5 denote the cyclic graph with 5 vertices. Then $\mathrm{MH}_{2,3}(C_5)$ is isomorphic to the free abelian group spanned by the oriented edges of C_5. In particular, it is nonzero.

Given an oriented edge (a_1, a_2) of C_5, we will list the vertices of C_5 as a_1, a_2, a_3, a_4, a_5 by starting at a_1, moving to a_2, and then continuing round the graph in the same direction. Then the generators of $\mathrm{MC}_{2,3}(C_5)$ all have one of the four forms

$$(a_1, a_2, a_4), \quad (a_1, a_3, a_4), \quad (a_1, a_2, a_5), \quad (a_1, a_3, a_2)$$

for a uniquely determined oriented edge (a_1, a_2), and they are all cycles. Similarly, the generators of $\mathrm{MC}_{3,3}(C_5)$ all have one of the four forms

$$(a_1, a_2, a_3, a_4), \quad (a_1, a_2, a_3, a_2), \quad (a_1, a_2, a_1, a_2), \quad (a_1, a_2, a_1, a_5)$$

for a uniquely determined oriented edge (a_1, a_2), and their boundaries are as follows.

$$\partial(a_1, a_2, a_3, a_4) = (a_1, a_2, a_4) - (a_1, a_3, a_4)$$
$$\partial(a_1, a_2, a_3, a_2) = (a_1, a_3, a_2)$$
$$\partial(a_1, a_2, a_1, a_2) = 0$$
$$\partial(a_1, a_2, a_1, a_5) = (a_1, a_2, a_5)$$

Thus $\mathrm{MH}_{2,3}(C_5)$ is freely generated by the homology classes of the generators (a_1, a_2, a_4), one for each oriented edge of C_5.

There is an alternative approach to defining magnitude homology that makes use of simplicial sets and filtered objects. We have chosen not to use that approach here in order to make as few technical requirements of the reader as possible, but it is discussed later in the paper. In Section 8 we explain how the magnitude chain complex $\mathrm{MC}_{*,l}(G)$ can be regarded as the normalised, reduced chain complex of a pointed simplicial set $M_l(G)$. In particular, in Remark 8.5 we explain how there is a filtration such that each set $M_l(G)$ arises as a filtration quotient, and in Remark 8.7 we discuss the spectral sequence arising from this filtration.

The next result gives the basic relationship between magnitude homology and magnitude. It is analogous to [6, Proposition 27] and [17, Theorem 13.3], which state that the graded Euler characteristic of the Khovanov homology and Heegaard-Floer

homology of a link are the Jones and Alexander polynomial, respectively. It is our justification for calling magnitude homology a categorification of magnitude.

Theorem 2.8. *Let G be a graph. Then*

$$\sum_{k,l \geqslant 0} (-1)^k \cdot \mathrm{rank}(\mathrm{MH}_{k,l}(G)) \cdot q^l = \#G.$$

Proof. Let χ denote the ordinary Euler characteristic of chain complexes. Then

$$\begin{aligned}
\sum_{k,l \geqslant 0} (-1)^k \cdot \mathrm{rank}(\mathrm{MH}_{k,l}(G)) \cdot q^l &= \sum_{l \geqslant 0} \chi(\mathrm{MH}_{*,l}(G)) \cdot q^l \\
&= \sum_{l \geqslant 0} \chi(\mathrm{MC}_{*,l}(G)) \cdot q^l \\
&= \sum_{k,l \geqslant 0} (-1)^k \cdot \mathrm{rank}(\mathrm{MC}_{k,l}(G)) \cdot q^l \\
&= \sum_{k \geqslant 0} (-1)^k \sum_{x_0 \neq \cdots \neq x_k} q^{d(x_0,x_1)+\cdots+d(x_{k-1},x_k)} \\
&= \#G.
\end{aligned}$$

Here the first and third inequalities are the definition of Euler characteristic of a graded abelian group, the second is a standard property of the Euler characteristic, and the fourth follows by counting the generators of $\mathrm{MC}_{k,l}(G)$. The final equality now follows by [**11**, Proposition 3.9]. $\square$

We now see some further basic properties of magnitude, these are illustrated in Table 1. The first proposition explains that the top two entries are the number of vertices and twice the number of edges.

Proposition 2.9. *Let G be a graph. Then $\mathrm{MH}_{0,0}(G)$ is the free abelian group on the vertices of G and $\mathrm{MH}_{1,1}(G)$ is the free abelian group on the oriented edges of G.*

Proof. The same properties hold trivially for chains, and all differentials involving the terms $\mathrm{MC}_{0,0}(G)$ and $\mathrm{MC}_{1,1}(G)$ are zero (having zero domain or range), so the properties hold for homology. $\square$

The next proposition explains why the table is lower triangular and why the diameter of the 5-cycle being 2 restricts the non-zero entries to being reasonably close to the diagonal.

Proposition 2.10. *Let G be a graph and suppose that $\mathrm{MH}_{k,l}(G) \neq 0$. Then*

$$k \leqslant l.$$

Furthermore, if G has diameter d, then

$$\frac{l}{d} \leqslant k$$

and moreover

$$\frac{l}{d} < k$$

if $d > 1$ and $l > 0$.

Proof. If $\mathrm{MH}_{k,l}(G) \neq 0$ then $\mathrm{MC}_{k,l}(G) \neq 0$, so there is a tuple $(x_0, \ldots, x_k)$ satisfying

$$l = d(x_0, x_1) + \cdots d(x_{k-1}, x_k).$$

Each of the summands is at least 1, since consecutive entries are distinct, and this gives the first inequality. If G has diameter d then each summand is at most d, and this gives the second inequality. For the final part suppose that $d > 1$, $l > 0$ and $k = l/d$. Let $(x_0, \ldots, x_k)$ be a generator of $\mathrm{MC}_{k,l}(G)$, so that $d(x_{i-1}, x_i) = d$ for all i. Since $d(x_0, x_1) \geqslant 2$ there is $y \neq x_0, x_1$ such that $d(x_0, y) + d(y, x_1) = d(x_0, x_1)$. Then $\partial_1(x_0, y, x_1, \ldots, x_k) = (x_0, x_1, \ldots, x_k)$ while $\partial_i(x_0, y, x_1, \ldots, x_k) = 0$ for $i = 2, \ldots, k - 1$. Thus $(x_0, \ldots, x_k) = \partial(-(x_0, y, x_1, \ldots, x_k))$. It follows that $\mathrm{MH}_{k,l}(G) = 0$. □

Let us conclude the section by verifying that the operators ∂ satisfy the relation $\partial \circ \partial = 0$. This is a routine consequence of inequality (2), but because similar arguments will appear several times in the rest of the paper, we give a detailed proof here.

Lemma 2.11. *For any graph G, any $k \geqslant 2$, and any $l \geqslant 0$, the composite*

$$\mathrm{MC}_{k,l}(G) \xrightarrow{\partial} \mathrm{MC}_{k-1,l}(G) \xrightarrow{\partial} \mathrm{MC}_{k-2,l}(G)$$

vanishes.

Proof. It is sufficient to show that for any generator $(x_0, \ldots, x_k)$ of $\mathrm{MC}_{k,l}(G)$, and any i, j in the range $0 \leqslant i < j \leqslant k$, we have

$$\partial_i \circ \partial_j(x_0, \ldots, x_k) = \partial_{j-1} \circ \partial_i(x_0, \ldots, x_k).$$

Each side of this equation is given by either $(x_0, \ldots, \widehat{x_i}, \ldots, \widehat{x_j}, \ldots, x_k)$ or 0. The left hand side is nonzero if and only if

$$\ell(x_0, \ldots, \widehat{x_i}, \ldots, x_k) = l \quad \text{and} \quad \ell(x_0, \ldots, \widehat{x_i}, \ldots, \widehat{x_j}, \ldots, x_k) = l,$$

and inequality (2) tells us that

$$\ell(x_0, \ldots, \widehat{x_i}, \ldots, \widehat{x_j}, \ldots, x_k) \leqslant \ell(x_0, \ldots, \widehat{x_i}, \ldots, x_k) \leqslant \ell(x_0, \ldots, x_k) = l,$$

so that the left hand side is nonzero if and only if

$$\ell(x_0, \ldots, \widehat{x_i}, \ldots, \widehat{x_j}, \ldots, x_k) = l.$$

A similar argument shows that the right hand side is nonzero if and only if the same condition holds. This completes the proof. □

3. Induced maps

The magnitude of a graph is an element of a set, the set of formal power series with integer coefficients. The magnitude homology of a graph, on the other hand, is an object of a category, the category of bigraded abelian groups. This categorification gives us the opportunity to make magnitude into a functor, and that is what we will do in this section.

In order to make graphs into the objects of a category we choose the following notion of morphism. Given graphs G and H, a *map of graphs* $f \colon G \to H$ is a map of

vertex sets $f\colon V(G) \to V(H)$ satisfying the condition

$$\{x, y\} \in E(G) \implies \{f(x), f(y)\} \in E(H) \text{ or } f(x) = f(y).$$

In words, a map of graphs is a map of vertex sets that preserves or contracts each edge. And in terms of distance, a map of graphs is a map of vertex sets for which $d_H(f(x), f(y)) \leqslant d_G(x, y)$ for all vertices $x, y \in V(G)$. From the metric space perspective these distance non-increasing maps are the correct ones to consider in the context of magnitude. Observe that if $f\colon G \to H$ is a map of graphs, then the inequality

$$\ell(f(x_0), \ldots, f(x_k)) \leqslant \ell(x_0, \ldots, x_k) \tag{3}$$

holds for any tuple $(x_0, \ldots, x_k)$ of vertices of G.

Definition 3.1 (Induced chain maps). Let $f\colon G \to H$ be a map of graphs. The *induced chain map*

$$f_\#\colon \mathrm{MC}_{*,*}(G) \longrightarrow \mathrm{MC}_{*,*}(H)$$

is defined on generators by

$$f_\#(x_0, \ldots, x_k) = \begin{cases} (f(x_0), \ldots, f(x_k)) & \text{if } \ell(f(x_0), \ldots, f(x_k)) = \ell(x_0, \ldots, x_k) \\ 0 & \text{otherwise.} \end{cases}$$

If $f\colon G \to H$ is a map of graphs then the relation $f_\# \circ \partial = \partial \circ f_\#$ holds, so that $f_\#$ is indeed a chain map. The proof, which we omit, is similar to that of Lemma 2.11, and makes use of inequalities (2) and (3).

Definition 3.2 (Induced maps in homology). Let $f\colon G \to H$ be a map of graphs. The *induced map in homology* is the map

$$f_*\colon \mathrm{MH}_{*,*}(G) \longrightarrow \mathrm{MH}_{*,*}(H)$$

induced by $f_\#$.

Proposition 3.3. *The assignment* $G \mapsto \mathrm{MH}_{*,*}(G)$, $f \mapsto f_*$ *is a functor from the category of graphs to the category of bigraded abelian groups.*

That the identity map of a graph induces the identity map in homology is immediate. To prove that $g_* \circ f_* = (g \circ f)_*$ holds for any maps of graphs $f\colon G \to H$ and $g\colon H \to K$ one proceeds as in Lemma 2.11, making use of both (3) and (2). The details are left to the reader.

Recall from Proposition 2.9 that $\mathrm{MH}_{0,0}(G)$ is the free abelian group on the set of vertices of G, and that $\mathrm{MH}_{1,1}(G)$ is the free abelian group on the set of oriented edges of G. The following result, whose proof is an immediate consequence of the definitions, describes the effect of induced maps in these degrees.

Proposition 3.4. *Let* $f\colon G \to H$ *be a map of graphs. Then* $f_*\colon \mathrm{MH}_{0,0}(G) \to \mathrm{MH}_{0,0}(H)$ *sends a vertex* x *to* $f(x)$. *And* $f_*\colon \mathrm{MH}_{1,1}(G) \to \mathrm{MH}_{1,1}(H)$ *sends an edge* $\{x, y\}$ *to* $\{f(x), f(y)\}$ *if that is an edge, and to 0 otherwise.*

Corollary 3.5. *Let* $f\colon G \to H$ *be a map of graphs. If* $f_*\colon \mathrm{MH}_{*,*}(G) \to \mathrm{MH}_{*,*}(H)$ *is an isomorphism, then* f *is an isomorphism of graphs.*

4. Disjoint unions

In this brief section we prove the additivity of magnitude homology with respect to disjoint unions. As an immediate corollary we get the additivity of the magnitude.

Proposition 4.1. *Let G and H be graphs and write $i\colon G \to G \sqcup H$ and $j\colon H \to G \sqcup H$ for the inclusion maps. Then the induced map*

$$i_* \oplus j_*\colon \mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(H) \longrightarrow \mathrm{MH}_{*,*}(G \sqcup H)$$

is an isomorphism.

Proof. Let $(x_0, \ldots, x_k)$ be a generator of $\mathrm{MC}_{k,l}(G \sqcup H)$. Since $\ell(x_0, \ldots, x_k) = l$, we have $d(x_{i-1}, x_i) < \infty$ for all i, so that $x_0, \ldots, x_k$ all belong to G or all belong to H. Consequently $i_\# \oplus j_\#$ is an isomorphism, and the result follows. $\square$

Corollary 4.2 (Leinster [**11**, Lemma 3.5]). *Let G and H be graphs. Then $\#(G \sqcup H) = \#G + \#H$.*

Proof. This follows from Proposition 4.1, Theorem 2.8, and the fact that $\chi(C_* \oplus D_*) = \chi(C_*) + \chi(D_*)$ for finitely generated graded abelian groups C_* and D_*. $\square$

5. Cartesian products

In this section we state a Künneth Theorem for magnitude homology with respect to the cartesian product of graphs and we give an example. The proof of the theorem is given is Section 8.

The *cartesian product* $G \square H$ of graphs G and H has vertex set $V(G) \times V(H)$, and has an edge from (x_1, y_1) to (x_2, y_2) if either $x_1 = x_2$ and $\{y_1, y_2\}$ is an edge in H, or $y_1 = y_2$ and $\{x_1, x_2\}$ is an edge in G. The metric on $G \square H$ is given by

$$d_{G \square H}((x_1, y_1), (x_2, y_2)) = d_G(x_1, x_2) + d_H(y_1, y_2) \tag{4}$$

for $(x_1, y_1), (x_2, y_2) \in V(G \square H)$.

Remark 5.1. The cartesian product is not the *categorical* product on the category of graphs, but Equation (4) tells us that it is the natural tensor product from the perspective of enriched category theory. Leinster's definition of the magnitude of a metric space [**9**] was motivated by Lawvere's observation [**7**] that metric spaces can be viewed as categories enriched over the poset of extended non-negative real numbers $[0, \infty]$ equipped with addition $+$ as the monoidal product. Viewing graphs as enriched categories in this way means that via Equation (4) we can see that $G \square H$ is exactly the tensor product of G and H [**9**, Section 1.4].

Definition 5.2 (The exterior product). Fix $l \geqslant 0$. The *exterior product* is the map

$$\square\colon \mathrm{MC}_{*,*}(G) \otimes \mathrm{MC}_{*,*}(H) \longrightarrow \mathrm{MC}_{*,*}(G \square H) \tag{5}$$

whose component

$$\square\colon \mathrm{MC}_{k_1,l_1}(G) \otimes \mathrm{MC}_{k_2,l_2}(H) \longrightarrow \mathrm{MC}_{k,l}(G \square H)$$

for $k_1, k_2 \geqslant 0$ with $k_1 + k_2 = k$ is defined by

$$(x_0, \ldots, x_{k_1}) \square (y_0, \ldots, y_{k_2}) = \sum_\sigma \mathrm{sign}(\sigma) \cdot ((x_{i_0}, y_{j_0}), \ldots, (x_{i_k}, y_{j_k})).$$

Here the sum ranges over all sequences $\sigma = ((i_0, j_0), \ldots, (i_k, j_k))$ for which $i_0 = j_0 = 0$, for which $0 \leqslant i_r \leqslant k_1$ and $0 \leqslant j_r \leqslant k_2$ for all r, and for which each term (i_{r+1}, j_{r+1}) is obtained from its predecessor (i_r, j_r) by increasing exactly one of the two components by 1. Given such a sequence, we define $\text{sign}(\sigma) = (-1)^n$ where n is the number of pairs (i, j) for which $i = i_r \implies j < j_r$. Compare with the discussion in [**3**, pp. 277-278]. The exterior product is a chain map, and so induces a map in homology that we indicate by the same symbol,

$$\square \colon \text{MH}_{*,*}(G) \otimes \text{MH}_{*,*}(H) \longrightarrow \text{MH}_{*,*}(G \square H). \tag{6}$$

Theorem 5.3 (The Künneth theorem for magnitude homology). *The exterior product in homology* (6) *fits into a natural short exact sequence*

$$0 \longrightarrow \text{MH}_{*,*}(G) \otimes \text{MH}_{*,*}(H) \xrightarrow{\square} \text{MH}_{*,*}(G \square H)$$
$$\longrightarrow \text{Tor}(\text{MH}_{*-1,*}(G), \text{MH}_{*,*}(H) \longrightarrow 0$$

that is non-naturally split. In particular, $\square$ becomes an isomorphism after tensoring with the rationals, and is an isomorphism if either $\text{MH}_{*,*}(G)$ or $\text{MH}_{*,*}(H)$ is torsion-free.

The proof of this theorem is deferred to Section 8.

Example 5.4 (The cyclic graph C_4). The magnitude homology of the cyclic graph C_4 with four vertices is

$$\text{MH}_{k,l}(C_4) = \begin{cases} \mathbb{Z}^{4(l+1)} & \text{if } k = l, \\ 0 & \text{otherwise.} \end{cases}$$

To see this, observe that $C_4 = K_2 \square K_2$. Example 2.5 shows that $MH_{k,l}(K_2)$ vanishes if $k \neq l$ and that it is free abelian on two generators if $k = l$. Since these groups contain no torsion the Künneth Theorem (Theorem 5.3) shows that

$$\text{MH}_{*,*}(C_4) \cong \text{MH}_{*,*}(K_2) \otimes \text{MH}_{*,*}(K_2),$$

or more explicitly that

$$\text{MH}_{k,l}(C_4) \cong \bigoplus_{\substack{k_1+k_2=k \\ l_1+l_2=l}} \text{MH}_{k_1,l_1}(K_2) \otimes \text{MH}_{k_2,l_2}(K_2),$$

and the claim follows.

Remark 5.5. We know of no graph G for which $\text{MH}_{*,*}(G)$ contains torsion.

6. The Mayer-Vietoris sequence

In this section we show that a Mayer-Vietoris theorem holds for so-called projecting decompositions of graphs. The long exact Mayer-Vietoris sequence actually breaks up into split short exact sequences. From this we will obtain Leinster's inclusion-exclusion principle. The proof is given in Section 9.

We begin by recalling some definitions of Leinster. Firstly, convexity is supposed to be reminiscent of the idea that in convex subset of $\mathbb{R}^n$ each pair of points is connected by a geodesic which is also contained in the subset.

Definition 6.1 (Convex [**11**, Definition 4.2]). A subgraph $U \subset X$ is called *convex* if $d_U(u, v) = d_X(u, v)$ for all $u, v \in U$.

Secondly, projecting to a convex subgraph is reminiscent of the idea that there is a 'nearest point map' from $\mathbb{R}^n$ to any convex subset.

Definition 6.2 (Projecting [**11**, Definition 4.6]). Let $U \subset X$ be a convex subgraph. We say that X *projects* to U if for every $x \in X$ that can be connected by an edge-path to some vertex of U, there is $\pi(x) \in U$ such that for all $u \in U$ we have

$$d(x, u) = d(x, \pi(x)) + d(\pi(x), u).$$

Thus $\pi(x)$ is the unique point of U closest to x. Writing X_U for the union of those components of X whose vertices admit an edge path to U, there is then a map $\pi \colon X_U \to U$ defined by $u \mapsto \pi(u)$.

Every even cyclic graph projects to any of its edges, whereas no odd graph projects to any of its edges. Projecting to U is stronger than each point having a closest point in U as can be seen by considering two adjacent edges as a subgraph of the 5-cycle graph.

Suppose that X is a graph that is the union of subgraphs G and H, such that $G \cap H$ is convex in $G \cup H$, and such that H projects to $G \cap H$. Leinster [**11**, Theorem 4.9] has shown that in this situation the magnitude satisfies the inclusion-exclusion principle $\#X = \#G + \#H - \#(G \cap H)$. We will categorify this to a Mayer-Vietoris sequence relating the magnitude homologies of $G \cap H$, G, H and $G \cup H$.

Definition 6.3 (Projecting decompositions). A *projecting decomposition* is a triple $(X; G, H)$ consisting of a graph X and subgraphs G and H such that the following properties hold.

- $X = G \cup H$
- $G \cap H$ is convex in X
- H projects to $G \cap H$

Given a projecting decomposition $(X; G, H)$, we write

$$i^G \colon G \to X, \quad i^H \colon H \to X, \quad j^G \colon G \cap H \to G, \quad j^H \colon G \cap H \to H$$

for the inclusion maps. A *decomposition map* $f \colon (X; G, H) \to (X'; G', H')$ is a map of graphs $f \colon X \to X'$ such that $f(G) \subset G'$ and $f(H) \subset H'$. It is a *projecting* decomposition map if $H_{G \cap H} = f^{-1}(H'_{G' \cap H'})$ and $f(\pi(h)) = \pi(f(h))$ for all $h \in H_{G \cap H}$.

Definition 6.4. Given a projecting decomposition $(X; G, H)$, let $\mathrm{MC}_{*,*}(G, H)$ denote the subcomplex of $\mathrm{MC}_{*,*}(G \cup H)$ spanned by those tuples $(x_0, \dots, x_k)$ whose entries all lie in G or all lie in H.

Theorem 6.5 (Excision for magnitude chains). *Let $(X; G, H)$ be a projecting decomposition. For all $l \geqslant 0$ the inclusion $\mathrm{MC}_{*,l}(G, H) \hookrightarrow \mathrm{MC}_{*,l}(G \cup H)$ is a quasi-isomorphism.*

This result is a version of excision for the magnitude chain complex, and is closely analogous to versions of excision that hold for the singular chain complex of a topological space, see for example [**3**, Proposition 2.21] or [**1**, Proposition 7.3]. The proof of Theorem 6.5 is deferred until Section 9. From excision we get the Mayer-Vietoris Theorem.

Theorem 6.6 (Mayer-Vietoris for magnitude homology)**.** *Let $(X; G, H)$ be a projecting decomposition. Then there is a split short exact sequence:*

$$0 \to \mathrm{MH}_{*,*}(G \cap H) \xrightarrow{(j_*^G, -j_*^H)} \mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(H)$$

$$\xrightarrow{i_*^G \oplus i_*^H} \mathrm{MH}_{*,*}(G \cup H) \to 0.$$

The sequence is natural with respect to decomposition maps, and the splitting is natural with respect to projecting decomposition maps.

Theorem 6.6 is stronger than one might anticipate, since it gives a short exact sequence in each homological degree, rather than the single long exact sequence that is familiar from the Mayer-Vietoris theorem for singular homology [**3**, p. 149]. In fact, our short exact sequence is obtained from a long exact sequence of Mayer-Vietoris type by splitting it into the short exact sequences. The splitting is possible due to our assumption that in a projecting decomposition $(X; G, H)$ the subgraph H projects onto $G \cap H$. This assumption is impossible to remove, as shown in Section A.2. The proof of Theorem 6.6 is also deferred until Section 9.

Corollary 6.7 (Inclusion-exclusion [**11**, Theorem 4.9])**.** *If $(X; G, H)$ is a projecting decomposition then $\#X = \#G + \#H - \#(G \cap H)$.*

Proof. From the short exact sequence in Theorem 6.6 it follows that

$$\chi(\mathrm{MH}_{*,l}(G \cap H)) - \chi(\mathrm{MH}_{*,l}(G) \oplus \mathrm{MH}_{*,l}(H)) + \chi(\mathrm{MH}_{*,l}(G \cup H)) = 0.$$

Since Euler characteristic is additive with respect to direct sums, this rearranges to give

$$\chi(\mathrm{MH}_{*,l}(G \cup H)) = \chi(\mathrm{MH}_{*,l}(G)) + \chi(\mathrm{MH}_{*,l}(H)) - \chi(\mathrm{MH}_{*,l}(G \cap H)).$$

Multiplying this equation by q^l, then summing over all $l \geqslant 0$ and applying Theorem 2.8, the claim follows. □

Corollary 6.8 (Magnitude homology of trees)**.** *Let T be a tree. Then*

$$\mathrm{MH}_{k,l}(T) \cong \begin{cases} \mathbb{Z}V(T) & \text{if } k = l = 0, \\ \mathbb{Z}\vec{E}(T) & \text{if } k = l > 0, \\ 0 & \text{if } k \neq l. \end{cases}$$

This isomorphism is natural with respect to maps of trees, where $\mathbb{Z}\vec{E}(T)$ is made into a functor of T by declaring that if $f\colon T \to S$ is a map, then $f_*\colon \vec{E}(T) \to \vec{E}(S)$ sends an oriented edge (x, y) to $(f(x), f(y))$ if $f(x) \neq f(y)$, and to 0 if $f(x) = f(y)$.

Proof. Let us write $F_{k,l}$ for the functor appearing on the right hand side of the desired isomorphism. There is a natural transformation $\theta\colon F_{k,l} \Rightarrow \mathrm{MH}_{k,l}$ given on generators by $\theta_T(x) = (x)$ if $k = l = 0$ and $x \in V(T)$, and by

$$\theta_T((x, y)) = (\underbrace{x, y, x, y, \dots}_{k+1})$$

if $k = l > 0$ and $(x, y) \in \vec{E}(T)$. We prove that θ_T is an isomorphism by induction on the number of edges of T. Observe that θ_T is trivially an isomorphism if T has no

edges or a single edge. In general, if T has two or more edges then we may write $T = T_1 \cup T_2$ where T_1, T_2 and $T_1 \cap T_2$ are subtrees of T. It is immediate that $T_1 \cap T_2$ is convex in T and that T_2 projects to $T_1 \cap T_2$, so that by Theorem 6.6 we have a short exact sequence

$$0 \to \mathrm{MH}_{k,l}(T_1 \cap T_2) \to \mathrm{MH}_{k,l}(T_1) \oplus \mathrm{MH}_{k,l}(T_2) \to \mathrm{MH}_{k,l}(T) \to 0.$$

There is an analogous short exact sequence in which $\mathrm{MH}_{k,l}$ is replaced with $F_{k,l}$, and it can be combined with the one above to form the rows of a commuting diagram whose vertical arrows are obtained using θ. The first two vertical arrows, $\theta_{T_1 \cap T_2}$ and $\theta_{T_1} \oplus \theta_{T_2}$, are isomorphisms by induction, and it follows that θ_T is an isomorphism as well. $\qquad\square$

Corollary 6.9 (Wedge sums). *Let G and H be graphs with chosen base vertices, and let $G \vee H$ denote the graph obtained by identifying the two base vertices to a single vertex P. Then the inclusion maps $a \colon G \to G \vee H$ and $b \colon H \to G \vee H$ induce an isomorphism*

$$a_* \oplus b_* \colon \mathrm{MH}_{*,*}(G) \oplus_{\mathrm{MH}_{*,*}(P)} \mathrm{MH}_{*,*}(H) \xrightarrow{\cong} \mathrm{MH}_{*,*}(G \vee H).$$

Proof. By considering $G \vee H$ as the union of G and H, and observing that H projects onto P, we obtain a projecting decomposition $(G \vee H; G, H)$. The result then follows from Theorem 6.6. $\qquad\square$

7. Diagonal graphs

We have seen that complete graphs (Example 2.5), discrete graphs (Example 2.6), and trees (Corollary 6.8) are all diagonal in the following sense.

Definition 7.1 (Diagonality). A graph is *diagonal* if its magnitude homology is concentrated on the diagonal. In other words G is diagonal if $\mathrm{MH}_{k,l}(G) = 0$ for $k \neq l$.

We have also seen that the five-cycle C_5 is not diagonal (Example 2.7). In this section we will give some rather general results that demonstrate diagonality in various situations. These will be enough for us to explain in general terms all the instances of diagonality that we have seen so far, and also several more, including complete multipartite graphs and the octahedral graph. The magnitude of a diagonal graph has coefficients that alternate in sign; our examples of diagonal graphs explain all instances of this phenomenon that are known so far.

Proposition 7.2 (Diagonal graphs and magnitude). *If G is diagonal then the coefficients of the magnitude $\#G$ alternate in sign, and $\#G$ determines $\mathrm{MH}_{*,*}(G)$ up to isomorphism.*

Proof. Let G be diagonal. By Theorem 2.8 we have

$$\#G = \sum_{l \geqslant 0} (-1)^l \cdot \mathrm{rank}(\mathrm{MH}_{l,l}(G)) \cdot q^l$$

so that the coefficients alternate in sign as claimed, and $\#G$ determines the quantities $\mathrm{rank}(\mathrm{MH}_{l,l}(G))$. Since the chain groups $\mathrm{MC}_{l+1,l}(G)$ are identically 0, it follows that the $\mathrm{MH}_{l,l}(G)$ are free abelian, and so determined by their ranks. This completes the proof. $\qquad\square$

Proposition 7.3. *The cartesian product of diagonal graphs is diagonal. A graph that admits a projecting decomposition into diagonal graphs is itself diagonal.*

Proof. The first claim follows from Theorem 5.3, together with the fact that if G is diagonal then each group $\mathrm{MH}_{l,l}(G)$ is torsion-free (see the proof of Proposition 7.2). The second follows from Theorem 6.6. $\qquad\square$

Definition 7.4. Let G and H be graphs. The *join* of G and H, denoted $G \star H$, is the graph obtained from $G \sqcup H$ by adding the edges $\{x, y\}$ for all $x \in V(G)$ and $y \in V(H)$.

Theorem 7.5. *Let G and H be graphs such that $G, H \neq \emptyset$. Then the join $G \star H$ is diagonal.*

The proof, which is rather lengthy, is deferred until Section 10.

Example 7.6 (Complete multipartite graphs). The complete multipartite graph with maximal independent subsets of size $n_1, \ldots, n_k$ is the iterated join $E_{n_1} \star \cdots \star E_{n_k}$, and so is diagonal by Theorem 7.5. This also gives another proof that complete graphs are diagonal, since they are iterated joins of copies of E_1, and that C_4 is diagonal, since it is $E_2 \star E_2$.

Example 7.7 (1-skeleta of platonic solids). The 1-skeleta of the tetrahedron, cube and octahedron are all diagonal, since they are K_4, $K_2 \square K_2 \square K_2$ and $E_2 \star E_2 \star E_2$ respectively. The 1-skeleton of the dodecahedron is not diagonal: its magnitude homology in bidegree $(2, 3)$ is nonzero, as one sees by adapting the reasoning of Example 2.7. On the other hand it appears from Sage computations (see Appendix A.4) that the 1-skeleton of the icosahedron is diagonal, though we cannot prove it using the techniques of this section.

8. Proof of the Künneth Theorem

We now give the proof of Theorem 5.3. While the proofs in the previous sections were complicated but not strictly speaking technical, the present proof is indeed technical, relying on the version of the Künneth theorem that applies to the homology of simplicial sets.

Definition 8.1 (The simplicial set $M_l(G)$). Let G be a graph and let $l \geqslant 0$. We define $M_l(G)$ to be the pointed simplicial set whose k-simplices are the $(k + 1)$-tuples $(x_0, \ldots, x_k)$ of length l, plus a basepoint simplex. (Adjacent entries are allowed to be equal.) The i-th face map deletes the i-th entry of a tuple if this preserves the length, and sends it to the basepoint otherwise. The i-th degeneracy doubles the i-th entry of a tuple. The faces and degeneracies all send basepoints to basepoints.

Observe that the non-degenerate, non-basepoint k-simplices of $M_l(G)$ are precisely the generators of $\mathrm{MC}_{k,l}(G)$.

Proposition 8.2 (A simplicial Künneth theorem)**.** *Let G and H be graphs and fix $l \geqslant 0$. Then the map of pointed simplicial sets*

$$\square \colon \bigvee_{l_1+l_2=l} M_{l_1}(G) \wedge M_{l_2}(H) \longrightarrow M_l(G \square H)$$

defined by

$$(x_0, \ldots, x_k) \square (y_0, \ldots, y_k) = ((x_0, y_0), \ldots, (x_k, y_k))$$

is an isomorphism.

Proof. That the map is simplicial and an isomorphism both follow from the observation that

$$\ell((x_0, y_0), \ldots, (x_k, y_k)) = \ell(x_0, \ldots, x_k) + \ell(y_0, \ldots, y_k)$$

for any tuple $((x_0, y_0), \ldots, (x_k, y_k))$ of vertices of $G \square H$. $\square$

Given a pointed simplicial set X, we write $\bar{N}_*(X)$ for the *normalised reduced chain complex* of X. This is given in degree k by the free abelian group on X_k, divided out by the span of the degenerate simplices and the basepoint. The differential $d \colon \bar{N}_k(X) \to \bar{N}_{k-1}(X)$ is given by $d = \sum_{i=0}^{k}(-1)^i d_i$, where d_i denotes the i-th face map (extended linearly). The following is immediate from the definitions.

Lemma 8.3. $\mathrm{MC}_{*,l}(G) = \bar{N}_*(M_l(G))$.

Given pointed simplicial sets X and Y, we define the *reduced normalised Eilenberg-Zilber map*

$$\nabla^{\bar{N}} \colon \bar{N}_*(X) \otimes \bar{N}_*(Y) \longrightarrow \bar{N}_*(X \wedge Y)$$

by

$$\nabla^{\bar{N}}(x \otimes y) = \sum_{\sigma} \mathrm{sign}(\sigma)(c \circ (x \times y) \circ \sigma)$$

for $x \in X_p$ and $y \in Y_q$ non-degenerate, non-basepoint. Here $c \colon X \times Y \to X \wedge Y$ denotes the collapse map while σ and $\mathrm{sign}(\sigma)$ are like those in Definition 5.2 except we are regarding x, y and σ as simplicial maps $x \colon \Delta[p] \to X$, $y \colon \Delta[q] \to Y$ and $\sigma \colon \Delta[p+q] \to \Delta[p] \times \Delta[q]$, so that $c \circ (x \times y) \circ \sigma$ is a simplicial map $\Delta[p+q] \to X \wedge Y$, or in other words an element of $(X \wedge Y)_{p+q}$. The following fact is presumably standard, but we do not know of a proof that applies to *reduced* normalised chains.

Proposition 8.4. *The Eilenberg-Zilber map $\nabla^{\bar{N}}$ is a quasi-isomorphism.*

Proof. Given a simplicial set Z, let us write $N_*(Z)$ for the normalised chains on Z, or in other words the standard chains on Z divided out by the span of the degenerate elements. See section 4 of [**2**]. Let U and V be simplicial sets. As in section 5 of [**2**], the standard Eilenberg-Zilber map reduces to a map

$$\nabla^N \colon N_*(U) \otimes N_*(V) \longrightarrow N_*(U \times V)$$

that is a chain homotopy equivalence. The definition of the Eilenberg-Zilber map is

given in line (5.3) of [2], and it is simple to use this to verify that

$$\nabla^N(u \otimes v) = \sum_{\sigma} \text{sign}(\sigma)(u \times v) \circ \sigma$$

for $u \in U_p$ and $v \in V_q$ non-degenerate, with the right-hand-side understood as in the definition of $\nabla^{\bar{N}}$. One sees that the diagram below commutes.

$$
\begin{array}{ccc}
N_*(X) \otimes N_*(Y) & \xrightarrow{\nabla^N} & N_*(X \times Y) \\
\downarrow & & \downarrow \\
\bar{N}_*(X) \otimes \bar{N}_*(Y) & \xrightarrow{\nabla^{\bar{N}}} & \bar{N}_*(X \wedge Y)
\end{array}
$$

The kernels of the vertical maps are $N_*(X) \otimes N_0(\text{pt}) + N_0(\text{pt}) \otimes N_*(Y)$ and $N_*(X \vee Y)$ respectively, and it is evident from the formula that ∇^N restricts to an isomorphism between these. Since ∇^N is a quasi-isomorphism, it follows that $\nabla^{\bar{N}}$ is as well. $\square$

Proof of Theorem 5.3. The composite

$$\bigoplus_{l_1+l_2=l} \text{MC}_{*,l_1}(G) \otimes \text{MC}_{*,l_2}(H) \xrightarrow{=} \bigoplus_{l_1+l_2=l} \bar{N}(M_{l_1}(G)) \otimes \bar{N}(M_{l_2}(H))$$

$$\xrightarrow{\oplus \nabla^{\bar{N}}} \bigoplus_{l_1+l_2=l} \bar{N}(M_{l_1}(G) \wedge M_{l_2}(H))$$

$$\xrightarrow{=} \bar{N}\left(\bigvee_{l_1+l+2=l} M_{l_1}(G) \wedge M_{l_2}(H) \right)$$

$$\xrightarrow{\bar{N}(\square)} \bar{N}(M_l(G \,\square\, H))$$

$$\xrightarrow{=} \text{MC}_{*,l}(G \,\square\, H)$$

consists of isomorphisms and one quasi-isomorphism, and so is itself a quasi-isomorphism. Unravelling the definitions shows that this composite is precisely the map $\square$. The result then follows by applying the Algebraic Künneth Theorem [3, Theorem 3B.5]. $\square$

Remark 8.5 ($M_l(G)$ as filtration quotients). We may realise each $M_l(G)$ as a filtration quotient of a filtered simplicial set, as follows. Define $\text{MS}(G)$ to be the simplicial set whose k-simplices are finite-length $(k+1)$-tuples $(x_0, \ldots, x_k)$ of vertices of G, in which the i-th face map deletes the i-th entry, and in which the i-th degeneracy doubles the i-th entry. Form the filtration

$$\text{MS}_0(G) \subset \text{MS}_1(G) \subset \text{MS}_2(G) \subset \cdots$$

of $\text{MS}(G)$ in which $\text{MS}_l(G)$ consists of all tuples of length at most l. Then $M_l(G) = \text{MS}_l(G)/\text{MS}_{l-1}(G)$. The simplicial set $\text{MS}(G)$ has one component for each component of G, and it is not difficult to show that each component is contractible.

Remark 8.6 (The simplicial approach to magnitude homology). Readers with the relevant background in abstract homotopy theory may find it more natural to think about magnitude homology using the pointed simplicial sets $M_l(G)$ introduced in this section, and indeed using the filtered simplicial set $\text{MS}(G)$ of the previous remark,

rather than using the definition of $\mathrm{MC}_{*,*}(G)$. We have chosen to downplay this simplicial approach in order to make the paper as accessible as possible, in particular to readers coming from graph theory and category theory. We expect that large parts of our work could be 'lifted' to the context of simplicial sets, however it is not clear that this would lead to any significant simplifications in the material covered here. Moreover, considering how little we know of magnitude homology (for example, we do not know any graphs whose magnitude homology contains torsion), it seems reasonable to limit ourselves to a homological approach at this stage.

Remark 8.7 (A spectral sequence). The filtered simplicial set $\mathrm{MS}(G)$ of the previous two remarks gives rise to a spectral sequence $(E^r_{*,*})_{r \geqslant 1}$ whose E^1-page is obtained from the homology of the filtration quotients of $\mathrm{MS}(G)$, and which converges to the homology of $\mathrm{MS}(G)$. To be precise, $E^1_{i,j} = \mathrm{MH}_{i+j,i}(G)$, $E^\infty_{i,j} = 0$ for $(i,j) \neq (0,0)$, and $E^\infty_{0,0} = \mathbb{Z}^c$ where c denotes the number of components of G.

9.   Proof of excision and Mayer-Vietoris

In this section we give the proofs of Theorem 6.5 (excision for magnitude chains) and Theorem 6.6 (Mayer-Vietoris for magnitude homology). To that end we fix throughout the section a projecting decomposition $(X; G, H)$. In this situation the pairs

$$G \cap H \subset X \qquad G \cap H \subset G \qquad G \cap H \subset H \qquad G \subset X \qquad H \subset X$$

are all convex. In the first case this is an assumption. In the second and third cases it is an immediate consequence of the first. And in the fourth and fifth cases it follows from [**11**, Lemma 4.3]. Thus the length of a tuple $(x_0, \ldots, x_k)$ of vertices of X is unambiguously defined: even if the vertices happen to all lie in G or H or $G \cap H$, the length is the same whichever graph one regards the tuple as belonging to.

Proof of Theorem 6.6, assuming Theorem 6.5. Fix $l \geqslant 0$. It follows from our remarks on lengths of tuples in X that the sequence

$$0 \to \mathrm{MC}_{*,l}(G \cap H) \xrightarrow{(j^G_\#, -j^H_\#)} \mathrm{MC}_{*,l}(G) \oplus \mathrm{MC}_{*,l}(H) \longrightarrow \mathrm{MC}_{*,l}(G,H) \to 0$$

is short exact. Taking the associated long exact sequence and using the isomorphism $H_*(\mathrm{MC}_{*,l}(G,H)) \cong \mathrm{MH}_{*,l}(G \cup H)$ induced by the quasi-isomorphism of Theorem 6.5, one obtains the following long exact sequence.

$$\cdots \to \mathrm{MH}_{*,*}(G \cap H) \xrightarrow{(j^G_*, -j^H_*)} \mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(H)$$

$$\xrightarrow{i^G_* + i^H_*} \mathrm{MH}_{*,*}(G \cup H) \xrightarrow{\partial} \mathrm{MH}_{*-1,*}(G \cap H) \to \cdots$$

Writing $H = A \sqcup B$ where A is the full subgraph consisting of vertices that can be joined to $G \cap H$ by an edge-path, and $\pi \colon A \to G \cap H$ for the projection map, it follows that the composite

$$\mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(H) \xrightarrow{\;=\;} \mathrm{MH}_{*,*}(G) \oplus \mathrm{MH}_{*,*}(A) \oplus \mathrm{MH}_{*,*}(B)$$

$$\longrightarrow \mathrm{MH}_{*,*}(A)$$

$$\xrightarrow{-\pi_*} \mathrm{MH}_{*,*}(G \cap H)$$

is left inverse to $(j_*^G, -j_*^H)$. Consequently the long exact sequence splits into the split short exact sequences of the statement. The naturality claims are immediately verified. $\square$

Now we move on to the proof of Theorem 6.5, which is our excision theorem for magnitude chains. While the statement of our theorem is closely analogous to versions of excision for singular chains, we know of no analogy between the proof we give here and standard proofs of excision in singular homology, which use barycentric subdivision as their fundamental tool. The proof occupies the remainder of this section.

Definition 9.1. Let $l \geqslant 0$ and let $a, b \in G \cup H$ be a pair of vertices not both contained in G, and not both contained in H. (Thus we must have $a \in G \setminus H$ and $b \in H \setminus G$, or *vice versa*.) Define $A_{*,l}(a, b)$ to be the subcomplex of $\mathrm{MC}_{*,l}(G \cup H)$ spanned by those tuples $(x_0, \ldots, x_k)$ for which $x_0 = a$, $x_k = b$, and $x_1, \ldots, x_{k-1} \in G \cap H$.

Lemma 9.2. *The complex $A_{*,l}(a, b)$ is acyclic.*

Proof. For the purposes of the proof we assume that $b \in H \setminus G$ and $a \in G \setminus H$, the proof in the other case being similar. Let us define a map

$$s \colon A_{*,l}(a, b) \longrightarrow A_{*+1,l}(a, b)$$

by

$$s(x_0, \ldots, x_k) = \begin{cases} (-1)^k (x_0, \ldots, x_{k-1}, \pi(x_k), x_k) & \text{if } x_{k-1} \neq \pi(x_k), \\ 0 & \text{if } x_{k-1} = \pi(x_k). \end{cases}$$

We claim that $\partial \circ s + s \circ \partial = \mathrm{Id}$, so that s is a chain homotopy from Id to 0, and in particular that $A_{*,l}(a, b)$ is acyclic. Applied to a generator $(x_0, \ldots, x_k)$, this amounts to the claim that

$$\sum_{i=1}^{k} (-1)^i \partial_i s(x_0, \ldots, x_k) + \sum_{i=1}^{k-1} (-1)^i s \partial_i (x_0, \ldots, x_k) = (x_0, \ldots, x_k).$$

For $i = 1, \ldots, k-2$ we have that $\partial_i s(x_0, \ldots, x_k) + s\partial_i(x_0, \ldots, x_k) = 0$, since ∂_i does not affect the last two entries and s does not affect the first $(k-1)$. It therefore remains to show that

$$(-1)^{k-1}\partial_{k-1}s(x_0, \ldots, x_k) + (-1)^k \partial_k s(x_0, \ldots, x_k) + (-1)^{k-1}s\partial_{k-1}(x_0, \ldots, x_k)$$

is equal to $(x_0, \ldots, x_k)$. We verify this on a case-by-case basis.

- If $x_{k-1} = \pi(x_k)$ then $x_{k-2} \neq \pi(x_k)$ and $d(x_{k-2}, x_{k-1}) + d(x_{k-1}, x_k) = d(x_{k-2}, x_k)$. Consequently the first two terms in the above sum vanish, while the third term is $(x_0, \ldots, x_k)$.
- If $x_{k-1} \neq \pi(x_k)$ and $d(x_{k-2}, x_{k-1}) + d(x_{k-1}, x_k) > d(x_{k-2}, x_k)$, so that in addition $d(x_{k-2}, x_{k-1}) + d(x_{k-1}, \pi(x_k)) > d(x_{k-2}, \pi(x_k))$, then the first and third terms in the sum vanish, while the second is $(x_0, \ldots, x_k)$.
- If $x_{k-1} \neq \pi(x_k)$ and $d(x_{k-2}, x_{k-1}) + d(x_{k-1}, x_k) = d(x_{k-2}, x_k)$, so that in addition $d(x_{k-2}, x_{k-1}) + d(x_{k-1}, \pi(x_k)) = d(x_{k-2}, \pi(x_k))$ and $x_{k-2} \neq \pi(x_k)$, then the sum above becomes

$$-(x_0, \ldots, x_{k-2}, \pi(x_k), x_k) + (x_0, \ldots, x_k) + (x_0, \ldots, x_{k-2}, \pi(x_k), x_k).$$

In all cases the claim holds. This completes the proof. $\qquad\square$

Definition 9.3. For what follows we require the following notion. If C_* is a chain complex and $j \geqslant 0$, then the *j-th suspension* $\Sigma^j C_*$ of C_* is the chain complex in which $(\Sigma^j C_*)_i = C_{i-j}$.

Definition 9.4. Let $l \geqslant 0$. Given $b \in G \cup H \setminus G \cap H$, we define a complex $B_{*,l}(b)$ and a subcomplex $\bar{B}_{*,l}(b)$ as follows. If $b \in G \setminus H$ then $B_{*,l}(b)$ is defined to be the subcomplex of $\mathrm{MC}_{*,l}(G \cup H)$ spanned by tuples of the form $(x_0, \dots, x_k)$ with $x_k = b$ and $x_0, \dots x_{k-1} \in H$, and $\bar{B}_{*,l}(b)$ is defined to be the subcomplex of $B_{*,l}(b)$ spanned by tuples $(x_0, \dots, x_k)$ in which $x_0, \dots, x_{k-1} \in G \cap H$. If $b \in H \setminus G$ then the definitions are obtained in the same way, interchanging the role of G and H.

Lemma 9.5. *Let $l \geqslant 0$ and let $b \in G \cup H \setminus G \cap H$. Then the complex $B_{*,l}(b)/\bar{B}_{*,l}(b)$ is acyclic.*

Proof. Without loss we assume that $b \in H \setminus G$, the proof in the other case being similar. For $i = 0, \dots, l$ let F_i denote the subcomplex of $B_{*,l}(b)$ spanned by tuples $(x_0, \dots, x_k)$ in which $x_i, \dots, x_{k-1} \in G \cap H$. (In the case $i \geqslant k$ we impose no condition.) Thus we obtain a filtration

$$\bar{B}_{*,l}(b) = F_0 \subset F_1 \subset \cdots \subset F_l = B_{*,l}(b)$$

and it will suffice for us to show that for each $i = 1, \dots, l$ the quotient F_i/F_{i-1} is acyclic.

Let us describe the complex F_i/F_{i-1}. Its generators are tuples $(x_0, \dots, x_k)$ with $x_k = b$, with $x_i, \dots, x_{k-1} \in G \cap H$, and with $x_{i-1} \in G \setminus H$. Here the first two conditions guarantee that $(x_0, \dots, x_k)$ is a generator of F_i, while the third guarantees that it lies outside F_{i-1}. The differential ∂ on F_i/F_{i-1} is induced by the differential ∂ on F_i, which is the alternating sum $\sum_{i=1}^{k-1} (-1)^i \partial_i$ of the operators ∂_i which omit a generator's i-th term if the length is preserved, and otherwise send it to 0. Reducing to F_i/F_{i-1} we find that the operators $\partial_1, \dots, \partial_{i-1}$ become trivial, while $\partial_i, \dots, \partial_{k-1}$ retain their previous description.

Using the description from the last paragraph, we see that there is an isomorphism

$$\bigoplus_{(x_0, \dots, x_{i-1})} \Sigma^{i-1} A_{*,l-l'}(x_{i-1}, b) \xrightarrow{\cong} F_i/F_{i-1}.$$

Here the direct sum is taken over all tuples $(x_0, \dots, x_{i-1})$ of elements of H with $x_{i-1} \in H \setminus G$, and $l' = \ell(x_0, \dots, x_{i-1})$. The isomorphism sends the generator $(x_{i-1}, y_i, \dots, y_k)$ of the summand $A_{k-i+1,l-l'}(x_{i-1}, b)$ corresponding to $(x_0, \dots, x_{i-1})$ to $(-1)^{(i-1)k}$ times the generator $(x_0, \dots, x_{i-1}, y_i, \dots, y_k)$ of F_i/F_{i-1}. This is a map of chain complexes since on F_i/F_{i-1} in degree k the maps $\partial_1, \dots, \partial_{i-1}$ vanish while the maps $\partial_i, \dots, \partial_{k-1}$ are intertwined with the maps $\partial_1, \dots, \partial_{k-i}$ on $A_{k-i+1,l-l'}(x_{i-1}, b)$. The map is an isomorphism since it restricts to bijection between the generators of the domain and the range. $\qquad\square$

Proof of Theorem 6.6. We wish to prove that the inclusion

$$\mathrm{MC}_{*,l}(G, H) \longrightarrow \mathrm{MC}_{*,l}(G \cup H)$$

is a quasi-isomorphism. For $i = 0, \dots, l$ let F_i denote the subcomplex of $\mathrm{MC}_{*,l}(G \cup H)$ spanned by the tuples $(x_0, \dots, x_k)$ for which $x_0, \dots, x_{k-i}$ either all lie in G or all lie

in H. (When $i > k$ we impose no condition.) Thus we have a filtration

$$F_0 \subset \cdots \subset F_l$$

with

$$F_0 = \mathrm{MC}_{*,l}(G, H) \quad \text{and} \quad F_l = \mathrm{MC}_{*,l}(G \cup H).$$

So it will suffice to prove that for $i = 1, \ldots, l$ the quotient F_i/F_{i-1} is contractible.

There is an isomorphism

$$\bigoplus_{(x_{k-i+1}, \ldots, x_k)} \Sigma^{i-1} B_{*,l-l'}(x_{k-i+1})/\bar{B}_{*,l-l'}(x_{k-i+1}) \xrightarrow{\cong} F_i/F_{i-1},$$

where the direct sum is taken over all tuples $(x_{k-i+1}, \ldots, x_k)$ of elements of $G \cup H$ with $x_{k-i+1} \in G \cup H \setminus G \cap H$, and $l' = \ell(x_{k_i+1}, \ldots, x_k)$. The isomorphism is given on the summand corresponding to $(x_{k-i+1}, \ldots, x_k)$ by sending a generator $(x_0, \ldots, x_{k-i+1})$ to the generator $(x_0, \ldots, x_k)$ of F_i/F_{i-1}. We omit the details of why this is an isomorphism; the argument is similar to the one appearing in the proof of Lemma 9.2. Lemma 9.5 shows that the domain of this isomorphism is acyclic, and it follows that F_i/F_{i-1} is acyclic. This completes the proof. $\qquad\square$

10. Proof that joins are diagonal

Let G and H be graphs satisfying $G, H \neq \emptyset$. In this section we will prove Theorem 7.5, which states that the join $G \star H$ is diagonal, or in other words that $\mathrm{MH}_{k,l}(G \star H) = 0$ for $k < l$. We begin by stating the following, which is an immediate consequence of the definition of $G \star H$.

Lemma 10.1. *Let a and b be vertices of $G \star H$. Then $d(a, b)$ can only take the values 0, 1 and 2. Moreover $d(a, b) = 2$ only if a and b are both in G or both in H.*

For i in the range $0 \leqslant i \leqslant l - 1$, let F_*^i denote the subcomplex of $\mathrm{MC}_{*,l}(G \star H)$ spanned by generators $(x_0, \ldots, x_k)$ satisfying $d(x_j, x_{j+1}) = 2$ for some $j \leqslant i$. Thus

$$F_*^0 \subset F_*^1 \subset \cdots \subset F_*^{l-1} \subset \mathrm{MC}_{*,l}(G \star H).$$

Observe that F_*^{l-1} is simply the span of all generators such that $k < l$. Thus

$$F_k^{l-1} = \begin{cases} \mathrm{MC}_{k,l}(G \star H) & \text{if } k < l \\ 0 & \text{if } k = l. \end{cases}$$

Lemma 10.2. *For $i = 0, \ldots, l - 1$ the inclusion $F_*^i/F_*^{i-1} \hookrightarrow \mathrm{MC}_{*,l}(G \star H)/F_*^{i-1}$ induces the zero map in homology.*

Proof of Theorem 7.5, assuming Lemma 10.2. Let us first prove by induction on $i = 0, \ldots, l - 1$ that the inclusion $F_*^i \hookrightarrow \mathrm{MC}_{*,l}(G \star H)$ induces the zero map in homology. The initial case $i = 0$ is an instance of Lemma 10.2. Assuming that the claim holds for i, let us prove it for $i + 1$. Since by hypothesis the inclusion $F_*^i \hookrightarrow \mathrm{MC}_{*,l}(G \star H)$ induces the zero map in homology, it follows that the quotient $\mathrm{MC}_{*,l}(G \star H) \to$

$\mathrm{MC}_{*,l}(G \star H)/F_*^i$ induces an injection in homology. It will therefore suffice to prove that the composite

$$F_*^{i+1} \to \mathrm{MC}_{*,l}(G \star H) \to \mathrm{MC}_{*,l}(G \star H)/F_*^i$$

induces the zero map in homology. But this composite can be rewritten as the composite

$$F_*^{i+1} \to F_*^{i+1}/F_*^i \to \mathrm{MC}_{*,l}(G \star H)/F_*^i$$

in which the second map induces the zero map in homology by Lemma 10.2.

Since $F_*^{l-1} \hookrightarrow \mathrm{MC}_{*,l}(G \star H)$ is an isomorphism in degrees $k < l$, and induces the zero map in homology, it follows that $\mathrm{MH}_{k,l}(G \star H) = 0$ for $k < l$. $\qquad\square$

We now work towards a proof of Lemma 10.2. Given a vertex x of $G \star H$, denote by $A_*(x,l)$ the subcomplex of $\mathrm{MC}_{*,l}(G \star H)$ generated by the tuples of the form $(x, x_1, \ldots, x_k)$ with $d(x, x_1) = 2$, and denote by $B_*(x,l)$ the subcomplex of $\mathrm{MC}_{*,l}(G \star H)$ generated by the tuples of the form $(x, x_1, \ldots, x_k)$.

Lemma 10.3. *There is a commutative diagram*

$$
\begin{array}{ccc}
\bigoplus \Sigma^i A_*(x_i, l-i) & \lhook\joinrel\longrightarrow & \bigoplus \Sigma^i B_*(x_i, l-i) \\
{\scriptstyle \alpha} \downarrow {\scriptstyle \cong} & & \downarrow {\scriptstyle \beta} \\
F_*^i/F_*^{i-1} & \lhook\joinrel\longrightarrow & \mathrm{MC}_{*,l}(G \star H)/F_*^{i-1}
\end{array}
$$

in which the direct sums are indexed by tuples $(x_0, \ldots, x_i)$ of vertices of $G \star H$ satisfying $d(x_{j-1}, x_j) = 1$ for $j = 1, \ldots, i$, and in which the upper map is the direct sum of the inclusion maps.

Proof. We define α on the summand corresponding to the tuple $(x_0, \ldots, x_i)$ to be the map

$$\bar{\alpha} \colon \Sigma^i A_*(x_i, l-i) \longrightarrow F_*^i/F_*^{i-1}$$

that sends a generator $(x_i, \ldots, x_k)$ to $(-1)^{ik}(x_0, \ldots, x_i, \ldots, x_k)$. To see that $\bar{\alpha}$ is a chain map, observe that in degree k the differential on $\Sigma^i A_*(x_i, l-i)$ is the sum $\sum_{j=1}^{k-i-1} (-1)^j \partial_j$, while on F_*^i/F_*^{i-1} it is the sum $\sum_{j=1}^{k-1} (-1)^j \partial_j$. However on F_*^i/F_*^{i-1} the maps $\partial_1, \ldots, \partial_i$ all vanish, and in addition one can verify directly that $\bar{\alpha} \circ \partial_j = (-1)^i \partial_{i+j} \circ \bar{\alpha}$ for $j = 1, \ldots, k-i-1$. It follows that $\bar{\alpha}$ is indeed a chain map. To see that α is an isomorphism, observe that the generators of F_*^i/F_*^{i-1} are precisely the tuples $(x_0, \ldots, x_k)$ in which $d(x_0, x_1) = \cdots = d(x_{i-1}, x_i) = 1$ and $d(x_i, x_{i+1}) = 2$, so that α in fact restricts to a bijection between the generators of its domain and range. The chain map β is obtained in an entirely analogous way, and commutativity of the resulting square is then evident. $\qquad\square$

Lemma 10.4. *The inclusion $A_*(x,l) \hookrightarrow B_*(x,l)$ induces the trivial map in homology.*

Proof of Lemma 10.2, assuming Lemma 10.4. Since the upper arrow of the commutative diagram of Lemma 10.3 induces the zero map in homology, so does the lower arrow. $\qquad\square$

We now work towards the proof of Lemma 10.4. In order to do so, we assume without loss that $x \in G$, and we fix a vertex $y \in H$. Then we define the *height* of a generator $(x, x_1, \ldots, x_k)$ of $A_*(x, l)$ to be the largest integer h such that

$$d(x, x_1) = 2, \ d(y, x_2) = 2, \ d(x, x_3) = 2, \ \ldots \ d(-, x_h) = 2$$

where the final $-$ denotes x if h is odd and y if h is even. Thus all generators have height at least 1, and the height of a generator is no more than its degree.

Lemma 10.5. *If $(x, x_1, \ldots, x_k)$ is a generator of $A_*(x, l)$ then $\partial_j(x, x_1, \ldots, x_k)$ is either 0 or a generator of height at most $j - 1$.*

Proof. Suppose not. Since $\partial_j(x, x_1, \ldots, x_k)$ is nonzero, it follows that

$$d(x_{j-1}, x_j) = 1, \ d(x_j, x_{j+1}) = 1 \text{ and } d(x_{j-1}, x_{j+1}) = 2.$$

In particular x_{j-1} and x_{j+1} both lie in G or both lie in H. On the other hand, since $\partial_j(x, x_1, \ldots, x_k)$ has height at least j, then (assuming without loss that j is even) we have that $d(x, x_{j-1}) = 2$ and $d(y, x_{j+1}) = 2$, so that x_{j-1} lies in G and x_{j+1} lies in H. This is a contradiction. $\square$

Proof of Lemma 10.4. For $i \geqslant 1$, let $s_i \colon A_*(x, l) \to B_{*+1}(x, l)$ be the map defined on generators by the rule

$$s_i(x, x_1, \ldots, x_k) = \begin{cases} (\underbrace{x, y, x, y, \ldots,}_{i+1 \text{ terms}} x_i, \ldots, x_k) & \text{if } i \leqslant h \\ 0 & \text{if } i > h \end{cases}$$

where h denotes the height of $(x, x_1, \ldots, x_k)$. In the first case, the assumption on the height guarantees that the term on the right has length exactly l. We have the following compatibilities between the s_i and the operators ∂_j.

- $\partial_j \circ s_i = 0$ for $1 \leqslant j < i$
- $\partial_{i+1} \circ s_i = \partial_{i+1} \circ s_{i+1}$ for $1 \leqslant i$
- $\partial_j \circ s_i = s_i \circ \partial_{j-1}$ for $i \geqslant 1$ and $j \geqslant i + 2$
- $\partial_1 \circ s_1$ is the inclusion $A_*(x, l) \hookrightarrow B_*(x, l)$
- $s_i \circ \partial_j = 0$ for $1 \leqslant j \leqslant i$

The first four follow by direct computation and the last follows from Lemma 10.5.

Now define $s \colon A_*(x, l) \to B_{*+1}(x, l)$ by $s = \sum_{i \geqslant 1} (-1)^i s_i$. We claim that s is a chain homotopy from the inclusion map to the zero map, or in other words that $s \circ \partial + \partial \circ s$ is the inclusion. From the properties listed above we have

$$\partial \circ s = \partial_1 \circ s_1 - \sum_{j \geqslant i+1, \, i \geqslant 1} (-1)^{i+j} s_i \circ \partial_j$$

and

$$s \circ \partial = \sum_{j \geqslant i+1, i \geqslant 1} (-1)^{i+j} s_i \circ \partial_j.$$

It follows that $\partial \circ s + s \circ \partial = \partial_1 \circ s_1$, which is the inclusion map. $\square$

Appendix A. Numerical examples

In this appendix we give various further computer calculated examples of ranks of magnitude homology groups. The reader can verify these computations using `rational_graph_homology_arxiv.py`, written in sage and python, which was uploaded to the arXiv with this paper.

A.1. Cyclic graphs

Cyclic graphs seem to have a clear pattern in their graph homology, and this pattern depends on whether the graph has an even or an odd number of vertices.

First consider graphs with odd numbers of vertices as exemplified by the 7-cyclic graph below. (The 5-cyclic graph was given in Table 1.)

	0	1	2	3	4	5	6	7	8	9	10	11
0	7											
1		14										
2			14									
3				14								
4			14		14							
5				42		14						
6					70		14					
7						98		14				
8					28		126		14			
9						112		154		14		
10							252		182		14	
11								448		210		14

The behaviour for a cyclic graph with n vertices where n is odd, seems to be as follows. The non-zero ranks are ordered in diagonal lines, so in the above the first diagonal starts at $k = 0$, $l = 0$ with $7, 14, 14, \ldots$ and the second diagonal starting at $k = 2$, $l = 4$ with $14, 42, 70, \ldots$: in general the ith diagonal starts at $k = 2(i - 1)$ and $l = (i - 1)(n + 1)/2$. Denote by $T_{i,j}^n$ the jth non-zero entry in the ith diagonal, then it would appear that these are given by the following recursion relation.

$$T_{1,1}^n = n; \quad T_{1,2}^n = 2n; \quad T_{i,j}^n = T_{i,j-1}^n + 2T_{i-1,j}^n.$$

If that is the case then $T_{i,j}^n/(n2^{i-1})$ is an integer independent of n.

On the other hand, the even case, as exemplified by the 8-cycle graph below, appears more straightforward.

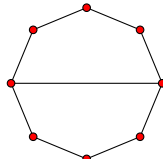|	0	1	2	3	4	5	6	7	8	9	10
0	8										
1		16									
2			16								
3				16							
4			8		16						
5				16		16					
6					16		16				
7						16		16			
8					8		16		16		
9						16		16		16	
10							16		16		16

It appears that for an n-cycle graph with n even, the first rank in each diagonal is n and the subsequent ranks are $2n$. The ith diagonal starts at $k = 2(i - 1)$ and $l = (i - 1)n/2$.

A.2. Projecting is necessary for Mayer-Vietoris

This example shows the necessity of the projecting condition in the Mayer-Vietoris short exact sequence of a convex decomposition of a graph. Consider the graph X pictured below. This is the union of two 5-cycle graphs along a common edge, i.e. along a 2-cycle. This is a convex decomposition of the graph, however, neither 5-cycle is projecting, as the 'apex' of each 5-cycle can't project. If this graph did have a Mayer-Vietoris short exact sequence then for each k and l we would have

$$\operatorname{rank} \operatorname{MH}_{k,l}(X) = 2 \cdot \operatorname{rank} \operatorname{MH}_{k,l}(C_5) - \operatorname{rank} \operatorname{MH}_{k,l}(C_2).$$

The 2-cycle is diagonal with $\operatorname{rank} \operatorname{MH}_{k,k}(C_2) = 2$ for all k. Comparing the table of ranks below with that for the 5-cycle in Table 1 we see that the first two diagonals are as you would expect if the above equation were satisfied, however the third diagonal is wrong, with the first differing entry being $\operatorname{rank} \operatorname{MH}_{2,4}(X) = 2$.

|	0	1	2	3	4	5	6	7	8	9	10
0	8										
1		18									
2			18								
3			20	18							
4			2	60	18						
5				12	100	18					
6					76	140	18				
7					8	236	180	18			
8					2	56	492	220	18		
9						16	280	844	260	18	
10							92	904	1292	300	18

A.3. Some symmetric cubic graphs

Here we include some classical graphs for further examples. These graphs all have large symmetry groups, and these act on the magnitude homology groups. Intriguingly the order of the automorphism group is showing up in the ranks of the homology group. This might be simply indicating that the automorphism group is acting freely transitively on the generators of those magnitude homology groups.

Some of these graphs have patterns in the ranks of the magnitude homology groups reminiscent of those for the cyclic graphs. We leave the reader to discover them.

The Möbius Kantor graph and the Pappus graph illustrate that the rank can sometimes decrease as you move down a diagonal.

A.3.1. Petersen graph

The automorphism group has order 120.

	0	1	2	3	4	5	6	7	8
0	10								
1		30							
2			30						
3			120	30					
4				480	30				
5					840	30			
6					1440	1200	30		
7						7200	1560	30	
8							17280	1920	30

A.3.2. Heawood graph

The automorphism group has order 336.

	0	1	2	3	4	5	6	7	8
0	14								
1		42							
2			42						
3			112	42					
4				336	42				
5					336	42			
6					896	336	42		
7						2688	336	42	
8							2688	336	42

A.3.3. Tutte Coxeter graph

The automorphism group has order 1440.

	0	1	2	3	4	5	6	7	8
0	30								
1		90							
2			90						
3				90					
4			480		90				
5				1440		90			
6					1440		90		
7						1440		90	
8					7680		1440		90

A.3.4. Moebius Kantor graph

The automorphism group has order 96.

	0	1	2	3	4	5	6	7	8
0	16								
1		48							
2			48						
3			112	48					
4				304	48				
5				48	288	48			
6					832	288	48		
7						1952	288	48	
8						656	1776	288	48

A.3.5. Pappus graph

The automorphism group has order 216.

	0	1	2	3	4	5	6	7	8
0	18								
1		54							
2			54						
3			108	54					
4				252	54				
5				108	216	54			
6					756	216	54		
7						1188	216	54	
8						1224	972	216	54

A.4. The icosahedral graph

This is the only graph which our calculations show to be diagonal, but for which we know of no proof that it is diagonal.

	0	1	2	3	4	5	6	7
0	12							
1		60						
2			240					
3				912				
4					3420			
5						12780		
6							47712	
7								178080

References

[1] Albrecht Dold. *Lectures on algebraic topology.* Classics in Mathematics. Springer-Verlag, Berlin, 1995. Reprint of the 1972 edition.

[2] Samuel Eilenberg and Saunders Mac Lane. On the groups of $H(\Pi, n)$. I. *Ann. of Math. (2)*, 58:55–106, 1953.

[3] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[4] Laure Helme-Guizon and Yongwu Rong. A categorification for the chromatic polynomial. *Algebr. Geom. Topol.*, 5:1365–1388, 2005.

[5] Edna F. Jasso-Hernandez and Yongwu Rong. A categorification for the Tutte polynomial. *Algebr. Geom. Topol.*, 6:2031–2049 (electronic), 2006.

[6] Mikhail Khovanov. A categorification of the Jones polynomial. *Duke Math. J.*, 101(3):359–426, 2000.

[7] F. William Lawvere. Metric spaces, generalized logic, and closed categories. *Rend. Sem. Mat. Fis. Milano*, 43:135–166 (1974), 1973.

[8] Tom Leinster. The Euler characteristic of a category. *Doc. Math.*, 13:21–49, 2008.

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

[10] Tom Leinster. Tutte polynomials and magnitude functions. blog post at the *n*-Category Café, `http://golem.ph.utexas.edu/category/2013/04/tutte_polynomials_and_magnitud.html`, April 2013. [Accessed 16 February 2015].

[11] Tom Leinster. The magnitude of a graph. Preprint, available at arXiv:1401.4623, 2014.

[12] Tom Leinster and Christina A. Cobbold. Measuring diversity: the importance of species similarity. *Ecology*, 93(3):477–489, 2012.

[13] Tom Leinster and Simon Willerton. On the asymptotic magnitude of subsets of Euclidean space. *Geom. Dedicata*, 164:287–310, 2013.

[14] J. P. May. *A concise course in algebraic topology*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1999.

[15] Mark W. Meckes. Positive definite metric spaces. *Positivity*, 17(3):733–757, 2013.

[16] Mark W. Meckes. Magnitude, Diversity, Capacities, and Dimensions of Metric Spaces. *Potential Anal.*, 42(2):549–572, 2015.

[17] Peter Ozsváth and Zoltán Szabó. An introduction to Heegaard Floer homology. In *Floer homology, gauge theory, and low-dimensional topology*, volume 5 of *Clay Math. Proc.*, pages 3–27. Amer. Math. Soc., Providence, RI, 2006.

[18] Marko Stošić. Categorification of the dichromatic polynomial for graphs. *J. Knot Theory Ramifications*, 17(1):31–45, 2008.

[19] Charles A. Weibel. *An introduction to homological algebra*, volume 38 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1994.

[20] Simon Willerton. On the magnitude of spheres, surfaces and other homogeneous spaces. *Geom. Dedicata*, 168:291–310, 2014.

Richard Hepworth   `r.hepworth@abdn.ac.uk`

Institute of Mathematics, University of Aberdeen, Aberdeen, AB24 3UE, United Kingdom

Simon Willerton   `s.willerton@sheffield.ac.uk`

School of Mathematics and Statistics, Hicks Building, University of Sheffield, Sheffield, S3 7RH, United Kingdom