Natural Homology

Jérémy Dubut[1,2], Éric Goubault[1], and Jean Goubault-Larrecq[2]

[1] LIX, Ecole Polytechnique, F-91128 Palaiseau Cedex
[2] LSV, ENS Cachan, F-94230 Cachan

Abstract. We propose a notion of homology for directed algebraic topology, based on so-called natural systems of abelian groups, and which we call natural homology. As we show, natural homology has many desirable properties: it is invariant under isomorphisms of directed spaces, it is invariant under refinement (subdivision), and it is computable on cubical complexes.
Keywords: directed algebraic topology, homology, path space, geometric semantics, persistent homology, natural system.

1 Introduction

The purpose of this paper is to introduce a satisfactory notion of homology for directed algebraic topology. Let us clarify.

From a mathematical point of view, algebraic topology is a well-established, and rich domain. Its purpose is to classify shapes (topological spaces), disregarding differences in shapes that can be obtained from each other by continuous deformations (homotopy equivalence). A particularly useful notion there is *homology*, which is a sound abstraction of homotopy equivalence. Soundness means, notably, that if two spaces are not homologous, then one cannot deform one into the other continuously, in whichever way we attempt this. Homology is also computable [22, 14] on finitely presented shapes (simplicial, resp. cubical sets), in sharp contrast with homotopy.

Directed algebraic topology is a variant of algebraic topology where the spaces also have a direction of time [12], and deformations must not only be continuous but also preserve the direction of time. Directed algebraic topology was born out of the so-called geometric semantics of concurrent processes (progress graphs [3], generally attributed to E. W. Dijkstra), and the higher-dimensional automaton model of true concurrency [18]. Imagine n concurrent processes, each with a local time $t_i \in [0, 1]$. A configuration is a point in $[0, 1]^n$, and a trajectory is a continuous *and monotonic* map from $[0, 1]$ to $[0, 1]^n$: monotonicity (a.k.a., directedness) reflects the fact that no process can go back in time. One can arguably consider as equivalent any two trajectories that are dihomotopic, namely that can be deformed into each other continuously, while respecting monotonicity at all times. This not only yields a geometric semantics for concurrency, but also one that is at the root of fast algorithms for state-space reduction, deadlock and unreachable states detection, and verification of coordination properties, as in e.g. [8, 6, 9, 10].

However intuitive the geometric semantics of processes may be, previous attempts at defining notions of homology suited to *directed* algebraic topology were somehow disappointing. We discuss them in Section 2. Our contribution is: 1. a new notion of directed homology, based on so-called natural systems of abelian groups, and which we call *natural homology*, and 2. the proofs of its basic properties, the most important probably being *invariance under refinement*—a central property of truly concurrent semantics [23]. We define and motivate natural systems (on pospaces) in Section 3, and the requirement for a notion of bisimilarity between them in Section 4. Finite cubical complexes provide finite presentations of pospaces, as explained in Section 5. We adapt the notion of natural homology to cubical complexes, and we prove the properties mentioned above in Section 6. We conclude in Section 7.

2 Related Work

Homology is a classical concept in (undirected) algebraic topology, and we shall only discuss the notions that various authors have proposed to fit the directed case. Non-abelian homology [17] may seem promising; so far, we are only aware of work by Krishnan in this direction [16]. All the other attempts we know [7, 11, 4, 15] have the same weakness: they are not precise enough.

By this we mean, *first*, that directed homology should not be invariant under (undirected) homotopy. If it were, it would be blind to the essential feature of directed algebraic topology: that directions are important. In that case, we may as well use classical homology theories, which are already homotopy invariants. This is the bane of early directed homology theories [7].

Second, the Hurewicz theorems in classical algebraic topology state that the loss of information we must pay when replacing homotopy groups by homology groups is limited. One trivial consequence of these results is that a space X whose first homotopy group is non-trivial (different from 0) also has a non-trivial first homology group $H_1(X)$. Similarly, we would like any dihomotopically non-trivial shape to have non-trivial directed homology. This fails in any of the remaining proposals [11, 4, 15], as we explain now.
Consider the *matchbox* example, due to Fahrenberg [4], shown on the right. The exploded view is on the left, the finished product on the right. Note that this is not a cube: the bottom face and the interior are missing. The matchbox is meant to stand on its tip (vertex s), and time goes up, that is, a point is before another one if and only if its altitude is smaller.

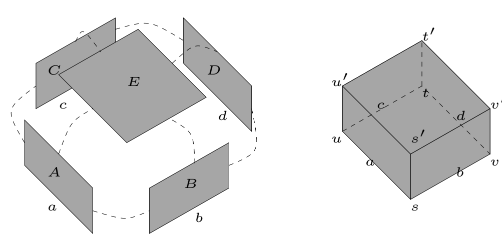

Fig. 1: Fahrenberg's matchbox example

Look at the edges a, b, c, d. The concatenation $a \star c$ of a and c is a directed path from s to t, and $b \star d$ is another one. A dihomotopy between these two would be a continuous map h from $[0,1]$ to the space of directed paths from s to t such that $h(0) = a \star c$ and $h(1) = b \star d$. If h existed, then for some α, $h(\alpha)$ would be a directed path going through t', and then down to t: this is impossible, since $h(\alpha)$ must remain directed for all values of α. Indeed, $a \star c$ and $b \star d$ are not dihomotopic. In particular, the matchbox has non-trivial dihomotopy. This contrasts with the undirected view: the matchbox is contractible, and in particular all its classical homotopy and homology groups are trivial (equal to 0). We now examine why this example is not dealt properly by the existing proposals for directed homology.

Grandis [11] defined a notion of directed homology by enriching the classical homology groups of the space at hand with a partial order. The idea is that the generators of homology groups are holes, and the partial order serves to remember which holes come before which others. Since the classical homology groups of the matchbox are trivial, so is the ordered homology group that Grandis defines. However, the matchbox is dihomotopically non-trivial.

Kahl's homology graphs [15], which are defined, similarly, as homology groups with extra relations, suffer from the same problem.

Finally, the matchbox was produced by Fahrenberg to show that his own notion of directed homology [4] is unsatisfactory. The problem is common to any notion of directed homology that is based on homology *groups*, or even on cancellative monoids (an assumption used in [17]). Let e be the directed edge from t to t'. The directed paths $a \star c \star e$ and $b \star d \star e$ *are* dihomotopic, hence must be in the same equivalence class of (directed) homology. Cancellation of e then implies that $a \star c$ and $b \star d$ must be equivalent with respect to directed homology.

Our solution avoids cancellation by working with so-called *natural systems* of abelian groups, and builds upon several prior strands of research. The natural systems themselves arise from Baues and Wirsching's work on the cohomology of small categories [1]. The view of directed homotopy (resp., homology) as being based on classical homotopy (resp., homology) of spaces of traces can be traced to Raussen [19], and the carrier morphism we use near the end has its origin in work by Fajstrup [5]. The idea of using natural systems, indexed by so-called traces, to organize information on several topological objects together is also present in [19], although Raussen did not apply this to homology.

The notion of bisimulation of such natural systems, which we define to forget about irrelevant differences between them, is novel. We shall see later why this is needed. To make the argument short, a natural system of homology groups is an immense picture of homology groups, indexed by traces, but we do not care about the whole picture, rather about the patterns of change we see in groups as we extend the traces. This is a very similar concern as in persistent homology (of undirected spaces) [2]. However, the latter uses a very simple, linear ordering of the indices, and we must deal with a much more complex situation.

3 Natural Homology of Pospaces, and Natural Systems

Let X be a *pospace*, i.e., a topological space X with a partial order $\leq$ whose graph is closed in X^2. A fundamental example is I, the interval $[0,1]$ with the usual ordering. A *path* from a to b in X is a continuous map $\pi\colon I \to X$ such that $\pi(0) = a$ and $\pi(1) = b$. It is a *dipath* (short for directed path) if and only if it is also monotonic. Following Raussen [19], a (directed) *trace* is the equivalence class $\langle\pi\rangle$ of a dipath π modulo reparametrization: a reparametrization is a monotonic continuous onto map φ from I to I, and π and π' are equivalent if and only if there are two reparametrizations φ, ψ such that $\pi \circ \varphi = \pi' \circ \psi$. Traces are dipaths, up to the speed at which we travel from time $t = 0$ to time $t = 1$, which is considered irrelevant. Given two (di)paths π from a to b and π' from b to c, the concatenation $\pi \star \pi'$ maps $t \in [0, 1/2]$ to $\pi(2t)$ and $t \in [1/2, 1]$ to $\pi'(2t - 1)$. This induces an associative operation $\star$ on the quotient space of traces.

A standard notion of classical algebraic topology is *homotopy*. Consider the n-cube I^n, and write ∂I^n for its boundary. Given a path-connected space Y, and fixing a so-called base point $y \in Y$, an n-loop in Y is a continuous map from I^n to Y that maps ∂I^n to y. In particular, for $n = 1$, a 1-loop is just a path from y to itself. A homotopy h between two n-loops λ, λ' is a continuous map from $I \times I^n$ to Y such that $h(0, \_) = \lambda$, $h(1, \_) = \lambda'$, and, for each α, $h(\alpha, \_)$ maps ∂I^n to the base point y. If such a homotopy exists, then one says that λ and λ' are *homotopic*—we can deform one continuously into the other. We let $\pi_n(Y)$ denote the set of equivalence classes of n-loops of Y modulo homotopy. It is useful to visualize the case $n = 1$, where loops modulo homotopy form a group $\pi_1(Y)$ under concatenation.

One can define *di*homotopy and *di*homology similarly, but one should be careful. For example, *directed* n-loops in a pospace are trivial. Instead, Raussen [19] proposes to consider $(n - 1)$-loops in the space $Y = Tr(X; a, b)$ of traces from a to b in X. For example, for $n = 1$, the points of Y are the traces from a to b, and any path between two such points $\langle\pi\rangle$ and $\langle\pi'\rangle$ is easily seen to be (up to reparametrization of π and π') a homotopy between π and π' that fixes the two endpoints a and b: a continuous map $h\colon I \times I \to X$ such that $h(0, \_) = \pi$, $h(1, \_) = \pi'$, $h(\_, 0) = a$, $h(\_, 1) = b$, and, for every value of the deformation parameter α, $h(\alpha, \_)$ is a dipath from a to b. The zeroth homology group $H_0(Y)$ of Y is of the form $\mathbb{Z}^k$ with k the number of equivalence classes of traces from a to b up to dihomotopy. In general, we may define the n-th directed homology group $\overrightarrow{H}_n(X; a, b)$ as the ordinary $(n - 1)$st singular homology group of $Tr(X; a, b)$.

While this looks like a perfect definition of directed homology, this is still unsatisfactory. Consider the following two pospaces. In each, time goes from left to right and from bottom to top, starting at $\mathbf{0}$ and ending at $\mathbf{1}$. The leftmost pospace is the geometric semantics of a PV-program [3] extended with global synchronization (written "$\bullet$"), namely $(\mathsf{P}a\mathsf{V}a \,||\, \mathsf{P}a\mathsf{V}a) \bullet (\mathsf{P}a\mathsf{V}a\mathsf{P}b\mathsf{V}b \,||\, \mathsf{P}b\mathsf{V}b\mathsf{P}a\mathsf{V}a)$.

The rightmost pospace is the geometric semantics of the PV-program $\mathsf{P}a\mathsf{V}a\mathsf{P}a\mathsf{V}a \;||\; \mathsf{P}a\mathsf{V}a\mathsf{P}a\mathsf{V}a$. The four squares we carved out are those regions of space where the two processes would have acquired the lock a—which is impossible.

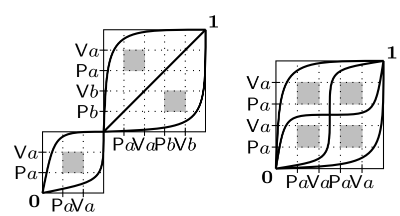

It is natural to compare pospaces X with two distinguished endpoints $\mathbf{0}$ and $\mathbf{1}$ by determining their dihomology groups $\overrightarrow{H}_n(X;\mathbf{0},\mathbf{1})$. For $n = 1$, the two pospaces above both have exactly six traces up to dihomotopy, shown as thick lines: the two pospaces have the same dihomology group for $n = 1$, namely $\mathbb{Z}^6$. For $n \geq 2$, they also have the same dihomology groups $\overrightarrow{H}_n(X;\mathbf{0},\mathbf{1})$, because the path-connected components of their trace spaces are contractible. Therefore, the $\overrightarrow{H}_n(\_;\mathbf{0},\mathbf{1})$ construction does not distinguish the two pospaces, although they visibly have very different behaviors.

However, when we zoom in, and look at different pairs of endpoints, the situation changes. Consider the $(\mathsf{P}a\mathsf{V}a \mid\mid \mathsf{P}a\mathsf{V}a) \bullet (\mathsf{P}a\mathsf{V}a\mathsf{P}b\mathsf{V}b \mid\mid \mathsf{P}b\mathsf{V}b\mathsf{P}a\mathsf{V}a)$ pospace again, but look at its dihomology group $\overrightarrow{H}_1(X;\mathbf{0},t)$, where t is shown on the right: this is equal to $\mathbb{Z}^4$. However, no trace space of the other pospace $(\mathsf{P}a\mathsf{V}a\mathsf{P}a\mathsf{V}a \mid\mid \mathsf{P}a\mathsf{V}a\mathsf{P}a\mathsf{V}a)$ has exactly four connected components, so $\mathbb{Z}^4$ cannot be a dihomology group of the latter. This detects an essential difference between the two pospaces.

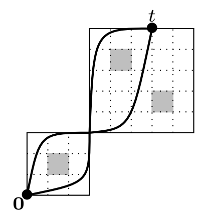

Given a trace $\langle\pi\rangle$, with π a dipath of X from a to b, we define $\overrightarrow{H}_n(X;\langle\pi\rangle) = \overrightarrow{H}_n(X;a,b)$. The family of groups $\overrightarrow{H}_n(X;\langle\pi\rangle)$, when $\langle\pi\rangle$ varies over traces, has extra structure: if α is a dipath from a' to a and β is a dipath from b to b', we obtain a continuous map from $Tr(X;a,b)$ to $Tr(X;a',b')$, which maps every trace $\langle\pi'\rangle$ to $\langle\alpha\star\pi'\star\beta\rangle$. We call *extensions* the pairs $(\langle\alpha\rangle,\langle\beta\rangle)$. Applying the H_{n-1} functor to the map $\langle\pi'\rangle \mapsto \langle\alpha\star\pi'\star\beta\rangle$, we obtain a morphism of groups $\overrightarrow{H}_n(X;\langle\pi\rangle)$ to $\overrightarrow{H}_n(X;\langle\alpha\star\pi\star\beta\rangle)$, which we denote by $\langle\alpha\star\_\star\beta\rangle$. This keeps track of how the homology picture formed by the traces from a to b inserts into the larger picture formed by the traces from the lower point a' to the higher point b'.

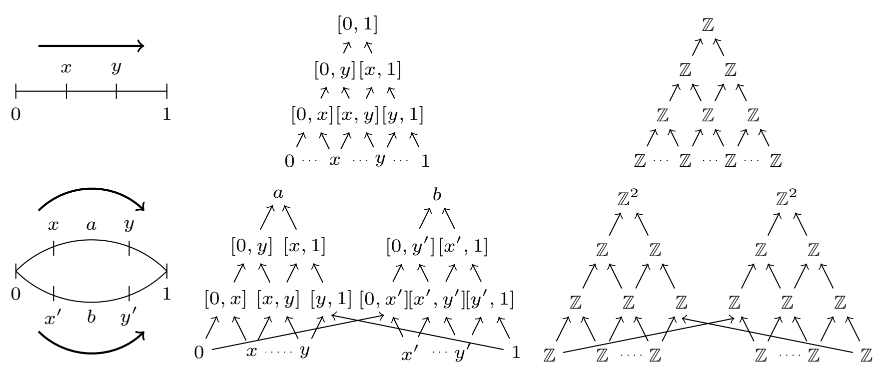

Fig. 2: Natural homology of two simple pospaces

We are ready to give formal definitions. Let X be a pospace, and Fc_X be the small category whose objects are traces of X, and whose morphisms are extensions. This is the *factorization category* [1] of the small category whose objects are points of X, and whose morphisms are traces. A *natural system (of abelian groups)* is by definition a functor from the factorization category of a small category (e.g. Fc_X) to the category **Ab** of abelian groups.

Definition 1 (Natural homology). *The* natural homology *of X is the natural system $\overrightarrow{H}_n(X)$ that, as a functor, maps every trace $\langle\pi\rangle$ to $\overrightarrow{H}_n(X;\langle\pi\rangle)$, and every extension $(\langle\alpha\rangle,\langle\beta\rangle)$ to $\langle\alpha\star\_\star\beta\rangle$.*

Figure 2 shows a few simple examples of natural homology systems $\overrightarrow{H}_1$. On top, we consider the pospace I itself. The middle diagram pictures the full subcategory of the (uncountable) category Fc_I whose objects are traces, which are identified to segments $[s,t]$ with $s,t\in\{0,1,x,y\}$ and $s\leq t$ (x and y are shown on the left). $[s,s]$ simplifies to s. The rightmost diagram pictures the collection of dihomology groups above each object of Fc_I. The bottom row is similar, and applies to two copies of I glued at 0 and 1.

4 Bisimilarity of Natural Systems

The natural homology $\overrightarrow{H}_n(X)$ of a pospace X is very fine-grained: it not only records local homology groups $\overrightarrow{H}_n(X;\langle\pi\rangle)$, but also for which traces they occur. If we wish to compare the natural homology of two pospaces, the latter should be unimportant. Just as with persistent homology [2], it is the patterns of change, between groups $\overrightarrow{H}_n(X;\langle\pi\rangle)$ when $\langle\pi\rangle$ is changed into $\langle\alpha\star\pi\star\beta\rangle$ by extension, that count, not the values of the trace $\langle\pi\rangle$.

We introduce a notion of bisimulation of natural systems, and more generally of **Ab**-valued functors, that smoothes this out. Given two small categories X, Y and two functors $F : X \longrightarrow \mathbf{Ab}$ and $G : Y \longrightarrow \mathbf{Ab}$, we call *bisimulation* between F and G any set R of triples (x,η,y) with x an object of X, y an object of Y and η an isomorphism of groups from Fx to Gy such that:

1. for every object x of X, R contains some triple of the form (x,η,y), and similarly for every object y of Y;
2. for every triple $(x,\eta,y)\in R$ and every morphism $i : x\longrightarrow x'$ in X, there is a triple $(x',\eta',y')\in R$ (hence η' is an isomorphism) and a morphism $j : y\longrightarrow y'$ in Y such that $\eta'\circ Fi = Gj\circ\eta$, and symmetrically,

   for every $(x,\eta,y)\in R$ and every morphism $j : y\longrightarrow y'$ of Y there is a triple $(x',\eta',y')\in R$ and a morphism $i : x\longrightarrow x'$ such that $\eta'\circ Fi = Gj\circ\eta$.

$$\begin{array}{cccccc} x & Fx & \xrightarrow{\eta} & Gy & y \\ i\downarrow & Fi\downarrow & & \downarrow Gj & \downarrow j \\ x' & Fx' & \xrightarrow[\eta']{} & Gy' & y' \end{array}$$

We say that F and G are *bisimilar* if and only if there is a bisimulation R between them. This is an equivalence relation.

A practical way of showing that two functors are bisimilar is by exhibiting an *open map* from one to the other. This arises from the theory of Joyal *et al.* [13]; the details are relatively unimportant here, and are relegated to Appendix A. The open maps from a functor $F : E\longrightarrow\mathbf{Ab}$ to a functor $G : X\longrightarrow\mathbf{Ab}$ are the pairs (Φ, σ) where Φ is a fibration from E to X, and σ is a natural isomorphism from F to $G \circ \Phi$. We say that $\Phi : E \longrightarrow X$ is a *fibration* if and only if: (1) Φ is surjective on objects, i.e., for every object x of X there is an object e of E such that $\Phi(e) = x$, and (2) for every object e of E, every morphism $f : \Phi(e) \longrightarrow x'$ in X lifts to a morphism $h : e \longrightarrow e'$ in E such that $\Phi(h) = f$ (in particular, $\Phi(e') = x'$).

We prove the following in Appendix B.

Proposition 1. *Two functors $F : X \longrightarrow \mathbf{Ab}$ and $G : Y \longrightarrow \mathbf{Ab}$ are bisimilar if and only if they are related by a span of open maps.*

We shall apply this to compare our natural homology of pospaces to a similar notion of natural homology of cubical complexes. We can think of the latter as a form of syntax for the latter. Their semantics is given by geometric realization, as we now explain.

5 Cubical Complexes and Their Geometric Realization

A cubical complex is a finite union of certain cubes of side-length 1 parallel to the axes in $\mathbb{R}^d$, whose vertices have integer coordinates [14]. Formally, let us define a (d-dimensional) *cubical complex* K as a finite set of *cubes* $(D, \boldsymbol{x})$, where $D \subseteq \{1, 2, \cdots, d\}$ and $\boldsymbol{x} \in \mathbb{Z}^d$, which is closed under taking past and future faces (to be defined shortly). The cardinality $|D|$ of D is the *dimension* of the cube $(D, \boldsymbol{x})$. Let $\mathbf{1}_k$ be the d-tuple whose kth component is 1, all others being 0. Each cube $(D, \boldsymbol{x})$ is *realized* as the geometric cube $\rho(D, \boldsymbol{x}) = I_1 \times I_2 \times \cdots \times I_d$ where $I_k = [x_k, x_k + 1]$ if $k \in D$, $I_k = [x_k, x_k]$ otherwise, matching the definition of [14].

When $|D| = n$, we write $D[i]$ for the ith element of D. For example, if $D = \{3, 4, 7\}$, then $D[1] = 3$, $D[2] = 4$, $D[3] = 7$. We also write $\partial_i D$ for D minus $D[i]$. Every n-dimensional cube $(D, \boldsymbol{x})$ has n *past faces* $\partial_i^0(D, \boldsymbol{x})$, defined as $(\partial_i D, x)$, and n *future faces* $\partial_i^1(D, \boldsymbol{x})$, defined as $(\partial_i D, x + \mathbf{1}_{D[i]})$, $1 \leq i \leq n$.

Together with these face operators, K exhibits the structure of a so-called *precubical set*, in the sense that the *precubical equations* $\partial_i^\alpha \partial_j^\beta = \partial_{j-1}^\beta \partial_i^\alpha$ ($1 \leq i < j$, $\alpha, \beta \in \{0, 1\}$) are satisfied. Precubical sets are a natural representation for truly concurrent processes, and occur as the main ingredient in the definition of *higher-dimensional automata* (HDA; see [18]). Cubical complexes are very particular precubical sets. Notably, they are non-looping in the sense of Fajstrup [5]. They are however enough for most purposes, including the definition of geometric semantics of finite PV-programs.

The *geometric realization* $\overrightarrow{Geom}(K)$ of a precubical set K is obtained, informally, by drawing it. For example, Fahrenberg's matchbox (Fig. 1) is really obtained by drawing a finite precubical set (a cubical complex, really) with 2-dimensional cubes A, B, C, D, and E, defined so that $\partial_1^0 A = \partial_1^0 B$ (the lower dashed connection in the exploded view), $\partial_2^0 A = a$, $\partial_2^0 B = b$, $\partial_1^0 a = \partial_1^0 b = s$, and so on. Formally, let $\overrightarrow{I}^n$ be the standard oriented cube $[0, 1]^n$, with the pointwise ordering. Form the coproduct $A = \sum_{e \in K} \overrightarrow{I}^{n_e}$ where n_e is the dimension of e, i.e., the disjoint union of as many copies of $\overrightarrow{I}^n$ as there are n-dimensional cubes e, for $n \in \mathbb{N}$; the elements of A are pairs $(e, \boldsymbol{a})$ where e is an n-dimensional cube in K and $\boldsymbol{a} \in [0,1]^n$, for some n. For convenience, for $\boldsymbol{a} = (a_1, a_2, \cdots, a_n)$, we write $\delta_i^\alpha \boldsymbol{a}$ for $(a_1, a_2, \cdots, a_{i-1}, \alpha, a_i, \cdots, a_n)$. Finally, we glue all these cubes together, by defining $\overrightarrow{Geom}(K)$ as $A/\equiv$, where $\equiv$ is the smallest equivalence relation such that $(\partial_i^\alpha e, \boldsymbol{a}) \equiv (e, \delta_i^\alpha \boldsymbol{a})$. We shall write $[e, \boldsymbol{a}]$ for the point obtained as the equivalence class of $(e, \boldsymbol{a})$.

For a cubical complex K, the element $[(D, \boldsymbol{x}), \boldsymbol{a}]$ (with $D \subseteq \{1, 2, \cdots, d\}$, $|D| = n$, $\boldsymbol{x} \in \mathbb{Z}^d$, $\boldsymbol{a} \in [0,1]^n$) of $\overrightarrow{Geom}(K)$ defines a point $\epsilon([(D, \boldsymbol{x}), \boldsymbol{a}]) = \boldsymbol{x} + \sum_{i=1}^n a_i \mathbf{1}_{D[i]}$. One checks easily that ϵ is a pospace isomorphism of $\overrightarrow{Geom}(K)$ onto the union of the cubes $\rho(D, \boldsymbol{x})$, $(D, \boldsymbol{x}) \in K$. This observation is needed to relate the notions of geometric realization of precubical *sets* (as used, say, in [5]) and of cubical *complexes* (as used in [14]).

6 Discrete Natural Homology of Cubical Complexes

Paralleling the notion of trace in a pospace, for example as in [5], there is a notion of *discrete trace* in a precubical set K. Given $a, b \in K$, say that a is a *past boundary* of b if and only if $a = \partial_{i_0}^0 \partial_{i_1}^0 \cdots \partial_{i_k}^0 b$ for some $k \geq 0$, i_0, i_1, …, i_k. For example, the edge a, the edge from s to s', and s, are past boundaries of A in the matchbox. *Future boundaries* are defined similarly, using the superscript 1 instead of 0: so the edge from u to u', the edge from s' to u', and u' itself, are future boundaries of A. We write $a \preceq b$ if and only if a is a past boundary of b or b is a future boundary of a. (Beware that this is not a transitive relation; we write $\preceq^*$ for its reflexive transitive closure.) A *discrete trace* from a to b in K is then a sequence $c_0 = a \preceq c_1 \preceq c_2 \preceq \cdots \preceq c_n = b$, $n \in \mathbb{N}$.

Abusing the Fc_X notation we used earlier for pospaces, let Fc_K be the small category whose objects are discrete traces. Its morphisms from a discrete trace from a to b to a discrete trace from a' to b' are the *discrete extensions*, namely pairs of discrete traces α from a' to a and β from b to b'. This is the factorization category of the small category whose objects are elements of K, and whose morphisms are discrete traces.

Note that we are not restricting a, b to be points, namely, of dimension 0; however, it is helpful to imagine, geometrically, that a full cube a stands for the point at its center. The construction is again due to Fajstrup [5]. Formally, for $a = (D, \boldsymbol{x})$, $n = |D|$, let $\hat{a}$ be the point $[a, \bullet]$ in $\overrightarrow{Geom}(K)$, where $\bullet = (\frac{1}{2}, \frac{1}{2}, \cdots, \frac{1}{2})$ is the center of the standard cube $\overrightarrow{I}^n$. Through the ϵ isomorphism, $\hat{a}$ is the point $\boldsymbol{x} + \sum_{i=1}^n \frac{1}{2} \mathbf{1}_{D[i]}$ in $\mathbb{R}^d$, the center of the cube $\rho(D, \boldsymbol{x})$.

Every discrete trace α from a to b, say of the form $c_0 = a \preceq c_1 \preceq c_2 \preceq \cdots \preceq c_n = b$, defines a trace $\hat{\alpha}$ from $\hat{a}$ to $\hat{b}$, obtained by concatenating the n straight lines $\widehat{c_0 c_1}$, $\widehat{c_1 c_2}$, …, $\widehat{c_{n-1} c_n}$. For a simple example, consider the cubical complex whose geometric realization is shown on Figure 3, left. There is a discrete trace α equal to $b \preceq A \preceq t'$, since $b = \partial_1^0 A$ is a past boundary of A and $t' = \partial_2^1 \partial_1^1 A$ is a future boundary of A. The corresponding trace $\hat{\alpha}$ is shown on the same

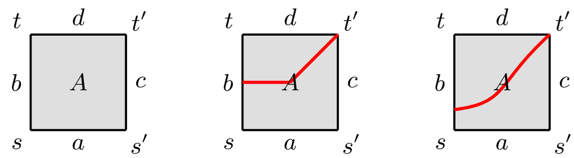

Fig. 3: From discrete traces to traces and vice versa

figure, middle. Formally, if c_{i-1} is a past boundary $\partial^0_{i_1}\partial^0_{i_2}\cdots\partial^0_{i_k}c_i$ of c_i, then $\hat{c}_{i-1} = [\partial^0_{i_1}\partial^0_{i_2}\cdots\partial^0_{i_k}c_i, \bullet] = [c_i, \boldsymbol{a}]$ where $\boldsymbol{a} = \delta^0_{i_k}\cdots\delta^0_{i_2}\delta^0_{i_1}\bullet$; define the dipath π by $\pi(t) = [c_i, (1-t)\boldsymbol{a} + t\bullet]$ for $t \in [0,1]$, and the trace $\widehat{c_{i-1}c_i}$ as $\langle\pi\rangle$. Similarly for future boundaries.

This allows us to transfer cubes a to points $\hat{a} \in \overrightarrow{Geom}(K)$, discrete traces α to traces $\hat{\alpha}$ in $\overrightarrow{Geom}(K)$, and also discrete extensions (α, β) to extensions $(\hat{\alpha}, \hat{\beta})$. We can now mimic the natural homology of a pospace in the discrete setting of a cubical complex K: given a discrete trace γ from a to b, let $\overrightarrow{h}_n(K;\gamma)$ be the $(n-1)$st singular homology group of $Tr(\overrightarrow{Geom}(K);\hat{a},\hat{b})$. This defines another natural system $\overrightarrow{h}_n(K)$, this time from Fc_K instead of Fc_X, to **Ab**: the discrete traces γ are mapped to $\overrightarrow{h}_n(K;\gamma)$, and discrete extensions (α,β) are mapped to $H_{n-1}(\langle\hat{\alpha}\star{}_{-}\star\hat{\beta}\rangle)$, mimicking the definition of $\overrightarrow{H}_n$.

For finite K, Raussen [21] shows that singular homology groups of trace spaces such as $Tr(\overrightarrow{Geom}(K);\hat{a},\hat{b})$ are computable, by computing a finite presentation of the trace spaces (a so-called prod-simplicial complex) from which we can compute homology using Smith normal form of matrices. As a consequence:

Proposition 2. *For a cubical complex K, for every $n \geq 1$, for all discrete trace γ of K, the nth discrete natural homology groups $\overrightarrow{h}_n(K;\gamma)$ are computable.*

By construction, the discrete natural homology group $\overrightarrow{h}_n(K;\gamma)$ is equal to the geometric homology group $\overrightarrow{H}_n(\overrightarrow{Geom}(K);\hat{\gamma})$. However (for finite K) the discrete natural homology functor $\overrightarrow{h}_n(K)$ only lists those for the finitely many discrete traces, while $\overrightarrow{H}_n(\overrightarrow{Geom}(K))$ lists one group for each of the uncountably many traces in $\overrightarrow{Geom}(K)$. The discrete functor $\overrightarrow{h}_n(K)$ also has to cater for finitely many discrete extension morphisms, whereas $\overrightarrow{H}_n(\overrightarrow{Geom}(K))$ has to map uncountably many extension morphisms to group homomorphisms. This makes quite a difference—but not one up to bisimilarity:

Theorem 1 (Discrete Nat. Homology≡Geometric Nat. Homology). *For every cubical complex K, there is an open map from the natural system $\overrightarrow{H}_n(\overrightarrow{Geom}(K))$ to the discrete natural system $\overrightarrow{h}_n(K)$. In particular, they are bisimilar.*

Before we describe the construction, notice that there is *no* open map in the other direction: remember that the open maps we consider have a fibration component, which must be surjective.

Proof. We need to define an open map $(\mathcal{C}, \sigma)$ from $\overrightarrow{H}_n(X)$, where $X = \overrightarrow{Geom}(K)$, to $\overrightarrow{h}_n(K)$. We start by building $\mathcal{C}$, which must be a fibration from Fc_X to Fc_K.

This is based on the notion of *carrier sequence* due to Fajstrup [5]. For a point s in $\overrightarrow{Geom}(K)$, there is a unique cube $e \in K$ of minimal dimension m such that s can be written as $[e, \boldsymbol{a}]$, $\boldsymbol{a} \in \overrightarrow{I}^m$. Write $\mathcal{C}(s)$ for this cube e, and call it the *carrier* of s. Every trace $\langle\pi\rangle$ in X gives rise to an ordered sequence of cubes $\mathcal{C}(\langle\pi\rangle)$ obtained as the carriers of $\pi(t)$, $t \in [0, 1]$, and removing consecutive duplicates. This is formally defined in [5]. By a compactness argument $\mathcal{C}(\langle\pi\rangle)$ is a finite sequence, in fact a discrete trace, called the *carrier sequence* of $\langle\pi\rangle$. For example, the carrier sequence of the trace on the right of Figure 3 is $b \preceq A \preceq t'$.

We use this to define our functor $\mathcal{C}$, on objects by letting $\mathcal{C}(\langle\pi\rangle)$ be defined as above, and on morphisms by letting $\mathcal{C}(\langle\alpha\rangle, \langle\beta\rangle) = (\mathcal{C}(\langle\alpha\rangle), \mathcal{C}(\langle\beta\rangle))$ for every extension $(\langle\alpha\rangle, \langle\beta\rangle)$. This is surjective on objects since $\mathcal{C}(\hat{\gamma}) = \gamma$ for every discrete trace γ. We now claim that $\mathcal{C}$ is a fibration, and this amounts to show that: given any trace $\langle\pi\rangle$ of $\overrightarrow{Geom}(K)$, with carrier sequence $c_0 \preceq \cdots \preceq c_k$, if the latter extends to a discrete trace $c_{-p} \preceq \cdots \preceq c_{-1} \preceq c_0 \preceq \cdots \preceq c_k \preceq c_{k+1} \preceq \cdots \preceq c_{k+q}$ in K, then $\langle\pi\rangle$ extends to some trace $\langle\alpha \star \pi \star \beta\rangle$ such that $\mathcal{C}(\langle\alpha \star \pi \star \beta\rangle) = c_{-p} \preceq \cdots \preceq c_{-1} \preceq c_0 \preceq \cdots \preceq c_k \preceq c_{k+1} \preceq \cdots \preceq c_{k+q}$. By induction, the cases $(p, q) = (1, 0)$ and $(p, q) = (0, 1)$ suffice to establish the property. Some care has to be taken: the extension paths are *not* concatenations of simple straight lines joining the extra points $\hat{c}_j$, $j \geq k$ or $j \leq 0$. As the picture on the right shows (for $(p, q) = (0, 2)$), the dipath β does not—and cannot—go through $\hat{c}_3$. Details of the construction are given in Appendix C, Lemma 1.

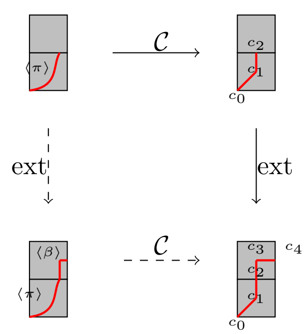

We now need to build a natural isomorphism $\sigma : \overrightarrow{H}_n(X) \longrightarrow \overrightarrow{h}_n(K) \circ \mathcal{C}$. In other words, we need to build group isomorphisms $\sigma_{\langle\pi\rangle} : \overrightarrow{H}_n(X; \langle\pi\rangle) \longrightarrow \overrightarrow{h}_n(K; \mathcal{C}(\langle\pi\rangle))$ that are natural, in the sense that, for every extension $(\langle\alpha\rangle, \langle\beta\rangle)$ of $\langle\pi\rangle$, and for $(\gamma, \delta) = \mathcal{C}(\langle\alpha\rangle, \langle\beta\rangle)$ the associated discrete extension, the following square commutes:

$$\begin{array}{ccc}
\overrightarrow{H}_n(X; \langle\pi\rangle) & \xrightarrow{\sigma_{\langle\pi\rangle}} & \overrightarrow{h}_n(K; \mathcal{C}(\langle\pi\rangle)) \\
{\scriptstyle \langle\alpha \star _ \star \beta\rangle}\downarrow & & \downarrow{\scriptstyle \langle\hat{\gamma} \star _ \star \hat{\delta}\rangle} \\
\overrightarrow{H}_n(X; \langle\alpha \star \pi \star \beta\rangle) & \xrightarrow[\sigma_{\langle\alpha\star\pi\star\beta\rangle}]{} & \overrightarrow{h}_n(K; \mathcal{C}(\langle\alpha \star \pi \star \beta\rangle))
\end{array}$$

Let π be from s to t. Every cube I^k has a lattice structure whose meet $\wedge$ is pointwise min and whose join $\vee$ is pointwise max. Write s as $[\mathcal{C}(s), \boldsymbol{a}]$, and let $s_- = [\mathcal{C}(s), \boldsymbol{a} \wedge \bullet]$. Recall that $\bullet = (\frac{1}{2}, \cdots, \frac{1}{2})$, and that $\widehat{\mathcal{C}(s)} = [\mathcal{C}(s), \bullet]$. Similarly, let $\widehat{\mathcal{C}(t)} = [\mathcal{C}(t), \bullet]$, and we define $t_+ = [\mathcal{C}(t), \boldsymbol{b} \vee \bullet]$, where $t = [\mathcal{C}(t), \boldsymbol{b}]$. The situation is illustrated in the two gray boxes to the right.

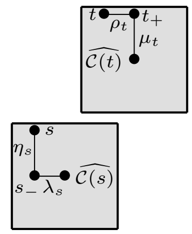

There are obvious dipaths η_s, λ_s, μ_t, ρ_t as displayed there, too. Those induce continuous maps between trace spaces by concatenation.
For example, there is a continuous map $\eta_s^* : Tr(X; s, t) \longrightarrow Tr(X; s_-, t)$ that sends each trace $\langle \pi' \rangle$ to $\langle \eta_s \star \pi' \rangle$. Similarly, $\lambda_s^*(\langle \pi' \rangle) = \langle \lambda_s \star \pi' \rangle$, and symmetrically, ${}^*\mu_t(\langle \pi' \rangle) = \langle \pi' \star \mu_t \rangle$, ${}^*\rho_t(\langle \pi' \rangle) = \langle \pi' \star \rho_t \rangle$. We show in Appendix D that each of these four maps is a homotopy equivalence, and therefore induce isomorphisms in homology. It remains to define $\sigma_{\langle \pi \rangle}$ as the composition $H_{n-1}({}^*\mu_t)^{-1} \circ H_{n-1}({}^*\rho_t) \circ H_{n-1}(\lambda_s^*)^{-1} \circ H_{n-1}(\eta_s^*)$ of those four isomorphisms. Naturality is, as usual, tedious but mechanical. □

The potential problem mentioned at the beginning of Section 4 is then solved: the uncountable natural homology of $\overrightarrow{Geom}(K)$ is reduced, through bisimilarity, to the finite, discrete natural homology of K.

A *dihomeomorphism* is a continuous monotonic bijection between pospaces whose inverse is also continuous and monotonic.

Corollary 1 (Invariance under dihomeomorphism). *For any cubical complexes K, K' whose geometric realizations are dihomeomorphic, $\overrightarrow{H}_n(\overrightarrow{Geom}(K))$ and $\overrightarrow{H}_n(\overrightarrow{Geom}(K'))$ are isomorphic, and $\overrightarrow{h}_n(K)$ and $\overrightarrow{h}_n(K')$ are bisimilar.*

Of particular importance to the field of true concurrency is invariance under refinement [23]. In our case, this means that if we replace certain n-cubes in K by unions of 2^n smaller cubes with all dimensions halved, then the result should have the same natural homology. Indeed, such a process is called *subdivision* in the literature, and it is well-known that if K' is a subdivision of K, then $\overrightarrow{Geom}(K)$ and $\overrightarrow{Geom}(K')$ are dihomeomorphic. Hence:

Corollary 2 (Invariance under subdivision). *Let K be a cubical complex, and K' be a subdivision of K. Then $\overrightarrow{h}_n(K)$, and $\overrightarrow{h}_n(K')$ are bisimilar.*

7 Conclusion

We have defined a promising notion of homology for directed algebraic topology. We have shown that our natural systems of homology are computable on cubical complexes. We have also introduced a notion of bisimilarity with respect to which those natural systems should be compared. Importantly, natural homology is invariant under subdivision. We showed this as a special case of a more general result: that the natural homology of a cubical complex is bisimilar to that of its geometric realization.

As a litmus test, does our natural homology pass the criteria we set forth in Section 2? Look again at Fahrenberg's matchbox (Fig. 1). Its discrete natural homology would be too big to fit on a page, however its $\overrightarrow{H}_1$ at the trace $a \star c$ is equal to $\mathbb{Z}^2$. In particular, it has non-trivial natural homology, in the strong sense that its natural homology is not bisimilar to any natural system consisting only of copies of $\mathbb{Z}$ (e.g., the natural homology of a filled-out cube). This is the first proposal that distinguishes the matchbox from a trivial pospace.

References

1. H.-J. Baues and G. Wirsching. Cohomology of small categories. *Journal of Pure and Applied Algebra*, 38(2-3):187–211, 1985.
2. G. Carlsson. Topology and data. *AMS Bulletin*, 46(2):255–308, 2009.
3. E. G. Coffman, M. J. Elphick, and A. Shoshani. System deadlocks. *Computing Surveys*, 3(2):67–78, 1971.
4. U. Fahrenberg. Directed homology. *Electronic Notes in Theoretical Computer Science*, 100:111–125, 2004.
5. L. Fajstrup. Dipaths and dihomotopies in a cubical complex. *Advances in Applied Mathematics*, 35(2):188–206, 2005.
6. L. Fajstrup, E. Goubault, E. Haucourt, S. Mimram, and M. Raußen. Trace Spaces: An Efficient New Technique for State-Space Reduction. In *21st European Symposium on Programming (ESOP)*, 2012.
7. E. Goubault. *Géométrie du parallélisme*. PhD thesis, Ecole Polytechnique, 1995.
8. E. Goubault. Geometry and concurrency: A user's guide. *Mathematical Structures in Computer Science*, 10(4):411–425, 2000.
9. E. Goubault and E. Haucourt. A Practical Application of Geometric Semantics to Static Analysis of Concurrent Programs. In *16th Intl. Conf. Concurrency Theory (CONCUR)*, pages 503–517, 2005.
10. E. Goubault, T. Heindel, and S. Mimram. A geometric view of partial order reduction. *Electronic Notes in Theoretical Computer Science*, 298, 2013.
11. M. Grandis. Inequilogical spaces, directed homology and noncommutative geometry. *Homology, Homotopy and Applications*, 6:413–437, 2004.
12. M. Grandis. *Directed Algebraic Topology, Models of non-reversible worlds*. Cambridge University Press, 2009.
13. A. Joyal, M. Nielsen, and G. Winskel. Bisimulation from open maps. *Information and Computation*, 127(2):164–185, 1996.
14. T. Kaczynski, K. Mischaikow, and M. Mrozek. Computing homology. *Homology, Homotopy and Applications*, 5(2):233–256, 2003.
15. T. Kahl. The homology graph of a HDA. arXiv:1307.7994, 2013.
16. S. Krishnan. Flow-cut dualities for sheaves on graphs. arXiv:1409.6712, 2014.
17. A. Patchkoria. On exactness of long sequences of homology semimodules. *Journal of Homotopy and Related Structures*, 1(1):229–243, 2006.
18. V. R. Pratt. Modeling Concurrency with Geometry. In D. S. Wise, editor, *18th Ann. ACM Symp. Principles of Programming Languages*, pages 311–322, 1991.
19. M. Raussen. Invariants of directed spaces. *Applied Categorical Structures*, 15, 2007.
20. M. Raussen. Trace spaces in a pre-cubical complex. *Topology and its Applications*, 156(9):1718–1728, 2009.
21. M. Raussen. Simplicial models for trace spaces II: General higher dimensional automata. *Algebraic and Geometric Topology*, 12(3):1741–1762, 2012.
22. F. Sergeraert. The computability problem in algebraic topology. *Advances in Mathematics*, pages 1–29, 1994.
23. R. van Glabeek. *Comparative concurrency semantics and refinement of actions*. PhD thesis, Centrum voor Wiskunder en Informatica, 1990.

A Bisimulations and open maps

It is useful to realize that our notion of bisimulation arises from the standard construction of bisimulations from open maps [13]. Here is how.

Let $\mathbf{C}$ be the category whose objects are functors from a small category X (which may vary) to the category $\mathbf{Ab}$ of abelian groups, and whose morphisms from $H : E \longrightarrow \mathbf{Ab}$ to $F : X \longrightarrow \mathbf{Ab}$ are pairs (Φ, σ), where Φ is a functor from E to X and σ is a natural isomorphism from H to $F \circ \Phi$. For every $n \in \mathbb{N}$, let $[n] = \{0, 1, \cdots, n-1\}$, and let $i_{mn} : [m] \to [n]$ be the inclusion map, $m \leq n$. As a poset, $[n]$ is a category, and i_{mn} is then a functor.

Consider the subcategory $\mathbf{P}$ (of so-called paths in that theory) whose objects are functors $F : [n] \longrightarrow \mathbf{Ab}$ and whose morphisms from $H : [m] \longrightarrow \mathbf{Ab}$ to $F : [n] \longrightarrow \mathbf{Ab}$ are of the form (i_{mn}, σ) for some σ.

The *open maps* with respect to $\mathbf{P}$ were defined by Joyal *et al.* [13] as those morphisms in $\mathbf{C}$ that have the right lifting property with respect to morphisms in $\mathbf{P}$. The latter means that in every commuting diagram made of solid lines below, there is a lifting (the dotted arrow) that makes the two triangles commute:

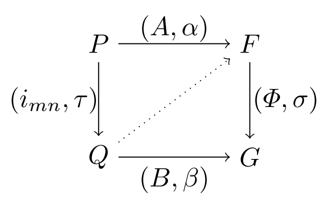

where $P : [m] \longrightarrow \mathbf{Ab}$, $Q : [n] \longrightarrow \mathbf{Ab}$, $F : E \longrightarrow \mathbf{Ab}$, and $G : X \longrightarrow \mathbf{Ab}$.

Proposition 3. *The open maps are exactly the morphisms* (Φ, σ) *of* $\mathbf{C}$ *where* Φ *is a fibration (and* σ *is a natural isomorphism of abelian groups).*

Our notion of a fibration $\Phi : E \longrightarrow X$ is most handily summed up by requiring that Φ is surjective on objects and the following lifting diagram is satisfied:

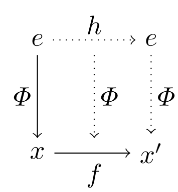

This reads: for every object e of E, every morphism $f : \Phi(e) \longrightarrow x'$ in X lifts to a morphism $h : e \longrightarrow e'$ in E such that $\Phi(h) = f$.

Proof. – Assume (Φ, σ) is open in the sense of [13]. Specialize the above diagram to the case where the four functors P, Q, F, G map every object to the trivial group 0. This has the effect that the natural isomorphisms σ, τ, α, β are all trivial, and can safely be ignored.

Look at the case $m = 0$, $n = 1$. For every object x of X, we may define B so that $B(0) = x$. There is a commuting diagram as above, where A is the empty functor, because $B \circ i_{01}$ and $\Phi \circ A$ are both the empty functor. The existence of a lifting, and notably the fact the the lower triangle commutes, implies the existence of an object e in E such that $\Phi(e) = x$. So Φ is surjective on objects.

Now look at the case $m = 1$, $n = 2$. Given any arrow $f : \Phi(e) \longrightarrow x'$ in X, define B so that it maps the unique arrow from 0 to 1 in [2] to f. In particular, B maps $0 = i_{12}(0)$ to $\Phi(e)$. Define A as mapping 0 to the object e. The lifting must map the unique arrow from 0 to 1 in [2] to a morphism $h : e \longrightarrow e'$ in E such that $\Phi(h) = f$. Therefore Φ is a fibration.

– Conversely, assume a morphism (Φ, σ) in $\mathbf{C}$ such that Φ is a fibration (and σ an isomorphism of groups). We claim that (Φ, σ) has the right lifting property with respect to morphisms (i_{mn}, τ) in $\mathbf{P}$. We assume that $(B, \beta) \circ (i_{mn}, \tau) = (\Phi, \sigma) \circ (A, \alpha)$, and wish to construct a lifting from Q to F.

We do this by induction on $n - m$. If $m = n$, then i_{mn} is the identity map, and we can just define the lifting as $(A, \alpha \circ \tau^{-1})$. Note that it is important that the second component of our morphisms be a group isomorphism.

If $m < n$, then we can factor (i_{mn}, τ) as the composition of $(i_{m(n-1)}, \tau) : P \longrightarrow Q_{|[n-1]}$ with $(i_{(n-1)n}, \mathrm{id}) : Q_{|[n-1]} \longrightarrow Q$, where $Q_{|[n-1]}$ is the restriction of Q to the subcategory $[n-1]$ of $[n]$, and id is the identity natural transformation. By induction hypothesis, we obtain a lifting (C, γ) as in the following diagram, and we wish to build the dotted arrow.

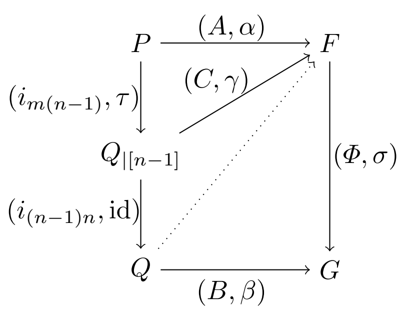

The dotted arrow should be a morphism (D, δ). Again, we look at the D part only, and leave the construction of δ to the end. For every $i \leq n - 1$, we define $D(i) = C(i)$, and, for every morphism $i \to j$ with $i \leq j \leq n - 1$ in $[n]$, we define $D(i \to j) = C(i \to j)$. We must then define $D(n)$ as some object e of E such that $\Phi(e) = B(n)$: this exists because Φ is surjective. We pick any such e. For every morphism $i \to n$ in $[n]$, either $i = n$ and we must set $D(n \to n) = \mathrm{id}$, or $i \leq n - 1$. In the latter case, $D(i \to n)$ will be determined uniquely as the composition of $D(i \to n - 1) = C(i \to n - 1)$ with $D(n - 1 \to n)$. There is no constraint on $D(n - 1 \to n)$ except that it must be a morphism h in E such that $\Phi(h) = B(n - 1 \to n)$. By induction hypothesis, $e = C(n - 1)$ is an object such that $\Phi(e) = B(n - 1)$, and since

Φ is a fibration, there is a morphism h in E such that $\Phi(h) = B(n-1 \to n)$: this is what we were looking for.

Finally, we construct δ. That means finding a group isomorphism $\delta(i)$ for every $i \leq n$ between $Q(i)$ and $\Phi(D(i))$ such that $\sigma(D(i)) \circ \delta(i) = \beta(i)$ for every $i \leq n$ (lower triangle) and $\delta(i) \circ \tau(i) = \alpha(i)$ for every $i \leq m$ (upper triangle). This forces us to define $\delta(i)$ as $\sigma(D(i))^{-1} \circ \beta(i)$ for every $i \leq n$, using the fact that σ is a natural *isomorphism*. The equation $\delta(i) \circ \tau(i) = \alpha(i)$ is then automatic for every $i \leq m$. And the naturality of δ follows from the naturality of σ^{-1} and β. □

The fact that being related by a span of fibrations is the same thing as being bisimilar in our sense is the topic of Proposition 1, and is proved in Appendix B.

B Proof of Proposition 1

Assume that $F : X \longrightarrow \mathbf{Ab}$ and $G : Y \longrightarrow \mathbf{Ab}$ are bisimilar, i.e., there exists a bisimulation R between F and G. We construct a span of open maps as follows.

Let E be the small category whose objects are elements of R, and whose morphisms from (x, η, y) to (x', η', y') are pairs (i, j) of a morphism $i : x \longrightarrow x'$ in X and of a morphism $j : y \longrightarrow y'$ in Y, such that the following diagram commutes:

$$\begin{array}{ccc} Fx & \xrightarrow{\ \eta\ } & Gy \\ {\scriptstyle Fi}\downarrow & & \downarrow{\scriptstyle Gj} \\ Fx' & \xrightarrow[\ \eta'\ ]{} & Gy' \end{array}$$

Define the tip H of the span between F and G as the functor $H : E \longrightarrow \mathbf{Ab}$ that maps every object $(x, \eta, y) \in R$ to Fx, and every morphism $(i, j) : (x, \eta, y) \longrightarrow (x', \eta', y')$ to $Fi : Fx \longrightarrow Fx'$.

We now build a morphism (Φ, σ) from H to F. We start by building $\Phi : E \longrightarrow X$. We define Φ as the functor that maps every object (x, η, y) to x and every morphism $(i, j) : (x, \eta, y) \longrightarrow (x', \eta', y')$ to $i : x \longrightarrow x'$. We verify that Φ is a fibration:

1. Φ is surjective on objects: this is condition 1 of the definition of R as a bisimulation.
2. Let $f : \Phi(e) \longrightarrow x'$ be a morphism of X. The object e must be a triple $(x, \eta, y) \in R$, and f is a morphism from x to x' in X. By condition 2 of the definition of R as a bisimulation, there is a triple $(x', \eta', y') \in R$ and a morphism $j : y \longrightarrow y'$ of Y such that the following diagram commutes:

$$\begin{array}{ccc} Fx & \xrightarrow{\eta} & Gy \\ {\scriptstyle Fi}\downarrow & & \downarrow{\scriptstyle Gj} \\ Fx' & \xrightarrow[\eta']{} & Gy' \end{array}$$

In particular, (i, j) is a morphism of E, from (x, η, y) to (x', η', y'). Moreover, $H(i, j) = i$.

For every $(x, \eta, y) \in R$, let $\sigma_{(x,\eta,y)} = \mathrm{id}_{Fx} : H(x, \eta, y) = Fx \longrightarrow F \circ \Phi(x, \eta, y) = F(x)$. Those are isomorphisms, and define a natural transformation $\sigma : H \longrightarrow F \circ \Phi$. It follows that (Φ, σ) is an open map from H to F.

We define the open map (Ψ, τ) from H to G similarly. Hence F and G are **P**-bisimilar in the sense of Joyal *et al.* [13].

Conversely, assume that F and G are bisimilar in the sense of Joyal *et al.* [13]. There is a span of open maps:

$$\begin{array}{ccccc} & & H & & \\ & {\scriptstyle (\Phi,\sigma)}\swarrow & & \searrow{\scriptstyle (\Psi,\tau)} & \\ F & & & & G \end{array}$$

with $H : E \longrightarrow \mathbf{Ab}$.

Consider the set R of triples $(\Phi e, \tau_e \circ \sigma_e^{-1}, \Psi e)$ with e an object of E. This is a set because the category E is small. Let us show that R is a bisimulation:

1. is a consequence of the fact that Φ and Ψ are surjective on objects.
2. Let $(\Phi e, \tau_e \circ \sigma_e^{-1}, \Psi e) \in R$ and $i : \Phi e \longrightarrow x'$ be a morphism in X. Since Φ is a fibration, there is a morphism $h : e \longrightarrow e'$ such that $\Phi h = f$, and in particular $\Phi e' = x'$. By construction, $(\Phi e', \tau_{e'} \circ \sigma_{e'}^{-1}, \Psi e')$ is in R and $\Psi h : \Psi e \longrightarrow \Psi e'$. It is sufficient to prove that :

$$\begin{array}{ccc} F \circ \Phi e & \xrightarrow{\tau_e \circ \sigma_e^{-1}} & G \circ \Psi e \\ {\scriptstyle F \circ \Phi h}\downarrow & & \downarrow{\scriptstyle G \circ \Psi h} \\ F \circ \Phi e' & \xrightarrow[\tau_{e'} \circ \sigma_{e'}^{-1}]{} & G \circ \Psi e' \end{array}$$

commutes. This is just the naturality diagram for $\tau \circ \sigma^{-1}$. The second part of condition 2 is symmetric.

It follows that R is a bisimulation, hence F and G are bisimilar. □

C The functor $\mathcal{C}$ is a fibration

In this Section, we assume that K is a cubical complex, and X is its geometric realization $\overrightarrow{Geom}(K)$.

Let us recall the fundamental properties of the carrier sequence, as defined by Fajstrup [5]. Given a dipath π of X, there is a unique sequence $c_0, c_1, \cdots, c_k$ of elements of K and a unique sequence of real numbers $0 = t_0 \leq t_1 \leq \cdots \leq t_k \leq t_{k+1} = 1$ (call them the *times of change*) such that:

- for every $1 \leq i \leq k$, $c_{i-1} \neq c_i$,
- for every $0 \leq i \leq k$, for every $t \in [t_i, t_{i+1}]$, $\pi(t)$ is a point of the form $[c, \boldsymbol{a}]$ with $c = c_i$,
- for every $0 \leq i \leq k$, for every $t \in (t_i, t_{i+1})$, $\mathcal{C}(\pi(t)) = c_i$,
- $\mathcal{C}(\pi(0)) = c_0$ and $\mathcal{C}(\pi(1)) = c_k$,
- for every $1 \leq i \leq k$, $\mathcal{C}(\pi(t_i)) \in \{c_{i-1}, c_i\}$ and if furthermore $t_i = t_{i+1}$ then $\mathcal{C}(\pi(t_i)) = c_i$.

The sequence $c_0, c_1, \cdots, c_k$ is the *carrier sequence of* π. Two dipaths that are equivalent modulo reparametrization have the same carrier sequence, so it is legitimate to call carrier sequence of a trace $\langle\pi\rangle$ the carrier sequence $\mathcal{C}(\pi)$ of π.

We now prove that $\mathcal{C}$ is a fibration. We have seen that $\mathcal{C}$ is surjective on objects, and we want to show that given any trace $\langle\pi\rangle$ in X, with carrier sequence $c_0 \preceq \cdots \preceq c_k$, if the latter extends to a discrete trace $c_{-p} \preceq \cdots \preceq c_{-1} \preceq c_0 \preceq \cdots \preceq c_k \preceq c_{k+1} \preceq \cdots \preceq c_{k+q}$ in K, then $\langle\pi\rangle$ extends to some trace $\langle\alpha\star\pi\star\beta\rangle$ such that $\mathcal{C}(\langle\alpha\star\pi\star\beta\rangle) = c_{-p} \preceq \cdots \preceq c_{-1} \preceq c_0 \preceq \cdots \preceq c_k \preceq c_{k+1} \preceq \cdots \preceq c_{k+\ell}$. The following lemma establishes the cases $p = 1, q = 0$ and $p = 0, q = 1$. Induction on p then q allows us to obtain the general result.

Lemma 1. *Let $\langle\pi\rangle$ be a trace in X ($= \overrightarrow{Geom}(K)$) with carrier sequence $c_0 \preceq c_1 \preceq \cdots \preceq c_k$.*

- *For every cube $c_{-1} \preceq c_0$, there is a dipath α in X such that $\mathcal{C}(\langle\alpha\star\pi\rangle) = c_{-1} \preceq c_0 \preceq c_1 \preceq \cdots \preceq c_k$.*
- *For every cube c_{k+1} such that $c_k \preceq c_{k+1}$, there is a dipath β in X such that $\mathcal{C}(\langle\pi\star\beta\rangle) = c_0 \preceq c_1 \preceq \cdots \preceq c_k \preceq c_{k+1}$.*

Proof. We examine the second case only: the other case is symmetric. Since $c_k \preceq c_{k+1}$, c_k can be a past boundary of c_{k+1}, or c_{k+1} can be a future boundary of c_k. We examine both cases:

- If c_k is a past boundary of c_{k+1}, say $c_k = \partial^0_{i_p} \cdots \partial^0_{i_0} c_{k+1}$, then by using the precubical equations we may require $i_0 > \ldots > i_p$. Writing $\pi(1)$ as $[c_k, \boldsymbol{a}]$, we also have $\pi(1) = [c_{k+1}, \delta^0_{i_0} \cdots \delta^0_{i_p} \boldsymbol{a}]$ by the definition of the geometric realization. Since $\mathcal{C}(\pi(1)) = c_k$, no component a_i of $\boldsymbol{a}$ is equal to 0 or 1. Let $\boldsymbol{b} = \delta^0_{i_0} \cdots \delta^0_{i_p} \boldsymbol{a}$: it follows that the components b_i of $\boldsymbol{b}$ that are equal to 0 are exactly those such that $i \in \{i_0, \cdots, i_p\}$. Let $\boldsymbol{a'}$ be the tuple whose ith component a'_i is $1/2$ if $b_i = 0$, and b_i otherwise. We define the dipath β by

$\beta(t) = [c_{k+1}, (1-t)\boldsymbol{b} + t\boldsymbol{a'})]$, $t \in [0,1]$. Note that β is indeed monotonic, because $b_i \leq a'_i$ for every i. One easily checks that $\beta(0) = \pi(1)$, and that the carrier sequence of $\langle\beta\rangle$ is $c_k \preceq c_{k+1}$: for $t = 0$, $\mathcal{C}(\beta(0)) = (\pi(1)) = c_k$, and, for $t \neq 0$, $\beta(t) = [c_{k+1}, (1-t)\boldsymbol{b} + t\boldsymbol{a'})]$ where no component of $(1-t)\boldsymbol{b} + t\boldsymbol{a'}$ is equal to 0 or 1, so its carrier $\mathcal{C}(\beta(t))$ is c_{k+1}. It follows that $\mathcal{C}(\langle\pi \star \beta\rangle) = c_0 \preceq c_1 \preceq \cdots \preceq c_k \preceq c_{k+1}$.

– If c_{k+1} is a future boundary of c_k, then c_{k+1} is of the form $\partial^1_{i_p} \ldots \partial^1_{i_0} c_k$ with $i_0 > \ldots > i_p$, and $\pi(1) = [c_k, \boldsymbol{a}]$ for some tuple $\boldsymbol{a}$ whose components a_i are all different from 0 or 1 (because $\mathcal{C}(\pi(1)) = c_k$). Let $\boldsymbol{b}$ be the tuple obtained from $\boldsymbol{a}$ by changing the ith component into 1 if and only if $i \in \{i_0, \cdots, i_p\}$. In other words, let $b_i = 1$ if $i \in \{i_0, \cdots, i_p\}$, $b_i = a_i$ otherwise. One can therefore write $\boldsymbol{b}$ as $\delta^1_{i_0} \cdots \delta^1_{i_p} \boldsymbol{b'}$, where $\boldsymbol{b'}$ is the tuple obtained from $\boldsymbol{b}$ by removing its components of indices $i_0, \ldots, i_p$. Define the dipath β by $\beta(t) = [c_k, (1-t)\boldsymbol{a} + t\boldsymbol{b}]$. This is monotonic because $a_i \leq b_i$ for every i. For $t \neq 1$, no component of $(1-t)\boldsymbol{a} + t\boldsymbol{b}$ is equal to 0 or 1, so $\mathcal{C}(\beta(t)) = c_k$, and for $t = 1$, $\beta(1) = [c_k, \boldsymbol{b}] = [c_{k+1}, \boldsymbol{b'}]$, which shows that $\mathcal{C}(\beta(1)) = c_{k+1}$ since no component of $\boldsymbol{b'}$ is equal to 0 or 1. Again, it follows that $\mathcal{C}(\langle\pi \star \beta\rangle) = c_0 \preceq c_1 \preceq \cdots \preceq c_k \preceq c_{k+1}$. ☐

D The construction of the natural isomorphism $\boldsymbol{\sigma}$

Let us make formal the construction of the dipath η_s. The other three are similar. This is a dipath from $s_- = [\mathcal{C}(s), \boldsymbol{a} \wedge \bullet]$ to $s = [\mathcal{C}(s), \boldsymbol{a}]$, and so we just let $\eta_s(t) = [\mathcal{C}(s), (1-t)(\boldsymbol{a} \wedge \bullet) + t\boldsymbol{a}]$.

Recall that η_s^* maps $\langle\pi\rangle \in Tr(X; s, t)$, to $\langle\eta_s \star \pi\rangle \in Tr(X; s_-, t)$.

Lemma 2. *The map η_s^* is a homotopy equivalence.*

Proof. By abuse of language, write $\eta_s^*(\pi)$ for the dipath $\eta_s \star \pi$ as well—we reason on spaces of dipaths first, then take a reparametrization quotient. Accordingly, let $P(X; s, t)$ denote the space of dipaths from s to t in $X = \overrightarrow{Geom}(K)$, with the usual compact-open topology. (The space $Tr(X; s, t)$ is a quotient of this space.)

Observe that η_s^* maps $P(X; s, t)$ to $P(X; s_-, t)$. We need to build a map $\nu : P(X; s_-, t) \longrightarrow P(X; s, t)$ such that $\eta_s^* \circ \nu$ and $\nu \circ \eta_s^*$ are homotopic to the identity.

For every dipath π from s to t, the carrier sequence $c_0, c_1, \cdots, c_k$ of $\eta_s^*(\pi)$ is equal to that of π. In the other direction, we shall define ν so that it also preserves the carrier sequence. This will turn out to be the crucial property that will allow us to conclude.

For every dipath π from s_- to t, with carrier sequence $c_0, c_1, \cdots, c_k$, and with times of change $0 = t_0 \leq t_1 \leq \cdots \leq t_k \leq t_{k+1} = 1$ (see Appendix C), we define $\nu(\pi)$ as follows. We abuse the notation $\vee$, and write $[c, \boldsymbol{a}] \vee_c [c, \boldsymbol{b}]$ for $[c, \boldsymbol{a} \vee \boldsymbol{b}]$. The three occurrences of c must be the same for this notation to make sense, but our intuition is best served by ignoring the c subscript to $\vee$, and to understand this as taking maxes, componentwise, in a local cube c. We then define $\nu(\pi)(u)$ for increasing values of u, inductively, as $s \vee_{c_0} \pi(u)$ for $u \in [t_0, t_1]$, as $\nu(\pi)(t_1) \vee_{c_1} \pi(u)$ for $u \in [t_1, t_2]$, $\ldots$, and finally as $\nu(\pi)(t_k) \vee_{c_k} \pi(u)$ for $u \in [t_k, t_{k+1}]$.

On $[t_0, t_1]$, $\nu(\pi)$ is a continuous monotonic map, with value $\nu(\pi)(0) = s \vee_{c_0} s_- = s$ at $u = t_0 = 0$, and with value $\nu(\pi)(t_1) = s \vee_{c_0} \pi(t_1)$ at $u = t_1$.

Let us show by induction on j that for every u with $0 \leq u \leq t_j$, $\mathcal{C}(\nu(\pi)(u)) = \mathcal{C}(\pi(u))$. For $j = 0$, this says that $\mathcal{C}(s) = \mathcal{C}(s_-)$, which is by construction of s_-. Otherwise, by induction hypothesis, for every u with $0 \leq u \leq t_j$, $\mathcal{C}(\nu(\pi)(u)) = \mathcal{C}(\pi(u))$. Let $t_j < u \leq t_{j+1}$. We can write $\pi(t_j)$ as $[c_j, (b_1, \ldots, b_m)]$ and $\pi(u)$ as $[c_j, (a_1, \ldots, a_m)]$, where $b_i \leq a_i$ for every i.

- If $u < t_{j+1}$, by the properties of the carrier sequence, $\mathcal{C}(\pi(u)) = c_j$, so with $0 < a_i < 1$ for every i. Since $b_i \leq a_i$, $b_i < 1$ for every i. Let us write $\nu(\pi)(t_j)$ as $[c_j, (b'_1, \ldots, b'_m)]$. Since $\mathcal{C}(\nu(\pi)(t_j)) = \mathcal{C}(\pi(t_j))$, $b_i = 1$ iff $b'_i = 1$. It follows that $b'_i < 1$ for every i. Therefore $0 < \max(a_i, b'_i) < 1$, so $\mathcal{C}(\nu(\pi)(u)) = c_k$.
- If $u = t_{j+1}$, we observe that $\max(a_i, b'_i)$ is equal to 1, resp. to 0, resp. in $(0, 1)$, if and only if a_i is. This observation is enough to conclude that $\mathcal{C}(\nu(\pi)(t_{j+1})) = \mathcal{C}(\pi(t_{j+1}))$, and is proved as follows. If $a_i = 1$, then $\max(a_i, b'_i) = 1$. If $a_i = 0$ then $b_i = 0$; moreover, since $\mathcal{C}(\nu(\pi)(t_j)) = \mathcal{C}(\pi(t_j))$, $b_i = 0$ iff $b'_i = 0$, so $b'_i = 0$, from which we obtain $\max(a_i, b'_i) = 0$. Finally, if $0 < a_i < 1$ then $b_i < 1$, and $b'_i < 1$ (since $\mathcal{C}(\nu(\pi)(t_j)) = \mathcal{C}(\pi(t_j))$, $b_i = 1$ iff $b'_i = 1$), so $0 < \max(a_i, b'_i) < 1$.

This finishes our argument that $c_0, \ldots, c_k$ is the carrier sequence of $\nu(\pi)$, with times of change $0 = t_0 \leq \ldots \leq t_{k+1} = 1$.

It remains to show that $\nu(\pi)(1) = t$. This is the only place where we need the ϵ mapping. The above argument works in general precubical sets, not just cubical complexes. On the contrary, we need the specific features of cubical complexes to show that $\nu(\pi)(1) = t$. We discuss this in a remark at the end of the section.

We know that $\mathcal{C}(t) = \mathcal{C}(\nu(\pi)(1)) = c_k$. Moreover, t is below $\nu(\pi)(1)$ in the ordering $\leq$ of the pospace $X = \overrightarrow{Geom}(K)$, because $\nu(\pi)(1) = \nu(\pi)(t_k) \vee_{c_k} \pi(1) = \nu(\pi)(t_k) \vee_{c_k} t$. Suppose that $\nu(\pi)(1) \not\leq t$. Because K is a cubical complex, we can make use of the ϵ isomorphism. From $\nu(\pi)(1) \not\leq t$, we obtain $\epsilon(\nu(\pi)(1)) \not\leq \epsilon(t)$. Let us write $\epsilon(\nu(\pi)(t_j))$ as $(x^j_1, \ldots, x^j_d)$ and $\epsilon(\pi(t_j))$ as $(y^j_1, \ldots, y^j_d)$. We show that $\epsilon(\nu(\pi)(t_j)) \not\leq \epsilon(t)$ by decreasing induction on j. The case $j = k + 1$ is by assumption. Suppose $\epsilon(\nu(\pi)(t_{j+1})) \not\leq \epsilon(t)$. There must be an index $m \in \{1, 2, \cdots, d\}$ such that $x^{j+1}_m > y^{k+1}_m$. It is easy to see that the identity $\epsilon([c, \boldsymbol{a}] \vee_c [c, \boldsymbol{b}]) = \epsilon([c, \boldsymbol{a}]) \vee \epsilon([c, \boldsymbol{b}])$ holds, where the right-hand $\vee$ is componentwise max in $\mathbb{R}^d$ (a property that is not usually implied by the mere fact that ϵ is an isomorphism). From that and $\nu(\pi)(t_{j+1}) = \nu(\pi)(t_j) \vee_{c_j} \pi(t_{j+1})$, we infer that $x^{j+1}_m = \max(x^j_m, y^{j+1}_m)$, hence $y^{j+1}_m \leq x^{j+1}_m$. But π restricts to a dipath from t_j to t, so $\epsilon(\pi(t_j)) \leq \epsilon(t)$, and therefore $y^{j+1}_m \leq y^{k+1}_m < x^{j+1}_m$. From $y^{j+1}_m < x^{j+1}_m$ and $x^{j+1}_m = \max(x^j_m, y^{j+1}_m)$, we obtain $x^{j+1}_m = x^j_m$, whence $x^j_m > y^{k+1}_m$. In particular, $\epsilon(\nu(\pi)(t_j)) \not\leq \epsilon(t)$.

Taking $j = 0$, this implies that $\epsilon(s) \not\leq \epsilon(t)$. This is impossible, since π is a dipath from s to t.

We have constructed a map ν such that π and $\nu(\pi)$ have same carrier sequence. We can now conclude by the following lemma:

Lemma 3. *Let $F, G : P(X; s, t) \longrightarrow P(X; s', t')$ such that:*

- *for every pair of dipaths p, q that are equivalent modulo reparametrization, $F(p)$ and $F(q)$ are equivalent modulo reparametrization—so F induces $\tilde{F}$: $Tr(X; s, t) \longrightarrow Tr(X; s', t')$, and similarly for G.*
- *for every π, $F(\pi)$ and $G(\pi)$ have the same carrier sequence.*

Then $\tilde{F}$ and $\tilde{G}$ are homotopic.

Proof. Let $C(X; s', t')$ is the subspace of $P(X; s', t') \times P(X; s', t')$ that consists of pairs of dipaths that have the same carrier sequence. The key ingredient consists in constructing a continuous map $\Gamma : I \times C(X; s', t') \longrightarrow P(X; s', t')$ in such a way that $\Gamma(0, (p, q)) = p$ and $\Gamma(1, (p, q)) = q$. Let $c_0, c_1, \cdots, c_k$ be the common carrier sequence to p and q, let $t_0 \leq t_1 \leq \cdots \leq t_{k+1}$ be the times of change for p, and $s_0 \leq s_1 \leq \cdots \leq s_{k+1}$ be the times of change for q. Define $u_i(t) = ts_i + (1-t)t_i$ for $t \in [0, 1]$, $0 \leq i \leq k+1$. For every $u \in [u_i(t), u_{i+1}(t)]$, define v as $\frac{u - u_i(t)}{u_{i+1}(t) - u_i(t)}$. (This is defined provided $u_i(t) \neq u_{i+1}(t)$; if this is not the case, let $v = 0$.) Then $p(v(t_{i+1} - t_i) + t_i)$ is of the form $[c_i, (a_1^u, \ldots, a_m^u)]$ and $q(v(s_{i+1} - s_i) + s_i)$ is of the form $[c_i, (b_1^u, \ldots, b_m^u)]$. We then define $\Gamma(t, p, q)(u) = [c_i, (1-t)a_j^u + tb_j^u]$.

We have to define a homotopy $H : I \times Tr(X; s, t) \longrightarrow Tr(X; s', t')$. It will be defined as the composition of:

- $\mathrm{id} \times \kappa : I \times Tr(X; s, t) \longrightarrow I \times P(X; s, t)$, where κ is a continuous map from $Tr(X; s, t)$ to $P(X; s, t)$, defined in such a way that $\langle \kappa(\langle \pi \rangle) \rangle = \langle \pi \rangle$ for every trace $\langle \pi \rangle$, therefore defining a canonical dipath representing a given trace. The existence of such a map is shown by Raussen in [20], as the composition $norm \circ \overrightarrow{s}$ of two more elementary maps.
- $\mathrm{id} \times (F, G) : I \times P(X; s, t) \longrightarrow I \times C(X; s', t')$, where (F, G) maps π to $(F(\pi), G(\pi))$.
- $\Gamma : I \times C(X; s', t') \longrightarrow P(X; s', t')$, as defined above.
- and $\langle \_ \rangle : P(X; s', t') \longrightarrow Tr(X; s', t')$, which maps each dipath to its trace.

We compute: $H(0, \langle \pi \rangle) = \langle \Gamma(0, (F(\kappa(\langle \pi \rangle)), G(\kappa(\langle \pi \rangle)))) \rangle = \langle F(\kappa(\langle \pi \rangle)) \rangle = \tilde{F}(\langle \pi \rangle)$. Similarly, $H(1, \_) = \tilde{G}$ and therefore H is an homotopy from $\tilde{F}$ to $\tilde{G}$. □

It only remains to prove that the construction is natural. The following diagram:

$$\begin{array}{ccc}
Tr(X; s, t) & \xrightarrow{\quad ({}^*\mu_t)^{-1} \circ {}^*\rho_t \circ (\lambda_s^*)^{-1} \circ \eta_s^* \quad} & Tr(X; \mathcal{C}(s), \mathcal{C}(t)) \\
{\scriptstyle \langle \alpha \star _ \star \beta \rangle} \downarrow & & \downarrow {\scriptstyle \langle \hat{\gamma} \star _ \star \hat{\delta} \rangle} \\
Tr(X; s', t') & \xrightarrow[\quad ({}^*\mu_{t'})^{-1} \circ {}^*\rho_{t'} \circ (\lambda_{s'}^*)^{-1} \circ \rho_{s'}^* \quad]{} & Tr(X; \mathcal{C}(s'), \mathcal{C}(t'))
\end{array}$$

is commutative modulo homotopy because of the previous lemma and so the same diagram is commutative in homology, which proves the naturality. □

Remark. Lemma 2 is false in general in a non-looping precubical set. Our result states that if there is a dipath from s to t, s_- has the same carrier than s and there is a dipath from s_- to s then the trace spaces $Tr(X; s, t)$ and $Tr(X; s', t)$ are homotopically equivalent—in particular, they have the same number of connected components. But let us consider the following non-looping precubical set:

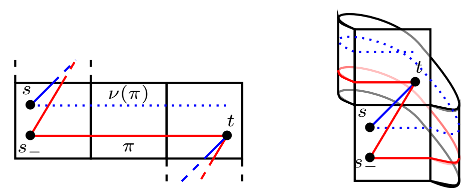

It has three squares (look at the view on the left), and the bottom face of the rightmost square is glued to the top face of the leftmost one. The glueing is displayed on the right. Consider now s, s_- and t as in the figure. $Tr(X; s, t)$ has one connected component (one of its element is drawn in plain blue) while $Tr(X; s', t)$ has two (an element of each is drawn in plain red). Hence Lemma 2 would fail if we allowed K to be a general non-looping precubical set, not just a cubical complex.

The argument we use to prove Lemma 2 works perfectly well in general non-looping precubical sets, except for one thing: it may be that $\nu(\pi)(1)$ does not coincide with t, and is strictly above. See the dotted blue line in the figure above to contemplate what $\nu(\pi)$ looks like in this example.

One may think of proving Theorem 1 by dispensing with Lemma 2, and finding another route. If this is true, this would be arduous. Notably, the twisted three-square counterexample above also disproves the fact that the pair $(\mathcal{C}, \sigma)$ we are constructing would be an open map, for K a general non-looping precubical set. We conjecture that $\overrightarrow{H}_n(\overrightarrow{Geom}(K))$ is not bisimilar to $\overrightarrow{h}_n(K)$ for general non-looping precubical sets K, hence that the assumption that K is a cubical complex is needed.

E Bonus

If you've read until now, you deserve knowing about a natural question we have not addressed in the paper. We plan to address this in more depth in future papers. The question is: can we decide whether two finite natural systems of abelian groups are bisimilar?

This seems to be a hard question. However, there is a variant of the question that has an easy answer, and which we describe next.

Everything we have done mentioned abelian groups. Abelian groups are $\mathbb{Z}$-modules, and one can think of generalizing by considering $\mathcal{R}$-modules instead, where $\mathcal{R}$ is a ring with unit. This is a classical trick in undirected homology, and is called homology *with coefficients in* $\mathcal{R}$. When $\mathcal{R}$ is a field, then $\mathcal{R}$-modules are vector spaces over $\mathcal{R}$. The computation of homology with coefficients in a field $\mathcal{R}$ is notably simpler than in abelian groups, because one does not have to care for torsion.

This is the view taken by the proponents of persistent homology, too, who always compute with coefficients in a field. Recall that our notion of natural homology has a much more complex structure, since its indexing category is not just a linear order.

Everything we did with abelian groups in the paper goes through by considering $\mathcal{R}$-modules instead. (The only thing that changes are the various "$\mathbb{Z}$" in diagrams, which should be replaced by $\mathcal{R}$.)

We can represent certain finite natural systems as follows. Call a *rational* natural system F a finite natural system of real vector spaces, presented as: a finite category X, a finite-dimensional real vector space for each object x of X, which we shall simply equate with some power of $\mathbb{R}$; and, for each morphism $f\colon x \to y$ in X, a rational matrix A_f. Bisimulations consist of triples (x, η, y), where η is a linear map between real vector spaces: namely, one representable as a matrix with real coefficients, not rational coefficients. Call them $\mathbb{R}$-bisimulations to make that clear.

Theorem 2. *Given two rational natural systems, it is decidable whether they are* $\mathbb{R}$*-bisimilar.*

Proof. A bisimulation between F, G can be represented as a finite set R of triples (x, P_{xy}, y) satisfying certain conditions, where P_{xy} are real matrices. We shall guess the pairs (x, y) and solve for the matrix P_{xy} and its inverse Q_{xy}. Call the corresponding set of pairs (x, y) the domain of R. We guess the domain of R and check that for every object x there is a y such that (x, y) is in the domain, and conversely. For each pair (x, y) in the domain, create two matrices of variables P_{xy}, Q_{xy}. For every morphism $i : x \longrightarrow x'$ on the left and every pair (x, y) in the domain, we guess a morphism $j : y \longrightarrow y'$ on the right such that (x', y') in the domain, and produce the equation $P_{x'y'}A_i = A_j Q_{xy}$ (and conversely, for each j, guessing an i). Finally, add the equations $P_{xy}Q_{xy} = 1$ and $Q_{xy}P_{xy} = 1$ for each pair (x, y) in the domain. Collect all equations, quantify existentially, and solve the resulting formula: indeed the first-order theory of reals is decidable. $\square$