

Polynomial-Time Homology for Simplicial Eilenberg–MacLane Spaces

Journal Article

Author(s): Krčál, Marek; Matoušek, Jiří; Sergeraert, Francis

Publication date: 2013-12

Permanent link: https://doi.org/10.3929/ethz-b-000412816

Rights / license: In Copyright - Non-Commercial Use Permitted

Originally published in: Foundations of Computational Mathematics 13(6), <u>https://doi.org/10.1007/s10208-013-9159-7</u>

The Journal of the Society for the Foundations of Computational Mathematics

FOUNDATIONS OF

COMPUTATIONAL MATHEMATICS

Polynomial-Time Homology for Simplicial Eilenberg–MacLane Spaces

Marek Krčál · Jiří Matoušek · Francis Sergeraert

Received: 30 January 2012 / Accepted: 24 April 2013 / Published online: 30 May 2013 © SFoCM 2013

Abstract In an earlier paper of Čadek, Vokřínek, Wagner, and the present authors, we investigated an algorithmic problem in computational algebraic topology, namely, the computation of all possible homotopy classes of maps between two topological spaces, under suitable restriction on the spaces.

We aim at showing that, if the dimensions of the considered spaces are bounded by a constant, then the computations can be done in polynomial time. In this paper we make a significant technical step towards this goal: we show that the Eilenberg– MacLane space $K(\mathbb{Z}, 1)$, represented as a simplicial group, can be equipped with *polynomial-time homology* (this is a polynomial-time version of *effective homology* considered in previous works of the third author and co-workers).

To this end, we construct a suitable *discrete vector field*, in the sense of Forman's *discrete Morse theory*, on $K(\mathbb{Z}, 1)$. The construction is purely combinatorial and it can be understood as a certain procedure for reducing finite sequences of integers, without any reference to topology.

Communicated by Peter Buergisser.

M. Krčál · J. Matoušek (🖂)

Department of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic e-mail: matousek@kam.mff.cuni.cz

J. Matoušek Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland

F. Sergeraert Institut Fourier, BP 74, 38402 St Martin, d'Hères Cedex, France



The Eilenberg–MacLane spaces are the basic building blocks in a *Postnikov sys*tem, which is a "layered" representation of a topological space suitable for homotopytheoretic computations. Employing the result of this paper together with other results on polynomial-time homology, in another paper we obtain, for every fixed k, a polynomial-time algorithm for computing the *k*th homotopy group $\pi_k(X)$ of a given simply connected space X, as well as the first k stages of a Postnikov system for X, and also a polynomial-time version of the algorithm of Čadek et al. mentioned above.

Keywords Computational homotopy theory · Eilenberg–MacLane space · Postnikov system · Effective homology

Mathematics Subject Classification 68U05 · 68W99 · 55S45 · 55S37

1 Introduction

Recently our co-authors and we [2] have developed an algorithm for a problem in computational algebraic topology (more precisely, in computational homotopy theory), namely, computing all homotopy classes of maps between two topological spaces X and Y (given as finite simplicial complexes, say), under certain natural conditions on X and Y.

Our original motivation was understanding the computational complexity of the \mathbb{Z}_2 -index of a given \mathbb{Z}_2 -space, which is a quantity appearing in various applications of topology in combinatorics and geometry (e.g., topological lower bounds for the chromatic number of a graph, or an algorithm for testing the embeddability of a given simplicial complex into \mathbb{R}^d). We hope to reach results in this direction in the future, and we also expect that the developed methods will be applicable for other natural problems (such as extendability of maps; as a concrete application, it was already possible to answer a question of Franek et al. [11] on testing nullhomotopy of maps into a sphere). For more information on this project we refer to [2, 3, 4].

Towards Polynomial-Time Homology The implementation of some of the operations in the algorithm of [2] relies on the methods of *effective homology*, initiated by the third author in [21] and further developed by him and his co-workers (see, e.g., [18–20]). These provide algorithmic solutions of many problems in algebraic topology, but so far no analysis of their running time was available, and for some parts the running time can actually be exponential.

One of our aims is to obtain *polynomial-time* algorithms for these tasks where possible, or alternatively, show computational hardness.

Let us stress that by "polynomial-time" we mean, throughout this paper, *polynomial-time for every fixed dimension*. Thus, assuming that the input to an algorithm is a space represented as a finite simplicial complex X, we want that the running time is polynomial in the number of simplices of X, but the polynomial may depend on the dimension k of X (and the dependence on k may be exponential or even worse). Of course, one could be even more ambitious and ask for a polynomial dependence on k as well; however, we do not expect such algorithms to exist, in view of computational hardness results [1, 4].

937

To integrate this effort with existing algorithms, we start with the framework of effective homology mentioned above, and we introduce an analogous definition of polynomial-time homology; see Sect. 2. In another paper [3], we show that various known constructions and operations on objects with effective homology have polynomial-time versions. With a repertoire of such operations, we also obtain a polynomial-time version of the algorithm of [2], as well as other algorithms, such as computing the higher homotopy group $\pi_k(X)$ in polynomial time for every fixed k, or computing the first k stages of a Postnikov system for X.

This Paper Here we make a significant step in this development. First we set up the framework of polynomial-time homology (modeled after effective homology mentioned above) and some tools of general applicability. Then, in the second part of the paper, we present our main technical result. The problem which we solve can be formulated purely combinatorially, although in this language it perhaps does not sound extremely natural: it is a question about reducing finite sequences of integers by certain simple operations. We will state it below, and no topological notion at all is required for understanding this problem and our solution.

However, to explain its role in computational topology, we first need to sketch some background information. A standard reference for this material is May [15]; a concise overview is given in [2], and more leisurely explanations can be found in [22] or [20].

A common technique in mathematics and in computer science is to decompose a general, presumably complicated object into simpler building blocks. For the purposes of understanding continuous mappings going *into* a given topological space Y, a suitable decomposition is a Postnikov system for Y; indeed, this is a crucial ingredient of the algorithm in [2].

We do not need to define the rather complicated notion of Postnikov system here; it suffices to say that its "building blocks" belong to a particular class of topological spaces, called *Eilenberg–MacLane spaces* and denoted by K(G, k), where G is an Abelian group and $k \ge 1$ is an integer. In the Eilenberg–MacLane spaces appearing in a Postnikov system for Y, the role of the group G is played by the homotopy groups $\pi_i(Y), i \geq 2.$

In topology, K(G, k) is defined as a topological space T whose homotopy groups satisfy $\pi_k(T) \cong G$ and $\pi_i(T) = 0$ for all $i \neq k$. It is determined uniquely up to homotopy equivalence (in the class of all CW complexes).

Generally speaking, the spaces K(G, k) are infinite-dimensional and they do not look like very simple objects (with the exception of $K(\mathbb{Z}, 1)$, which is homotopy equivalent to the circle S^1). However, they are in some sense the simplest possible spaces concerning maps going into them. These spaces are of basic importance in algebraic topology, and a lot of work has been devoted to studying their properties, and in particular, computing their homology and cohomology (Serre [23] and H. Cartan [5] are two of the most famous classical works; see, e.g., Clément [6] for an overview and some computational aspects). We also refer to Romero and Rubio [16] for an algorithmic study of K(G, 1) for noncommutative groups G.

For the intended algorithmic use, we need a particular representation of K(G, k); namely, we need it represented as a particular kind of a simplicial set (simplicial sets will be briefly introduced in Sect. 2 below), a so-called *Kan simplicial set*. We use the standard Eilenberg–MacLane simplicial model for K(G, k); see [8, Chap. III], [15, Chap. V].

For the algorithms, we need to equip the simplicial Eilenberg–MacLane spaces with polynomial-time homology. The K(G, k) we may encounter can have any finitely generated Abelian group as G, and any positive integer as k.

However, in this paper we will deal only with $K(\mathbb{Z}, 1)$, which serves as a base case, while the other K(G, k) can be obtained from it using several operations. First, for direct products of groups, we have $K(G \times H, k) \cong K(G, k) \times K(H, k)$, and so, with a general product operation available, we may assume that *G* is cyclic. Second, a general construction, known as the *classifying space* (actually, in the simplicial setting, we deal with the so-called \overline{W} -construction), allows one to pass from K(G, k) to K(G, k+1), so indeed k = 1 is the important base case. Finally, polynomial-time homology for $K(\mathbb{Z}/m\mathbb{Z}, 1)$ can be obtained from that for $K(\mathbb{Z}, 1)$ using another operation, namely, computing the *base space of a fibration*. These reductions are discussed in [20], and polynomial-time versions are discussed in [3]; here we just wanted to provide a quick explanation of why the $K(\mathbb{Z}, 1)$ case deserves special attention.¹

The Combinatorial Problem About Integer Sequences The *k*-dimensional simplices of the standard simplicial model of $K(\mathbb{Z}, 1)$, k = 0, 1, ... can be represented by *k*-term sequences of integers. With the traditional "bar notation", such a sequence is written as

$$\sigma = [a_1 \mid a_2 \mid \dots \mid a_k], \quad a_1, a_2, \dots, a_k \in \mathbb{Z}.$$
 (1)

In the rest of this introduction, a "*k*-dimensional simplex" will thus be synonymous with a "*k*-term sequence of integers".

For our problem we consider only *nondegenerate simplices*, represented by sequences with *no zero terms*. Thus, from now on, we always assume that all the a_i are nonzero.

For each k, there are k + 1 face operators $\partial_0, \partial_1, \ldots, \partial_k$, which map k-term sequences to k - 1 term sequences: ∂_0 deletes the first component, ∂_k deletes the last component, and for $i = 1, 2, \ldots, k - 1$, ∂_i reduces the number of components by one by adding together the *i*th and (i + 1)st component. More formally, with σ as above,

$$\partial_0 \sigma = [a_2 | \dots | a_k], \qquad \partial_k \sigma = [a_1 | \dots | a_{k-1}],$$
$$\partial_i \sigma = [a_1 | \dots | a_{i-1} | a_i + a_{i+1} | a_{i+2} | \dots | a_k], \qquad 1 \le i \le k-1$$

The goal is to divide the set of all possible finite sequences σ of nonzero integers into three classes S, T, and C (the *source simplices*, *target simplices*, and *critical*

¹Curiously, $K(\mathbb{Z}, 1)$ as a topological space almost cannot be simpler—as we mentioned, it is homotopy equivalent to the circle S^1 , and other Eilenberg–MacLane spaces are much more complicated. But we need to work with the Kan simplicial model of $K(\mathbb{Z}, 1)$ as introduced above, which has infinitely many simplices in every dimension $k \ge 1$. As we will see, for effective (or polynomial-time) homology, it is not sufficient to know, for example, that $H_2(K(\mathbb{Z}, 1)) = 0$, but we need to be able to actually compute "witnesses" for it; that is, given a 2-cycle z_2 on $K(\mathbb{Z}, 1)$, compute a 3-chain for which z_2 is its boundary. This problem would be trivial for the standard simplicial representation of S^1 with one vertex and one edge, but it is not trivial for the considered Kan model of $K(\mathbb{Z}, 1)$.

simplices), and construct a bijection $V : S \to T$ (which will be called a *discrete vector field*), such that for every $\sigma \in S$, we have $\sigma = \partial_i V(\sigma)$ for exactly one *i*. We also require certain additional properties, which we explain next.

With S, T, C, and V as above, let us consider a sequence (simplex) $\tilde{\sigma} \in S$ of some dimension k, and let us say that a simplex τ (of dimension k or k + 1) is *reachable* from $\tilde{\sigma}$ if it can be reached from $\tilde{\sigma}$ by finitely many moves, where the allowed moves are

- passing from a current simplex $\sigma \in S$ to the simplex $\tau = V(\sigma) \in T$, and
- passing from a current simplex $\tau \in T$ to a simplex $\sigma = \partial_i \tau \in S \cup C$ such that $\tau \neq V(\sigma)$, where $i \in \{0, 1, ..., k\}$.

With these definitions, it is required that

- (i) for every k, C contains only finitely many k-dimensional simplices; and
- (ii) starting with any $\tilde{\sigma}$, we can never make an infinite sequence of allowed moves; that is, we can reach only finitely many simplices, and we also cannot get into a cycle.

Moreover, we measure the *size* of a simplex $\sigma = [a_1|\cdots|a_k]$ as the total number of bits needed to write down a_1, \ldots, a_k ; more formally, we set $size(\sigma) := \sum_{i=1}^k size(a_i)$ and $size(a) := 1 + \lfloor \log_2(|a|+1) \rfloor$. Then we also require that

(iii) For every k-dimensional simplex $\tilde{\sigma}$, the sum of size(σ) over all σ reachable from $\tilde{\sigma}$ is bounded by a polynomial (depending on k) in size($\tilde{\sigma}$).

To illustrate these definitions, let us present a classical vector field V_{EML} due to Eilenberg and Mac Lane, which satisfies (i) and (ii) (and yields effective homology for $K(\mathbb{Z}, 1)$) but not (iii).

There are only two critical simplices, the 0-dimensional [] (the empty sequence) and the 1-dimensional [1].² The set S of source simplices consists of the sequences with $a_1 \neq 1$, while T contains the sequences with $a_1 = 1$ (the two critical simplices are exceptions to this rule).

For $\sigma = [a_1 | \cdots | a_k] \in S$, $a_1 \neq 1$, the vector field V_{EML} is defined by

$$V_{\text{EML}}(\sigma) := \begin{cases} [1 \mid a_1 - 1 \mid a_2 \mid \dots \mid a_k] & \text{for } a_1 > 1, \\ [1 \mid a_1 \mid a_2 \mid \dots \mid a_k] & \text{for } a_1 < 0. \end{cases}$$

It can be checked that, for any starting $\tilde{\sigma}$, the sequence of moves is determined uniquely (there is no branching), and that it always terminates after finitely many steps.

It is easy to see that, for a positive integer a, the sequence of moves starting from [a] is $[a] \rightarrow [1|a-1] \rightarrow [a-1] \rightarrow [1|a-2] \rightarrow [a-2] \rightarrow \cdots$; there are about a moves, and this is *exponential* in the number of bits of a. Thus, condition (iii) above indeed fails.

²This actually corresponds to the topological fact that the considered $K(\mathbb{Z}, 1)$, as a topological space, is homotopy equivalent to S^1 ; [] represents a vertex, and [1] an edge glued to that vertex by both ends, forming an S^1 .

We will provide a solution satisfying (i)–(iii) in Sect. 4. Before that, we introduce simplicial sets, polynomial-time homology, and discrete vector fields in general.

2 Simplicial Sets with Polynomial-Time Homology

Simplicial Sets A simplicial complex is a way of specifying a topological space in purely combinatorial terms, and also a way of presenting a topological space as an input to an algorithm; we assume that the reader is basically familiar with this concept.

A simplicial set can be regarded as a generalization of a simplicial complex; it is more complicated, but more powerful and flexible. The algorithms we consider use simplicial sets as the main data type for representing topological spaces and their maps. A friendly introduction to simplicial sets is [12], and another introductory treatment can be found in [22]; older compact sources are, e.g., [7, 15], and [13] is a more modern and comprehensive treatment.

Similar to a simplicial complex, a simplicial set is a space built of vertices, edges, triangles, and higher-dimensional simplices, but simplices are allowed to be glued to each other and to themselves in more general ways. For example, one may have several 1-dimensional simplices connecting the same pair of vertices, a 1-simplex forming a loop, two edges of a 2-simplex identified to create a cone, or the boundary of a 2-simplex all contracted to a single vertex, forming an S^2 .



Another new feature of a simplicial set, in comparison with a simplicial complex, is the presence of *degenerate simplices*. For example, the edges of the triangle with a contracted boundary (in the last example above) do not disappear, but each of them becomes a degenerate 1-simplex.

A simplicial set X is represented as a sequence $(X_0, X_1, X_2, ...)$ of mutually disjoint sets, where the elements of X_k are called the *k*-simplices of X (we note that, unlike for simplicial complexes, a simplex in a simplicial set need not be determined by the set of its vertices; indeed, there can be many simplices with the same vertex set). For every $k \ge 1$, there are k + 1 mappings $\partial_0, \ldots, \partial_k : X_k \to X_{k-1}$ called *face operators*; the intuitive meaning is that for a simplex $\sigma \in X_k$, $\partial_i \sigma$ is the face of σ opposite to the *i*th vertex. Moreover, there are k + 1 mappings $s_0, \ldots, s_k : X_k \to X_{k+1}$ (opposite direction) called the *degeneracy operators*; the approximate meaning of $s_i \sigma$ is the degenerate simplex which is geometrically identical to σ , but with the *i*th vertex duplicated. A simplex is called *degenerate* if it lies in the image of some s_i ; otherwise, it is *nondegenerate*. We write X^{ndg} for the set of all nondegenerate simplices of X.

There are natural axioms that the ∂_i and the s_i have to satisfy, but we will not list them here, since we will not really use them. Moreover, the usual definition of

simplicial sets uses the language of category theory and is very elegant and concise; see, e.g., [13, Sect. I.1].

Every simplicial set X specifies a topological space |X|, the *geometric realization* of X. It is obtained by assigning a geometric k-dimensional simplex to each nondegenerate k-simplex of X, and then gluing these simplices together according to the face operators; we refer to the literature for the precise definition.

There is a canonical way of converting a simplicial complex to a simplicial set; basically, one just needs to add appropriate degenerate simplices.

We have already given a relatively sophisticated example of a simplicial set, namely, $K(\mathbb{Z}, 1)$, or more precisely, the standard Eilenberg–MacLane representation of $K(\mathbb{Z}, 1)$ as a Kan simplicial set³ as defined in the introduction (except that we have not yet specified the degeneracy operators, which are very simple: s_i inserts 0 after the *i*th component of a sequence).

Representing Infinite Simplicial Sets In many areas where computer scientists seek efficient algorithms, both the input objects and the intermediate results in the algorithms are finite, and they can be explicitly represented in the computer memory; this is the case, e.g., for algorithms dealing with graphs or with matrices.

In contrast, in the algorithms for homotopy-theoretic questions considered here and in related works, we need to deal with infinite objects. For example, even if the input is a finite simplicial complex, its Postnikov system (mentioned in the introduction) is made of Eilenberg–MacLane spaces, such as $K(\mathbb{Z}, 1)$, represented as Kan simplicial sets, and these are necessarily infinite. More concretely, as we have seen, $K(\mathbb{Z}, 1)$ has infinitely many simplices in each dimension $k \ge 1$, and thus we cannot explicitly store even the part up to some fixed dimension.

For algorithmic purposes, we thus represent a simplicial set X by a collection of several algorithms, which allow us to access certain information about X, without having all of it explicitly stored in memory. (In computer science, this is also called a *black box* or *oracle* representation of X, and in the terminology of object-oriented programming, we can think of X as an instance of a *class* "simplicial set".) A similar representation is used for other kinds of infinite topological or algebraic objects as well.

Locally Effective Simplicial Sets For some computations, it may be sufficient to represent X by a black box providing only "local" information about X, and in that case, in accordance with the terminology in earlier papers, e.g., [17, 19, 20], we speak of a *locally effective representation*.

Concretely, let *X* be a simplicial set, and suppose that some computer representation ("encoding") for the simplices of *X* has been fixed. For example, in the case of $K(\mathbb{Z}, 1)$, we can fix the representation of the simplices of $K(\mathbb{Z}, 1)$ by integer sequences, and represent the integers in the sequences by the standard binary encoding.

E°⊂ D

³We will not define a Kan simplicial set, but we just mention a key property, which is the reason why these simplicial sets are essential to the considered algorithms. Namely, if X is a simplicial set and Y is a Kan simplicial set, then every *continuous* map $|X| \rightarrow |Y|$ is homotopic to a *simplicial* map $X \rightarrow Y$. Thus, continuous maps into Y have a combinatorial representation, describing them up to homotopy.

We say that *X* is a *locally effective simplicial set* if algorithms are available that, given (an encoding of) a *k*-simplex σ of *X* and $i \in \{0, 1, ..., k\}$, computes the simplex $\partial_i \sigma$, and similarly for the degeneracy operators s_i . Briefly speaking, the face and degeneracy operators should be computable maps.

Computing Global Information Suppose that we want to compute some "global" information about a given simplicial set X, for example, the *k*th homology group $H_k(X)$. Then a locally effective representation of X is typically insufficient, and we need to augment it in some way.

Of course, in the particular example with the homology groups, we could insist that *X* be augmented with a black box that, given *k*, returns some representation of $H_k(X)$. The problem is that *X* may not be given to us directly; rather, we may need to construct it from other simplicial sets by a sequence of various operations. For example, in the introduction we mentioned that the Eilenberg–MacLane spaces K(G, k) can be constructed starting with $K(\mathbb{Z}, 1)$ and applying operations of several kinds, such as product or classifying space.⁴ Then, for example, a black box for computing the homology groups of *X* is not in itself sufficient to compute the homology groups of the classifying space of *X*.

The third author and his co-authors have developed a more sophisticated way of augmenting a locally effective simplicial set X with homological information, which is captured in the notion of a *simplicial set with effective homology*. These simplicial sets do possess a black box for computing homology groups, but they are also equipped with additional information, which makes them stable under a large repertoire of operations: if we apply some of the "classical" operations, such as product, classifying space, loop space, etc. to simplicial sets with effective homology, the result is again a simplicial set with effective homology (and in particular, it has a black box for computing homology groups).⁵

It may be useful to keep in mind that, since a simplicial set is represented by a black box, operations on such simplicial sets are performed by *composition of algorithms*; i.e., the black box for the new simplicial set operates by calling the black boxes of the old sets and processing the values returned by them.⁶

For defining a simplicial set with effective homology, and their polynomial-time counterpart, we need to recall some notions concerning chain complexes.

Chain Complexes For our purposes, a *chain complex* C_* is a sequence $(C_k)_{k=-\infty}^{\infty}$ of *free* \mathbb{Z} -modules (i.e., free Abelian groups), together with a sequence $(d_k : C_k \rightarrow$

⁴As another, perhaps more sophisticated example, we can mention the computation of the homotopy group $\pi_k(X)$ for a 1-connected simplicial set X: for this, given X, one first produces another simplicial set X' from X, by a sequence of operations that "kill" the first k - 1 homotopy groups, and then $\pi_k(X)$ is computed as $H_k(X')$ using the *Hurewicz isomorphism*.

 $^{^{5}}$ One can also consider other kinds of objects with effective homology, such as chain complexes, but for concreteness, we will stick to simplicial sets.

⁶This feature makes it very natural to implement algorithms from this area using *functional programming languages*, as was done for the package *Kenzo*; see, e.g., [14].

 $C_{k-1})_{k=-\infty}^{\infty}$ of group homomorphisms.⁷ The C_k are the *chain groups*, their elements are called *k-chains*, and the d_k the *differentials*. The differentials have to satisfy $d_{k-1}d_k = 0$ for every *k* (here $d_{k-1}d_k$ denotes the composition of maps). We also recall that the *k*th homology group $H_k(C_*)$ of the chain complex C_* is defined as the factor-group ker $d_k/\operatorname{im} d_{k+1}$.

For every simplicial set X, there is a canonically associated chain complex, which is used to define the homology groups $H_k(X)$. Actually, there are two natural possibilities, depending on whether degenerate simplices are taken into account. We use the *normalized* chain complex, which is based solely on the nondegenerate simplices. We reserve the simple notation $C_*(X)$ for it.

Thus, $C_k(X)$ denotes the free Abelian group over X_k^{ndg} , the set of all *k*-dimensional nondegenerate simplices (in particular, $C_k(X) = 0$ for k < 0). This means that a *k*-chain is a formal sum

$$c = \sum_{\sigma \in X_k^{\mathrm{ndg}}} \alpha_{\sigma} \cdot \sigma,$$

where the α_{σ} are integers, only finitely many of them nonzero. The differentials are defined in a standard way using the face operators: for *k*-chains of the form $1 \cdot \sigma$, which constitute a basis of $C_k(X)$, we set $d_k(1 \cdot \sigma) := \sum_{i=0}^k (-1)^i \cdot \partial_i \sigma$ (some of the $\partial_i \sigma$ may be degenerate simplices; then they are ignored in the sum), and this extends to a homomorphism in a unique way ("linearly").

We note that if X is a locally effective simplicial set, then the k-chains of $C_*(X)$ are finite objects; a k-chain c can be represented by a list of the k-simplices σ on which c is nonzero, and of the corresponding coefficients α_{σ} . Then the differentials are computable maps.

However, if X_k^{ndg} is infinite, then $C_k(X)$ has infinite rank, and we cannot use it directly for computing homology groups. The solution adopted in effective homology is to have, together with a locally effective simplicial set X, a *reduction* from $C_*(X)$ to an "effective" chain complex EC_* , for which each chain group EC_k has a finite rank.

Reductions Let C_* , \tilde{C}_* be two chain complexes. To define a reduction from C_* to \tilde{C}_* , we first recall two other standard notions from homological algebra: A *chain map* $f: C_* \to \tilde{C}_*$ is a sequence $(f_k)_{k=-\infty}^{\infty}$ of homomorphisms $f_k: C_k \to \tilde{C}_k$ compatible with the differentials, i.e., $f_{k-1}d_k = \tilde{d}_k f_k$. If $f, g: C_* \to \tilde{C}_*$ are two chain maps, then a *chain homotopy* of f and g is a sequence $(h_k)_{k=-\infty}^{\infty}$ of homomorphisms $h_k: C_k \to \tilde{C}_{k+1}$ such that $f - g = \tilde{d}_{k+1}h_k + h_{k-1}d_k$.

Now a reduction ρ from C_* to \tilde{C}_* consists of three maps f, g, h, such that



⁷These chain complexes are over \mathbb{Z} ; more generally, one considers chain complexes over a commutative ring *R*, where the C_k are *R*-modules. These are needed, among others, for homology with coefficients in *R*. But for our purposes, homology with integer coefficients suffices; if needed, homology groups with other coefficients can be computed using universal coefficient theorems. Alternatively, all of the theory can be built with coefficients from a fixed ring *R*, provided that *R* is equipped with sufficiently strong algorithmic primitives.

- $f: C_* \to \tilde{C}_*$ and $g: \tilde{C}_* \to C_*$ are chain maps;
- the composition fg: C̃_{*} → C̃_{*} is equal to the identity id_{C̃*}, while the composition gf: C_{*} → C_{*} is chain-homotopic to id_{C*}, with h: C_{*} → C_{*} providing the chain homotopy; and
- fh = 0, hg = 0, and hh = 0.

The notion of reduction goes back to Eilenberg and Mac Lane [8, Sect. 12], who called it a *contraction*.⁸ It is routine to check that if there is a reduction from C_* to \tilde{C}_* , then C_* and \tilde{C}_* have isomorphic homology groups in each dimension. Reductions can also be composed, as follows: if (f, g, h) is a reduction from C_* to \tilde{C}_* and (f', g', h') is a reduction from \tilde{C}_* to \tilde{C}_* , then (f'f, gg', h + gh'f) is a reduction from C_* to \tilde{C}_* .

Effective Homology We are getting close to stating the definition of a simplicial set with effective homology. The last step is to define what we mean by an effective chain complex EC_* . We assume that, first, EC_* is *locally effective*, meaning that each chain group EC_k has some distinguished basis Bas_k , *k*-chains are represented as linear combinations of elements of Bas_k (and thus they can be added and subtracted algorithmically), and there is an algorithm for evaluating the differentials d_k . Second, EC_* is *effective*, which means, in addition to the above, that there is an algorithm that, given *k*, outputs the list of elements of the distinguished basis Bas_k ; in particular, this implies that each EC_k has a finite rank r_k . We note that by combining the construction of Bas_k with the ability to evaluate the differential d_k , we can compute the matrix of d_k with respect to the distinguished bases Bas_k and Bas_{k-1} .

We can now define a *simplicial set with effective homology* as a locally effective simplicial set X together with an effective chain complex EC_* and a reduction ρ from $C_*(X)$ to EC_* , where the three maps f, g, h from the definition of reduction are computable.⁹

In this paper we will not have the opportunity to demonstrate the usefulness of effective homology in algorithms; we refer to, e.g., [3, 20, 22] for examples of applications.

Polynomial-Time Homology The meaning of polynomial-time homology for the simplicial set $K(\mathbb{Z}, 1)$ considered in this paper is defined in a straightforward way: we want the face and degeneracy operators to be computable in polynomial time (which is obvious in this particular case), and $K(\mathbb{Z}, 1)$ should be equipped with effective homology as above in such a way that, for every k, the maps f_k , g_k , h_k are computable in polynomial time, with the polynomial possibly depending on k as usual.

⁸They did not require the condition hh = 0, but simple transformation converts a reduction without this condition into another one satisfying it.

⁹In [20] and in other papers, effective homology is defined in a more general way, using *strong equivalence* of chain complexes instead of just a reduction. A strong equivalence of C_* and \tilde{C}_* means that there is an auxiliary chain complex A_* and reductions of A_* to both C_* and \tilde{C}_* . However, here the simpler notion using a single reduction suffices, and this only makes the result formally stronger, since a reduction is a special case of a strong equivalence.

We stress that since we deal with a *single* effective chain complex EC_* , the ranks r_k depend only on k and thus, for k fixed, they are constants. The matrix of the differential d_k in EC_* , too, is a *constant-size* object.

However, our setting with $K(\mathbb{Z}, 1)$ is somewhat unusual in the analysis of algorithms: We are dealing with a single simplicial set, fixed once and for all, which does not depend on any input. This is an exceptional setting; most algorithms work with objects that do depend on the input. To draw an analogy from a different area, the setting of the present paper can be compared to seeking an algorithm for computing the *n*th digit of the number π , while the more usual case would be to consider algorithms for evaluating arithmetic expressions with arbitrary precision, where we start with integer numbers as inputs and apply addition, subtraction, multiplication, division, roots and functions like exp, ln or arcsin.

To have an example from the area considered here, in an algorithm for computing with a given topological space X, say specified as a finite simplicial complex, we may need polynomial-time homology for the Eilenberg–MacLane space $K(\mathbb{Z}^n, 1)$, where n is a parameter depending on X. Then we want that in the corresponding effective chain complex for $K(\mathbb{Z}^n, 1)$, the ranks r_2, r_3 , etc. each depend polynomially on n. (Of course, for this to be useful, we also need that n depends at most polynomially on the size of X.)

This example suggests that, in order to have a generally useful notion of polynomial-time homology, we need to define it formally for a whole family, typically infinite, of simplicial sets. Here we present this issue briefly, referring to [3] for a more detailed discussion.

Let \mathcal{I} be a set, typically countable, such that each element $I \in \mathcal{I}$ has some agreedupon computer representation (i.e. encoding by a finite string of bits). A *simplicial set parameterized by* \mathcal{I} is a mapping X that assigns a simplicial set X(I) to each $I \in \mathcal{I}$. We also assume that the simplices of each X(I) have some encoding by bit strings. Then we define a *locally polynomial-time simplicial set* as a simplicial set X parameterized by some \mathcal{I} such that the face and degeneracy operators on a k-simplex σ of X(I) can be evaluated in time polynomial in $size(I) + size(\sigma)$, where the polynomial may depend on k (and $size(\cdot)$ denotes the number of bits in the encoding).

Quite analogously, we define a chain complex $C_* = (C(I)_* : I \in \mathcal{I})$ parameterized by a set \mathcal{I} . We say that such a C_* is *locally polynomial-time* if each $C(I)_*$ is a locally effective chain complex (and in particular, it has a distinguished basis $Bas(I)_k$, and *k*-chains are represented w.r.t. this basis), and for each fixed *k*, the differential $(d_I)_k$ on $C(I)_k$ can be evaluated in time polynomial in size(*I*) plus the size of the input *k*-chain. We observe that addition and subtraction of *k*-chains are polynomial-time operations automatically.

We say that a simplicial set X parameterized by a set \mathcal{I} is equipped with *polynomial-time homology* if the following hold.

- X is locally polynomial-time.
- There is a locally polynomial-time chain complex EC_* , also parameterized by \mathcal{I} , such that, for each fixed k, the distinguished basis $Bas(I)_k$ of $EC(I)_k$ can be computed in time polynomial in size(I), and in particular, the rank $r(I)_k$ is bounded by such a polynomial.



• For every $I \in \mathcal{I}$, there is a reduction ρ_I from $C_*(X(I))$ to $EC(I)_*$, where the maps $(f_I)_k, (g_I)_k, (h_I)_k$ of ρ_I are all computable in time bounded by a polynomial in size(*I*) plus the size of the input *k*-chain; the polynomial may depend on *k*.

3 Polynomial-Time Homology from a Discrete Vector Field

Discrete Morse theory, developed by Forman [9] (also see [10]), belongs among fundamental tools in combinatorial topology. For us, the important point is that a suitable *discrete vector field* on a simplicial set¹⁰ X can be used to equip X with effective homology; this is an implication of one of Forman's results, as was observed by Romero and Sergeraert [17] (they also generalized Forman's construction by dropping a certain finiteness condition). Here we review the definitions, more or less repeating in a general setting the definitions given for $K(\mathbb{Z}, 1)$ in the introduction. Then we formulate a sufficient condition on the vector field so that the construction provides polynomial-time homology for X.

Discrete Vector Fields Let *X* be a simplicial set. For a simplex $\tau \in X$, it may happen that two face operators give the same simplex, i.e., $\partial_i \tau = \partial_j \tau$, $i \neq j$ (geometrically, this means that the two faces of the simplex τ are "glued together"). We say that σ is a *regular face* of τ if $\sigma = \partial_i \tau$ for *exactly one* index *i*.

A *discrete vector field* V on a simplicial set X is a set of ordered pairs (directed edges) of the form (σ, τ) , where $\sigma, \tau \in X^{ndg}$, σ is a regular face of τ , and for every two distinct pairs $(\sigma, \tau), (\sigma', \tau') \in V$, all of $\sigma, \tau, \sigma', \tau'$ are distinct.

Given a discrete vector field V, the nondegenerate simplices of X are classified into three subsets S, T, and C as follows:

- S are the *source simplices*; these are simplices σ such that $(\sigma, \tau) \in V$ for some τ .
- \mathcal{T} are the *target simplices*; these are simplices τ such that $(\sigma, \tau) \in V$ for some σ .
- *C* are the *critical simplices*; these are the remaining simplices, not occurring in any edge of *V*.

Often it is useful to regard V as a bijective mapping $V : S \to T$, as we did in the introduction. Thus, for $(\sigma, \tau) \in V$, we sometimes write $\tau = V(\sigma)$ and $\sigma = V^{-1}(\tau)$.

In a drawing of a simplicial set, the pairs (σ, τ) of a vector field can be indicated by arrows pointing from σ into τ , as in Fig. 1.

Admissible Vector Fields and the $V\partial$ -Graph The vector fields useful in discrete Morse theory, as well as in our context, have an extra property. For defining it, we first introduce an auxiliary directed graph, as drawn in Fig. 2, which we call the $V\partial$ -graph.

The vertex set of the $V\partial$ -graph is X^{ndg} . In the drawing, the empty circles correspond to source simplices, the full circles to target simplices, and the critical simplices are marked by double circles.

¹⁰In [17], vector fields are considered in somewhat greater generality, on *algebraic cell complexes*. Here it is sufficient to stay in the perhaps more intuitive setting of vector fields on simplicial sets.



Fig. 2 The V∂-graph corresponding to Fig. 1

The edges of the $V\partial$ -graph are of two kinds: first, those belonging to V (drawn bold and pointing upwards), and second, all edges of the form (τ, σ) , where τ is a *target* simplex, σ is a face of τ and a *source* or *critical* simplex, and $(\sigma, \tau) \notin V$ (these edges point downwards).¹¹ These edges correspond to the "allowed moves" defined in the introduction.

We call the vector field V *admissible* if the V ∂ -graph contains no directed cycle and no infinite directed path. The field in Fig. 1 is admissible, for example.

One of Forman's results says that an admissible vector field V can be used to "simplify" the underlying simplicial set X: by a sequence of suitable collapsing operations, which is defined based on V, one obtains a cell complex (no longer necessarily a simplicial set), which is homotopy equivalent to X but typically much smaller—its cells correspond only to the critical simplices.

We will not use this result directly (and thus we do not formulate it precisely). Rather, we build on a related result (obtained implicitly by Forman with an additional finiteness assumption, and explicitly and in general in [17]), asserting that an admissible vector field provides a reduction of the normalized chain complex $C_*(X)$ to a suitable chain complex C_*^{crit} . In this chain complex, each C_k^{crit} is the free Abelian group on the set of all *k*-dimensional critical simplices. The differentials in C_*^{crit} are



¹¹In a simplicial set, it may happen that σ is a "multiple" face of τ . i.e., $\sigma = \partial_i \tau$ holds for several indices *i*. In such case, we connect τ to σ with multiple edges in the V ∂ -graph, one edge for each such index *i*.

defined based on V, and they are locally effective assuming that X and V are locally effective in a natural sense.

Polynomially Bounded Vector Fields We need a polynomial-time version of this result. Let *V* be an admissible vector field *V* on a locally polynomial-time simplicial set; we assume that both *X* and *V* are parameterized by a set \mathcal{I} , as in the definition of a locally polynomial-time simplicial set.¹² For $\sigma \in X^{ndg}$, let reach_V(σ) (or just reach(σ) if *V* is understood) denote the set of all simplices reachable from σ by a directed path in the *V* ∂ -graph.

Let us say that V is *polynomially bounded* if the following hold:

- (PBV1) An algorithm is available that, given $I \in \mathcal{I}$ and a simplex $\sigma \in X(I)_k^{\text{ndg}}$, classifies σ as source, target, or critical. In the source case, it also returns the simplex $V(\sigma)$. The running time is polynomial in size(I) + size(σ) for every fixed k.
- (PBV2) For every fixed k and every $\sigma \in X(I)_k^{\text{ndg}}$, the sum of encoding sizes of all simplices in reach_V(σ) is bounded by a polynomial in size(I) + size(σ).

Theorem 3.1 If X is a locally polynomial-time simplicial set and V is a polynomially bounded vector field on X such that, for every k, the sum of the encoding sizes of all k-dimensional critical simplices is polynomially bounded (in size(I)), then X can be turned into a simplicial set with polynomial-time homology.

Proof The proof essentially follows by inspecting the work of Forman [9] (mainly Sects. 7 and 8) and making simple observations about the computation of the relevant maps. For the reader's convenience, we provide a self-contained presentation; this seems simpler and not much longer than referring to the appropriate claims in Forman's paper, introducing his notation, etc. Our presentation is, similar to that of Forman, mainly in a combinatorial language. We refer to [17] for two other, more algebraic variants of essentially the same proof.

Throughout the proof, we keep the parameterization of *X* and *V* by \mathcal{I} implicit.

To provide the desired reduction from $C_* := C_*(X)$, we need to define the target chain complex C_*^{crit} and provide the three maps f, g, h as in the definition of a reduction. We begin with introducing several auxiliary maps and checking some of their properties.

The vector field *V* induces a sequence $V_{\#} = (V_{\#k})_{k=-\infty}^{\infty}$ of homomorphisms $V_{\#k}$: $C_k \to C_{k+1}$, as follows: for a source *k*-simplex σ , we have $V_{\#k}(1 \cdot \sigma) := (-1)^{i+1} \cdot V(\sigma)$, where *i* is the unique index with $\sigma = \partial_i V(\sigma)$, and for σ target or critical, we have $V_{\#k}(1 \cdot \sigma) := 0$.

Next, we introduce a chain map $\Phi: C_* \to C_*$ by

$$\Phi := 1 + V_{\#}d + dV_{\#},$$

F∘⊏¬ ∰ Springer ⊔□

¹²Of course, for the main result of this paper, polynomial-time homology for $K(\mathbb{Z}, 1)$, parameterization is not needed, but we need it if we want to have a general tool for obtaining polynomial-time homology from a vector field.



Fig. 3 Forming the image $\Phi(1 \cdot \sigma)$

where 1 stands for the identity chain map and d is the differential of C_* . It is easy to check that Φ is a chain map: indeed, $d\Phi = d + dV_{\#}d + ddV_{\#} = d + dV_{\#}d = \Phi d$ (using dd = 0).

For the proof, it is important to understand how Φ works. We will thus discuss how the image $\Phi(1 \cdot \sigma)$ is formed, depending on the type of a *k*-simplex σ .

1. The simplest case is σ a *target* simplex; see Fig. 3 left. Then $V_{\#}(1 \cdot \sigma) = 0$, and thus $\Phi(1 \cdot \sigma) = 1 \cdot \sigma + \sum_{i=0}^{k} V_{\#k-1}((-1)^i \cdot \partial_i \sigma)$. So we consider all faces σ' of σ , with the appropriate signs, and apply $V_{\#}$ to them. Only the σ' that are sources may contribute to the image (and then (σ, σ') are edges of the $V\partial$ -graph), and $\Phi(1 \cdot \sigma)$ is supported only on target simplices.

Moreover, we observe that, crucially, the coefficient of σ in $\Phi(1 \cdot \sigma)$ is 0; indeed, if *j* is the unique index with $V^{-1}(\sigma) = \partial_j \sigma$, then we have $V_{\#}((-1)^j \cdot \partial_j \sigma) = (-1)^{j+1}(-1)^j \cdot \sigma = -1 \cdot \sigma$, which cancels out with the $1 \cdot \sigma$ coming from the 1 in the definition of Φ . (Here we rely on the condition that $V^{-1}(\sigma)$ is a *regular* face of σ from the definition of discrete vector field, since we need the coefficient of $V^{-1}(\sigma)$ in $d(1 \cdot \sigma)$ to be invertible, i.e., equal to ± 1 .)

Summarizing, $\Phi(1 \cdot \sigma)$ consists of the target simplices reachable from σ in exactly two steps in the V ∂ -graph, with appropriate signs.

- 2. For σ a *critical* simplex we find, by a similar reasoning, that $\Phi(1 \cdot \sigma)$ consists of σ with coefficient 1, plus all the (target) simplices reachable from σ in exactly two steps in the $V\partial$ -graph, again with appropriate signs.
- 3. Finally, for σ a source, both the dV_# and the V_#d terms may make a nonzero contribution to Φ(1 · σ). For dV_# (going first up, then down), we get, with appropriate signs, all the *source* simplices reachable from σ in exactly two steps in the V∂-graph, with σ itself cancelled out, plus some additional target and critical simplices (here we do not follow the edges of the V∂-graph—that is why the arrows are dotted in the picture). For V_#d (first down, then up), we get only target simplices.

Next, we define $\Phi^{\infty} = \lim_{N \to \infty} \Phi^N$ as the *stabilization* of Φ ; that is, given a *k*-chain *c*, we compute $\Phi(c)$, $\Phi(\Phi(c))$, etc., until we reach a chain \tilde{c} with $\Phi(\tilde{c}) = \tilde{c}$, and we set $\Phi^{\infty}(c) := \tilde{c}$.



To check that the iterations of Φ indeed stabilize after finitely many steps, it suffices to consider the case $c = 1 \cdot \sigma$, and then the stabilization follows easily from the above discussion of the action of Φ (and from the admissibility of the vector field *V*). Moreover, we can see that the chains in im Φ^{∞} are supported only on critical and target simplices.

We also need to check that Φ^{∞} is computable in polynomial time. In order to compute $\Phi^{\infty}(1 \cdot \sigma)$ (which is sufficient), we just compute the iterations $\Phi^{N}(1 \cdot \sigma)$, N = 1, 2, ..., until they stabilize. We observe that each simplex in the support of some $\Phi^{N}(1 \cdot \sigma)$ can be reached from σ by following a directed path in the $V\partial$ -graph, then possibly going to a face of the current simplex (a step corresponding to a dotted arrow in Fig. 3), and then again following a directed path in the $V\partial$ -graph. Hence, by the polynomial boundedness of the vector field V, the stabilization occurs for N at most polynomially large, and the sum of the encoding sizes of all simplices in the supports of all chains encountered along the way is also polynomially bounded (essentially by the square of the bound in condition (PBV2)).

Each *coefficient* in the chain $\Phi^{N+1}(1 \cdot \sigma)$ is the sum of O(k) coefficients in $\Phi^N(1 \cdot \sigma)$. So each coefficient in $\Phi^N(1 \cdot \sigma)$ is bounded by $\exp(O(N))$, and hence its size (number of bits) is at most O(N). Therefore, Φ^{∞} is indeed polynomial-time computable.

Now we define an auxiliary chain complex C_*^{Φ} ; we set $C_k^{\Phi} := \operatorname{im} \Phi_k^{\infty} \subseteq C_k$. Equivalently, as is easily seen, $C_k^{\Phi} = \{c \in C_k : \Phi(c) = c\}$. The differential of C_k^{Φ} is the restriction of the differential of C_* (this works since Φ is a chain map). Let $i : C_k^{\Phi} \to C_*$ be the inclusion (which is a chain map).

Next, we come to the definition of C_*^{crit} ; as was announced above, the chain group C_k^{crit} is the free Abelian group (\mathbb{Z} -module) with the set of the *k*-dimensional critical simplices in *X* as a basis. It remains to define the differential.

First we let $j_k : C_k^{\Phi} \to C_k^{crit}$ be the homomorphism that restricts a chain $c \in C_k^{\Phi}$ to the critical simplices (i.e., for $c = \sum_{\sigma \in X_k} \alpha_{\sigma} \cdot \sigma$, we set $j_k(c) := \sum_{\sigma \in X_k \cap C} \alpha_{\sigma} \cdot \sigma$). We observe that Φ_k^{∞} , viewed as a homomorphism $C_k^{crit} \to C_k^{\Phi}$, is an inverse to j_k . Indeed, from the description of Φ given above, it is easy to see that for σ critical, $\Phi^{\infty}(1 \cdot \sigma) = 1 \cdot \sigma + c'$ for some c' supported on target simplices, and from this the claim follows.

Hence each C_k^{crit} is isomorphic to C_k^{Φ} , and the differential d^{crit} of C_*^{crit} can be defined so as to make j and Φ^{∞} mutually inverse chain isomorphisms; explicitly, $d^{\text{crit}} := j d \Phi^{\infty}$. This finishes the definition of the target chain complex for the desired reduction; it is clear that the matrices of the differential d^{crit} are polynomially computable, provided that the total encoding size of the critical simplices is polynomial in each dimension.

It remains to define the maps f, g, h in the reduction. The following diagram summarizes the relevant chain complexes and maps defined so far, plus f, g, h:





As the diagram suggests, we put $f := j\Phi^{\infty}$ and $g := i\Phi^{\infty}$. Then, since j and Φ^{∞} are mutually inverse and $\Phi^{\infty}i = 1$, we have fg = 1, as required by the definition of a reduction, and $gf = i\Phi^{\infty}$.

The chain homotopy h of $i\Phi^{\infty}$ with the identity (Forman uses the letter L for this map) is now defined as the stabilization of the maps

$$-V_{\#}(1 + \Phi + \Phi^2 + \dots + \Phi^N), \quad N = 1, 2, \dots$$

To see that these iterations indeed stabilize on each chain $1 \cdot \sigma$, we recall that for sufficiently large N, $\Phi^N(1 \cdot \sigma)$ is supported only on critical and target simplices, and $V_{\#}$ sends such chains to 0. By essentially the same argument as that for the computability of Φ^{∞} , we also see that each h_k is computable in polynomial time.

We need to verify that *h* is the required chain homotopy, i.e., $dh + hd = 1 - i\Phi^{\infty}$. This is a simple formal calculation (showing where the formula for *h* comes from), which we leave to the reader (also see [9, proof of Th. 7.3]).

As the last step, we want to check the conditions fh = 0, hg = 0, and hh = 0. To this end, we note that the chains in mh are supported only on target simplices. Moreover, if c is a chain supported only on target and critical simplices, then $\Phi(c)$ has the same property, and hence h(c) = 0. These two properties immediately give hh = 0. Similarly, im $g = im \Phi^{\infty}$ is supported only on target and critical simplices, and hence hg = 0. Finally, we have seen that Φ^{∞} maps target simplices to 0, and so does $f = j\Phi^{\infty}$, which gives fh = 0 and concludes the proof of Theorem 3.1.

4 A Polynomially Bounded Vector Field for $K(\mathbb{Z}, 1)$

Here we finally get to the combinatorial core of the paper; we will provide a polynomially bounded vector field for $K(\mathbb{Z}, 1)$.

A Simple Composition of Vector Fields For the sake of presentation, it will be easier to split the vector field into two parts. Roughly speaking, the first part will get rid of all negative components in the considered sequences $[a_1|\cdots|a_k]$, and the second part will do the rest.

Here is the way of "splitting into two parts" in a general setting. Let *X* be a simplicial set, let V_1 be a vector field on *X*, with the set C_1 of critical simplices, and suppose that C_1 is closed under the face operators (each face of a critical simplex is again critical, or degenerate). Let *Y* be the simplicial subset of *X* induced by C_1 (i.e., its nondegenerate simplices are the critical simplices of V_1), and let V_2 be a vector field on *Y*.

Then we can define a "composition" V of V_1 and V_2 in the obvious way; formally, if we regard a vector field a set of ordered pairs, we simply set $V := V_1 \cup V_2$. Clearly, V is a vector field, and it is easily seen that V_1 , V_2 admissible imply V admissible, and similarly for polynomial boundedness.

In the case of $X = K(\mathbb{Z}, 1)$, the role of *Y* will be played by the simplicial set whose simplices are the integer sequences with all terms *nonnegative*. With some abuse of the usual notation, we will denote this simplicial set by $K(\mathbb{N}, 1)$.



The first vector field will be denote by V_{bs} and called the *bubblesort field*, since directed paths in its V ∂ -graph resemble the computation of a sorting algorithm called Bubblesort. Its critical simplices are integer sequences with all entries positive.

The second vector field is defined on $K(\mathbb{N}, 1)$, and it has only two critical simplices [] and [1], the same as the Eilenberg–MacLane field V_{EML} . We call it the *bit-chipping field* and denote it by V_{bch} .

Let us remark that one can consider composition of vector fields in a more general and more flexible setting, as is done in [17], but for our purposes, the simple notion above suffices.

4.1 The Bubblesort Field

Translating Positive Sequences to Sorted Sequences To define the vector field V_{bs} , it is convenient to consider a different representation of the simplices of $K(\mathbb{Z}, 1)$. Namely, we represent a k-dimensional simplex $\sigma = [a_1|\cdots|a_k]$ by a (k + 1)-tuple (b_0, b_1, \ldots, b_k) , where $b_0 \in \mathbb{Z}$ can be chosen arbitrarily and $b_i := b_{i-1} + a_i$, $i = 1, 2, \ldots, k$. Thus, each σ is represented as an equivalence class of (k + 1)-tuples of integers, where two (k + 1)-tuples are equivalent if their difference is of the form (a, a, \ldots, a) (all components equal). We denote the equivalence class of (b_0, \ldots, b_k) by $[b_0, \ldots, b_k]$.

This correspondence between simplices of the form $[a_1|\cdots|a_k]$ and equivalence classes of (k + 1)-tuples is obviously bijective. Nondegenerate simplices $[a_1|\cdots|a_k]$, i.e., those with no zero component, translate to $[b_0, \ldots, b_k]$ with $b_{i-1} \neq b_i$, $i = 1, 2, \ldots, k$.

A (nondegenerate) simplex from $K(\mathbb{N}, 1)$ corresponds to $[b_0, \ldots, b_k]$ with strictly increasing components, i.e., $b_0 < b_1 < \cdots < b_k$. The face operators become extremely simple in this notation: ∂_i corresponds to deleting the *i*th component.

The Field As was already announced, the critical simplices of V_{bs} are the $[b_0, \ldots, b_k]$ with $b_0 < \cdots < b_k$. If $\sigma = [b_0, \ldots, b_k]$ is not critical, we look at the smallest ℓ such that $b_{\ell} > b_{\ell+1}$; let us call it the *leading index* of σ . Let us write $v = b_{\ell}$ and $u = b_{\ell+1}$. We consider the maximal contiguous segment in the sequence b_0, b_1, \ldots starting at the ℓ th position and containing only v's and u's; formally, we take the largest $m \ge \ell + 1$ such that $b_i \in \{u, v\}$ for all $i = \ell, \ell + 1, \ldots, m$, and either $b_{m+1} \notin \{u, v\}$ or m = k. We call $b_{\ell}, b_{\ell+1}, \ldots, b_m$ the *leading alternating segment* of σ (indeed, there can be no two consecutive u's or v's, since this would mean that σ is degenerate), and we denote it by LAS (σ) .

Then we let σ be a source if LAS(σ) ends with u, and otherwise, σ is a target. For a source σ , still with u, v, m as above, we set

$$\tau = V_{\rm bs}(\sigma) := [b_0, \dots, b_m, v, b_{m+1}, \dots, b_k], \tag{2}$$

i.e., $V_{\rm bs}$ inserts another v just after LAS(σ).

With $\tau = V_{bs}(\sigma)$ as in the just given definition, we have $\sigma = \partial_{m+1}\tau$, and m+1 is easily seen to be the only index *i* with $\sigma = \partial_i \tau$ (thus, σ is a regular face of τ). Moreover, σ can be uniquely reconstructed from τ (delete the last element of LAS(τ)), and so V_{bs} is indeed a discrete vector field.

Next, we observe that once we show that V_{bs} is admissible, it becomes obvious that it is also polynomially bounded. This is because the boundary operators only delete components and the vector field duplicates them, and so any simplex reachable from a given k-dimensional σ is made of the components of σ . Hence at most $(k + 1)^{k+1}$ distinct source simplices are reachable from σ , which is a constant for k fixed.

It remains to prove admissibility, which is tricker than it might seem. Let us consider a source simplex $\sigma = [b_0, \ldots, b_{\ell-1}, v, \ldots, u, b_{m+1}, \ldots, b_k], b_0 < b_1 < \cdots < b_{\ell-1} < v > u$, where the part between the v and u is the LAS. We set $\tau = V_{\text{bs}}(\sigma)$, and ask for which *i*'s the simplex $\sigma' = \partial_i \tau$ can again be a source simplex (in this case we say that σ' arises from σ by a *double move*).

If $LAS(\sigma') = LAS(\tau)$, then σ' is a target simplex, and so ∂_i must change $LAS(\tau)$. It cannot delete elements from the middle of $LAS(\tau)$, since the result would be degenerate, and it cannot delete the final v, since this was inserted by V_{bs} .

Thus, one possibility is $i = \ell$, in which case σ' is obtained from σ by appending v to the end of the LAS and deleting the initial v of the LAS. Let us call this a *switching double move*. This is the "intended" type of double moves that do the bubble-sorting, provided that the LAS has length 2; for example, $\sigma = [3, 1, 2]$ is transformed to $\sigma' = [1, 3, 2]$. A switching double move may also occur for LAS(σ) of length 4 or more, if the deletion of the initial v creates a new LAS; i.e., if $b_{\ell-1} > u$. An example is $\sigma = [2, 3, 1, 3, 1], \sigma' = [2, 1, 3, 1, 3]$.

However, there is a second, less obvious possibility for a double move: if the sequence $[b_{m+1}, b_{m+2}, ...]$ following LAS (τ) has the form [x, u, v, u, v, ..., u, y, ...], $x, y \notin \{u, v\}$, or the form [x, u, v, ..., u], then we can also have i = m + 2. In this case, ∂_{m+2} deletes the component following the LAS, and produces a longer LAS. We call this an *appending double move*. For example, for $\sigma = [2, 3, 1, 4, 1, 3, 1]$, the switching double move yields $\sigma' = [2, 1, 3, 4, 1, 3, 1]$ and the appending one yields $\sigma' = [2, 3, 1, 3, 1, 3, 1]$.

If we follow a sequence of directed edges in the $V\partial$ -graph starting at some source simplex $\tilde{\sigma}$, and if all source simplices encountered along the way have LAS of length 2, then the path has a bounded length, since all the double moves are switching in this case, and each of them decreases the number of inversions (i.e., pairs (i, j) with i < j and $b_i > b_j$) in the current source simplex.

The following lemma shows that if $LAS(\tilde{\sigma})$ has length greater than 2, then every sequence of double moves starting at $\tilde{\sigma}$ finishes after a finite number of steps, and this already implies the admissibility of V_{bs} . All the difficulty of the lemma is in getting the statement right; the proof is routine.

Lemma 4.1 Let $\tilde{\sigma} = [b_0, \dots, b_{\ell-1}, b_{\ell} = v, u, v, \dots, u, \dots]$, $b_0 < \dots < b_{\ell} > u$, be a source simplex with LAS($\tilde{\sigma}$) of length greater than 2. Then every source σ obtainable from $\tilde{\sigma}$ by a sequence of double moves has the following structure: $[\beta_0, \beta_1, \dots, \beta_{\ell}, \gamma]$, where each β_i is a block of length $k_i \ge 1$ starting with b_i and possibly continuing with u, b_i, u, b_i, \dots (alternations of b_i and $u, u < b_i$), and γ is a possibly empty block that does not start with u. The sequence $(k_0, k_1, \dots, k_{\ell})$ has the form

$$(\underbrace{1,1,\ldots,1}_{j},k_{j},k_{j+1},\ldots,k_{\ell}),$$



where $k_j \ge 2$ is even, while all of the other k_i are odd, and there is at least one $k_i \ge 3$. In each double move of a sequence starting at $\tilde{\sigma}$, either j decreases, or it stays the same and k_i increases. Thus, each such sequence is finite.

Proof The initial $\tilde{\sigma}$ clearly has the claimed form. Let us assume that σ is of this form, and let a source σ' be obtained from it by a double move.

We have $LAS(\sigma) = \beta_j$, of even length $k_j \ge 2$. If the double move is switching, then

$$\sigma' = [b_0, b_1, \ldots, b_{j-1}, u, \underbrace{b_j, \ldots, u, b_j}_{k_j-1}, \beta_{j+1}, \ldots, \gamma].$$

If we had j = 0 or $b_{j-1} < u$, then LAS(σ') would be either the block b_j, \ldots, u, b_j of odd length $k_j - 1$ (for $k_j \ge 4$), or, for $k_j = 2$, another $\beta_i, i > j$, of odd length $k_i \ge 3$ (guaranteed to exist by the inductive assumption). In both cases σ' would be target, and so $b_{j-1} > u$. Then σ' has the claimed structure $[\beta'_0, \ldots, \beta'_\ell, \gamma]$, with j' = j - 1, $\beta'_i = \beta_i$ for all $i \notin \{j - 1, j\}, \beta'_{j-1} = b_{j-1}, u$ of length $k'_{j-1} = 2$, and β'_j of odd length $k'_i = k_j - 1$. So j has decreased.

For an appending double move, we distinguish two cases. For $j < \ell$, there is at least one more block β_{j+1} following β_j in σ , with $k_{j+1} \ge 3$ (since β_{j+1} must have an *u* to append to β_j), and we have

$$\sigma' = [b_0, b_1, \dots, b_{j-1}, \underbrace{b_j, \dots, u, b_j, u}_{k_j+2}, \underbrace{b_{j+1}, u, \dots, b_{j+1}}_{k_{j+1}-2}, \beta_{j+2}, \dots, \gamma].$$

This is the claimed structure with j' = j, $k'_{j} = k_{j} + 2$, and $k'_{j+1} = k_{j+1} - 2$.

Finally, if $j = \ell$, then γ has to start with x, u, \ldots , and here we get $j' = j = \ell$ and $k'_{\ell} \ge k_{\ell} + 2$ (depending on the number of u, v alternations in γ following x).

A Lower Bound Although the bubble-sorting process itself is only quadratic, it turns out that $|\text{reach}_{V_{bs}}(\tilde{\sigma})|$ for a suitable source simplex $\tilde{\sigma}$ may indeed be exponential in k, and thus the bound $(k + 1)^{k+1}$ claimed above is not so far off the mark. Mainly to illustrate the behavior of the vector field V_{bs} , we indicate the lower bound via a concrete example without proof. Namely, from

$$\tilde{\sigma} = [2, 3, 4, 5, 6, 7, 1, 7, 1, 7, 1, 7, 1, 7, 1, 7, 1, 7, 1]$$

we can reach source simplices such as [2, 1, 2, 1, 3, 1, 3, 4, 5, 1, 5, 1, 5, 6, 7, 1, 7, 1, 7]. Such simplices have 6 blocks (denoted by β_0, \ldots, β_5 in the proof above), and we can choose the block lengths at will, with the obvious restrictions (the total length is fixed, and the block lengths are all odd except for the first one). In an analogous construction with 6 replaced by an arbitrary integer *b* we take k = 3b and obtain a lower bound exponential in *k*.

4.2 The Bit-Chipping Field

Here we return to the "bar" notation $[a_1|a_2|\cdots|a_k]$, and we will consider only simplices of $K(\mathbb{N}, 1)$, which means $a_i \ge 1$ for all *i*.

The Anatomy of a Simplex Let $\sigma = [a_1|a_2|\cdots|a_k]$ be a nondegenerate k-simplex of $K(\mathbb{N}, 1)$. We introduce the following terminology.

- Let $p = p(\sigma) \in \{0, 1, ..., k\}$ be the largest index such that $a_1, ..., a_p$ are all powers of 2 and $a_1 \le a_2 \le \cdots \le a_p$. The sequence $a_1|a_2|\cdots|a_p$ is called the *nondecreasing dyadic part* of σ . If $1 \le p < k$ and $a_p > a_{p+1}$, then p is called the *peak* of σ ; otherwise, σ has no peak.
- Let q = q(σ) ∈ {0, 1, ..., k} be the largest index such that a₁, ..., a_q are all powers of 2 (thus, q ≥ p). The sequence a₁|a₂|···|a_q is called the *dyadic part* of σ. If q = k, then σ is called *fully dyadic*. If, on the other hand, q < k, then q + 1 is the *breakpoint* of σ and a_{q+1} is the *breakpoint value* of σ (which is not a power of 2). The sequence a_{q+2}|a_{q+3}|···|a_k is the *right part* of σ.

Here are two concrete examples:



The Vector Field We define a vector field V_{bch} on $K(\mathbb{N}, 1)$. There are two types of source simplices.

(a) The first type of source simplices are the simplices that are not fully dyadic and have no peak. Thus, all of the dyadic part is nondecreasing (i.e., p = q; we also admit p = q = 0) and the breakpoint value is larger than the last element of the dyadic part. Explicitly, they are of the form

$$\sigma = \left[2^{i_1} \mid 2^{i_2} \mid \dots \mid 2^{i_q} \mid b \mid a_{q+2} \mid \dots \mid a_k\right],$$

 $2^{i_1} \leq 2^{i_2} \leq \cdots \leq 2^{i_q} < b$. In this case we set

$$V_{\rm bch}(\sigma) = \tau := \left[2^{i_1} \mid 2^{i_2} \mid \dots \mid 2^{i_q} \mid {\rm lpow}(b) \mid {\rm ltrim}(b) \mid a_{q+2} \mid \dots \mid a_k\right], \quad (3)$$

where pow(b) is the largest power of 2 not exceeding b, and ltrim(b) := b - pow(b). That is, τ is obtained by splitting the breakpoint value b into two components, pow(b) and ltrim(b); informally, we can think of this as "chipping off" the leading bit of b.

We observe that each target simplex τ as defined above has a peak, namely, $p(\tau) = q(\sigma) + 1$, and in particular, τ has a nonempty dyadic part (but it may happen that the dyadic part of τ is longer than the nondecreasing dyadic part, since $\operatorname{ltrim}(b)$ may be a power of two).

(b) The second type of source simplices are the fully dyadic simplices $\sigma = [2^{i_1} | 2^{i_2} | \cdots | 2^{i_k}]$ with $2^{i_1} \le 2^{i_2} \le \cdots \le 2^{i_{k-1}} < 2^{i_k}$ with $i_k \ge 1$ (this last condition is important only for k = 1). In this case we set

$$\tau = V_{\rm bch}(\sigma) := \left[2^{i_1} \mid 2^{i_2} \mid \dots \mid 2^{i_{k-1}} \mid 2^{i_k-1} \mid 2^{i_k-1} \mid 2^{i_k-1} \right]; \tag{4}$$

i.e., we split the last component of σ into two equal halves.

Lemma 4.2 *This definition indeed yields a vector field, and the only critical simplices are* [] *and* [1].

Proof Let us consider an arbitrary simplex τ . If it is not fully dyadic and is not a source simplex, then it has a peak, and thus it has the form $\tau = [2^{i_1} | \cdots | 2^{i_p} | c_{p+1} | \cdots | c_{k+1}]$ with $2^{i_1} \leq \cdots \leq 2^{i_p} > c_{p+1}$. This equals $V_{bch}(\sigma)$ for $\sigma = [2^{i_1} | \cdots | 2^{i_p} | (2^{i_{p-1}} | 2^{i_p} + c_{p+1} | c_{p+2} | \cdots | c_{k+1}]$. Thus, τ is a target simplex and there is exactly one edge $(\sigma, \tau) \in V_{bch}$. Moreover, we have $\sigma = \partial_p \tau$, while $\partial_j \tau \neq \sigma$ for $j \neq p$, so σ is a regular face of τ as needed.

Next, if τ is fully dyadic and has a peak p, i.e., $\tau = [2^{i_1} | \cdots | 2^{i_{k+1}}], 2^{i_1} \leq \cdots \leq 2^{i_p} > 2^{i_{p+1}}$, then τ is again a target simplex with $\tau = V_{\text{bch}}(\sigma)$ for $\sigma = [2^{i_1} | \cdots | 2^{i_{p-1}} | 2^{i_p} + 2^{i_p+1} | 2^{i_{p+1}} | \cdots | 2^{i_{k+1}}]$ (here $2^{i_p} + 2^{i_{p+1}}$ is the breakpoint value). Again, j = p is the only index with $\partial_j \tau = \sigma$.

The last remaining case is a fully dyadic τ with no peak, which must be nondecreasing. If it is not a source simplex, then either we have one of the cases [], [1], or $k \ge 2$ and the last two components of τ are equal, which means that τ is of the form (4) and σ can again be uniquely reconstructed from it. We have $\sigma = \partial_j \tau$ for the unique index j = d - 1.

*Preparations for Analyzing V*_{bch} It will be convenient to work mainly with the target simplices. Thus, given a target simplex τ , we let t-reach(τ) \subset reach(τ) be the set of all target simplices reachable from τ .

First we will classify all possible target simplices τ' reachable from a given target simplex τ by two steps in the V ∂ -graph; in other words, the τ' of the form $V_{bch}(\partial_j \tau)$ for some *j*. This is a straightforward, if somewhat lengthy, case analysis. The subsequent proofs of admissibility and polynomial boundedness will use this classification. It would be nice to avoid considering so many cases, but one needs to be careful in the analysis: for several other candidate vector fields we have tried, "most" cases apparently worked fine, but those fields failed in what seemed like minor details.

Lemma 4.3 Let $\tau = [a_1|a_2|\cdots|a_k]$ be a k-dimensional target simplex. If τ is not fully dyadic, we can write it in the form

$$[2^{i_1} | 2^{i_2} | \cdots | 2^{i_p} | 2^{i_{p+1}} | \cdots | 2^{i_q} | b | a_{q+2} | \cdots | a_k],$$

where b is not a power of 2, $2^{i_1} \leq \cdots \leq 2^{i_p}$, $p \geq 1$, $p \leq q \leq k-1$, and either $2^{i_p} > 2^{i_{p+1}}$ (if p < q) or $2^{i_p} > b$ (for p = q). Let τ' be a target simplex of the form $V_{\text{bch}}(\partial_j \tau)$ for some j, where $\sigma = \partial_j \tau$ is a (k-1)-dimensional source simplex. Then τ' has one of the following forms:

(A) If p = 1 and $2^{i_2} \leq \cdots \leq 2^{i_q} < b$, then we can have

 $\tau' = \left[2^{i_2} \mid \dots \mid 2^{i_q} \mid \mathsf{Ipow}(b) \mid \mathsf{ltrim}(b) \mid a_{q+2} \mid \dots \mid a_k\right]$

(we drop the first component and split b). Example: $\tau = [2^2|1|2|7], \tau' = [1|2|2^2|3].$

(B) If $i_j < i_{j+1}$ for some $j, 1 \le j \le p-1$, then we can have

$$\tau' = \left[2^{i_1} | \cdots | 2^{i_{j-1}} | 2^{i_{j+1}} | 2^{i_j} | 2^{i_{j+2}} | \cdots | 2^{i_q} | b | a_{q+2} | \cdots | a_k\right]$$

(the entries 2^{i_j} and $2^{i_{j+1}}$ are swapped). Example: $\tau = [1|2^2|2|7], \tau' = [2^2|1|2|7]$. (C) If $q \ge p+2$, $i_p - 1 = i_{p+1} = i_{p+2} < i_{p+3} \le \cdots \le i_q$, and $2^{i_q} < b$, then we can have

$$\tau' = \left[2^{i_1} | \cdots | 2^{i_p} | 2^{i_p} | 2^{i_{p+3}} | \cdots | 2^{i_q} | \operatorname{lpow}(b) | \operatorname{ltrim}(b) | a_{q+2} | \cdots | a_k\right]$$

(two components following the peak are merged and b is split). Example: $\tau = [2|1|1|2|7], \tau' = [2|2|2|2^2|3].$

(D) If $q \ge p + 2$ and $i_{p+2} \ge i_p > i_{p+1}$, then we can have

$$\tau' = \left[2^{i_1} | \cdots | 2^{i_p} | 2^{i_{p+2}} | 2^{i_{p+1}} | 2^{i_{p+3}} | \cdots | 2^{i_q} | b | a_{q+2} | \cdots | a_k\right]$$

(the entries $2^{i_{p+1}}$ and $2^{i_{p+2}}$ are swapped). Example: $\tau = [2|1|2^2|7], \tau' = [2|2^2|1|7].$

(E) If q = p + 1, $b' = 2^{i_{p+1}} + b$ satisfies $b' \ge 2^{i_p}$, and b' is not a power of 2, then we can have

$$\tau' = [2^{i_1} | \cdots | 2^{i_p} | \operatorname{lpow}(b') | \operatorname{ltrim}(b') | a_{q+2} | \cdots | a_k].$$

Example: $\tau = [2^3|2|7], \tau' = [2^3|2^3|1].$

(F) If the situation is as in (E) except that $b' = 2^i$ is a power of 2, then we can have

$$\tau' = V_{\rm bch}([2^{i_1} | \cdots | 2^{i_p} | 2^i | a_{q+2} | \cdots | a_k])$$

(note that here we do not write out τ' explicitly, since there are still several cases to distinguish depending on the right part of τ , but we will not need to discuss them explicitly). Example: $\tau = [2^3|1|7|19]$, $\tau' = [2^3|2^3|2^4|3]$.

(G) If $q = p \le k - 2$, $b' := b + a_{q+2} \ge 2^{i_p}$, and b' is not a power of 2, then we can have

$$\tau' = [2^{i_1} | \cdots | 2^{i_p} | \operatorname{Ipow}(b') | \operatorname{Itrim}(b') | a_{q+3} | \cdots | a_k].$$

Example: $\tau = [2^3|7|4], \tau' = [2^3|2^3|3].$

(H) If the conditions are as in (G) except that $b' = 2^i$ is a power of 2, then we can have

$$\tau' = V_{\rm bch} \left(\left[2^{i_1} | \cdots | 2^{i_p} | 2^i | a_{q+3} | \cdots | a_k \right] \right)$$

(as in (F), we need not write out τ' explicitly). Example: $\tau = [2^3|7|1|7]$, $\tau' = [2^3|2^3|2^2|3]$.

(I) If q = p = k - 1 and either p = 1 or $i_{p-1} < i_p$, then we can have

$$\tau' = \left[2^{i_1} | \cdots | 2^{i_{p-1}} | 2^{i_p-1} | 2^{i_p-1} \right].$$

Example: $\tau = [2|2^3|7], \tau' = [2|2^2|2^2].$



If $\tau = [2^{i_1} | \cdots | 2^{i_k}]$ is fully dyadic, then either p < k (τ has a peak), or p = k (τ is nondecreasing) and $i_{k-1} = i_k$. In the peak case, we have the following possibilities for $\tau' = V_{\text{bch}}(\partial_j \tau)$:

(dA) If p = 1 and $i_2 \le i_3 \le \cdots \le i_{k-1} < i_k$, we can have

 $\tau' = \left[2^{i_2} \mid \dots \mid 2^{i_{k-1}} \mid 2^{i_k-1} \mid 2^{i_k-1}\right]$

(deleting the first entry of τ and splitting the last).

(dB) For $1 \le j \le p-1$ and $i_j < i_{j+1}, \tau'$ can be obtained by swapping 2^{i_j} and $2^{i_{j+1}}$. (dC) If $i_p - 1 = i_{p+1} = i_{p+2} < i_{p+3} \le \cdots \le i_{k-1} < i_k$, we can have

$$\tau' = \left[2^{i_1} | \cdots | 2^{i_p} | 2^{i_p} | 2^{i_{p+3}} | \cdots | 2^{i_{k-1}} | 2^{i_k-1} | 2^{i_k-1}\right]$$

(merging two equal entries and splitting the last).

(dD) For $i_{p+2} \ge i_p > i_{p+1}$, τ' can be obtained from τ by swapping $2^{i_{p+1}}$ and $2^{i_{p+2}}$.

Finally, if a fully dyadic τ has no peak, we have the possibility (dB) for τ' and the following additional one:

(dI) If k = 2 or $i_{k-2} < i_{k-1}$, then we can have

$$\tau' = \left[2^{i_1} | \cdots | 2^{i_{k-2}} | 2^{i_{k-1}-1} | 2^{i_{k-1}-1}\right]$$

(drop the last component and split the previous one).

Proof As was already mentioned, the proof is totally straightforward and could probably be left to the reader. Yet, since getting used to the definitions and notation probably needs some practice, we chose to present the proof.

As in the lemma, we first consider τ not fully dyadic. If $\sigma = \partial_j \tau$ is a source simplex, then it has no peak, and thus the operation ∂_j has to "destroy" the peak of τ in some way. In particular, we have $j \le p + 1$, for otherwise, the peak of τ is also present in $\partial_j \tau$. We just need to discuss the values of j in this range.

For j = 0, ∂_0 removes the first coordinate, and this may destroy the peak only for p = 1. For p = 1, σ is a source iff $2^{i_{p+1}} \leq \cdots \leq 2^{i_q} < b$ (this condition is void for q = 1), and if this holds, then τ' is as in (A).

If $1 \le j \le p-1$, $\sigma = [2^{i_1}|\cdots|2^{i_{j-1}}|2^{i_j}+2^{i_{j+1}}|2^{i_{j+2}}|\cdots|2^{i_q}|b|\cdots]$. In this case, if $i_j = i_{j+1}$, then $2^{i_j} + 2^{i_{j+1}}$ is a power of two, σ necessarily has a peak, and thus it is not a source. So $i_j < i_{j+1}$; then σ is a source and $2^{i_j} + 2^{i_{j+1}}$ is the breakpoint value, and τ' is as in (B).

Next, we consider j = p. Here the *p*th component of σ is $2^{i_p} + 2^{i_{p+1}}$ (for q > p) or $2^{i_p} + b$ (for p = q). In both of these cases the *p*th component is not a power of 2 (since *p* was the peak of τ), hence *p* is the breakpoint of σ , and so $V_{bch}(\sigma) = \tau$. Therefore, j = p does not contribute any τ' .

Finally, we need to discuss j = p + 1. Here the sum of the two entries of τ following the peak must greater or equal to 2^{i_p} (and, in particular, $p \le k - 2$), for otherwise, p would be a peak in σ . We consider three cases, depending on how many of these two entries are powers of 2.

First, if $q \ge p + 2$, then the peak is followed by $2^{i_{p+1}}$ and $2^{i_{p+2}}$ in τ . If $2^{i_{p+1}} + 2^{i_{p+2}} = 2^{i_p}$, then $i_{p+1} = i_{p+2} = i_p - 1$. Then σ begins with $[2^{i_1}|\cdots| 2^{i_p}|2^{i_p}|2^{i_{p+3}}|\cdots|2^{i_q}|b|\cdots]$, and since it has no peak, the dyadic part is nondecreasing. Then τ' is as in (C). If, on the other hand $2^{i_{p+1}} + 2^{i_{p+2}} > 2^{i_p}$, then $2^{i_{p+1}} + 2^{i_{p+2}}$ is not a power of 2. Then τ' is as in (D).

Second, we can have q = p + 1 (still with j = p + 1). Then the entry of σ following 2^{i_p} is $b' = 2^{i_{p+1}} + b$, which has to be at least 2^{i_p} . If b' is not a power of two, then τ' is as in (E), and otherwise, we get (F).

Third, we can have q = p. If $q \le k - 2$, then the *p*th entry of σ is followed by $b' := b + a_{q+2}$, which has to be at least 2^{i_p} . If b' is not a power of two, then τ' is as in (G), and otherwise, we get (H).

There is still one remaining case for j = p + 1, namely, when p = k - 1; then ∂_j just deletes the last coordinate and σ is fully dyadic. Then σ is a source precisely when p = 1 or $i_{p-1} < i_p$, and we have τ' as in (I).

It remains to consider the case of $\tau = [2^{i_1} | \cdots | 2^{i_k}]$ fully dyadic; thus, q = k. First we assume that τ has a peak $p \le k - 1$. Then most of the analysis as above applies.

For j = 0, we find that $\partial_0 \tau$ is a source iff p = 1 and $i_2 \le i_3 \le \cdots \le i_{k-1} < i_k$, and then we have τ' as in (dA).

For $1 \le j \le p - 1$, arguing as in the not fully dyadic case above, for $i_j < i_{j+1}$ we get τ' by swapping 2^{i_j} and $2^{i_{j+1}}$ as in (dB). The case j = p again brings no τ' .

For j = p + 1, we have essentially the first of the three cases of the analogous analysis for the not fully dyadic case $(q = k \ge p + 2)$. For $i_p - 1 = i_{p+1} = i_{p+2} < i_{p+3} \le \cdots \le i_{k-1} < i_k$, we obtain (dC), and for $i_{p+2} \ge i_p > i_{p+1}$ we get (dD) (a swap).

Finally, we may have τ without a peak, which means that $\tau = [2^{i_1}|\cdots|2^{i_{k-2}}|2^{i_{k-1}}|]$, $i_1 \leq \cdots \leq i_{k-1}$ (see case (b) of the definition of V_{bch}). Here ∂_0 and ∂_{k-1} bring no τ' (since $V_{bch}(\partial_0 \tau) = V_{bch}(\partial_{k-1}\tau) = \tau$). For $1 \leq k \leq k-2$ and $i_j < i_{j+1}$, we get a τ' by swapping 2^{i_j} and $2^{i_{j+1}}$ as in (dB). For j = k, ∂_k drops the last component, and if $i_{k-2} < i_{k-1}$, we get a τ' by splitting the last component as in (dI).

Acyclicity Given Lemma 4.3, admissibility of V_{bch} can be proved quickly. Here we will check only acyclicity of the $V\partial$ -graph, since the non-existence of infinite paths will be a side-product of the proof of polynomial boundedness below.

Lemma 4.4 The V∂-graph contains no directed cycle.

Proof If $\tau' = V_{bch}(\partial_j \tau)$ is obtained from τ as in Lemma 4.3, then for τ not fully dyadic, one of the following can happen:

- q(τ') > q(τ), i.e., the length of the dyadic part increases. This is always the case in (F), (G), (H), and (I), and it may also happen in (A) and (C).
- 2. $q(\tau') = q(\tau)$ and the breakpoint value decreases. This happens in (A) and (C) (unless q drops) and also in (E). The latter is not entirely obvious, since we need to check that $\operatorname{ltrim}(2^{i_{p+1}} + b) < b$, but this holds since $\operatorname{ltrim}(2^{i_{p+1}} + b) \le 2^{i_{p+1}} + b 2^{i_p}$, and $2^{i_p} > 2^{i_{p+1}}$.
- 3. $q(\tau') = q(\tau)$, the breakpoint value stays the same, and the dyadic part becomes lexicographically larger. This happens in (B) and (D), since the swaps move a larger component forward.



If τ is fully dyadic, then so is τ' , and either the sum of components of τ' is smaller than that of τ (cases (dA) and (dI)), or the sums of components are equal and τ' is lexicographically larger than τ (cases (dB), (dC), and (dD)).

This implies that there can be no directed cycle.

We remark that an alternative proof of Lemma 4.4 can go along the following lines: If $\tau = [a_1|\cdots|a_k]$ is not fully dyadic, then it can be shown that *either* $ones(\tau') < ones(\tau)$, where $ones(\tau)$ is the total number of 1's in a_1, \ldots, a_k written in binary, *or* $ones(\tau') = ones(\tau)$ and the sequence (i_1, \ldots, i_p) is lexicographically (strictly) larger than $(i'_1, \ldots, i'_{p'})$, where $2^{i_1}|\cdots|2^{i_p}$ is the dyadic nondecreasing part of τ , and similarly for $2^{i'_1}|\cdots|2^{i'_{p'}}$ and τ' .

Polynomial Boundedness Condition (PBV1), polynomial computability of the vector field, is clearly satisfied for V_{bch} , and so we need to check (PVB2); i.e., we need a polynomial bound on the total encoding size of all simplices reachable from a given simplex σ . Obviously, we can focus only on target simplices: it suffices to provide, for every target simplex $\tilde{\tau}$, a polynomial bound on $\sum_{\tau \in t-\text{reach}(\tilde{\tau})} \text{size}(\tau)$ in terms of size($\tilde{\tau}$).

Moreover, it is easy to see that neither the application of V_{bch} nor the face operators ∂_i can increase the sum of the components of the simplex. Thus, $size(\tau) \leq size(\tilde{\tau})$ for every $\tau \in t\text{-reach}(\tilde{\tau})$, and it is enough to bound the *number* of simplices in t-reach($\tilde{\tau}$).

Thus, let us fix a target simplex $\tilde{\tau}$ and set $n := \text{size}(\tilde{\tau})$. Our goal is a polynomial bound, in terms of n, on $|\text{t-reach}(\tilde{\tau})|$.

First we observe that fully dyadic simplices are easily accounted for. Indeed, a fully dyadic simplex $[2^{i_1}|\cdots|2^{i_k}] \in \text{t-reach}(\tilde{\tau})$ is specified by $i_1, \ldots, i_k \in \{0, 1, \ldots, n-1\}$, and so there are at most n^k such simplices.

So we consider only the $\tau \in \text{t-reach}(\tilde{\tau})$ that are not fully dyadic. Let us write $\tilde{\tau} = [\tilde{a}_1|\cdots|\tilde{a}_k]$ and $\tau = [2^{i_1}|\cdots|2^{i_q}|b|a_{q+2}|\cdots a_k]$, where $q = q(\tau)$ is the length of the dyadic part and b is the breakpoint value.

We would like to show that with $\tilde{\tau}$ fixed, there are only polynomially many possibilities for τ . First, as was noted above, the number of choices for the dyadic part of τ is polynomially bounded.

Second, it turns out that all of the right part of τ is inherited from $\tilde{\tau}$, i.e., $a_i = \tilde{a}_i$ for all $i \ge q + 2$. This "stability of the right part" is not hard to prove inductively using Lemma 4.3, and it will be the first part of the key lemma below.

Thus, the last thing to do is showing that there are only polynomially many possibilities for the breakpoint value *b* of τ , and this is the most tricky part of the proof. We will distinguish two cases: if $b = \tilde{a}_{q+1}$, i.e., *b* is "inherited" from $\tilde{\tau}$, then we call τ a *raw* simplex, and otherwise, τ is *processed*.

The following lemma shows that if τ is processed, then its breakpoint value belongs to a certain inductively defined set, which is of polynomial size. In order that the proof goes through, we need to strengthen the inductive hypothesis: namely, we need that for a processed τ , the breakpoint value is smaller than the maximum entry of the dyadic part. This will play a role only in a single case among those in Lemma 4.3, namely (E); while all the other cases are natural and straightforward, (E) seems to work only by a small miracle. **Lemma 4.5** (Key lemma) Let $\tau \in \text{t-reach}(\tilde{\tau})$ be as above. Then $a_i = \tilde{a}_i$ for all $i \ge q + 2$, i.e., the right part of τ coincides with the corresponding segment of $\tilde{\tau}$. Moreover, if τ is processed, then $b < \max(2^{i_1}, \ldots, 2^{i_q})$, and $b \in B_{q+1}$, where the sets B_1, \ldots, B_k are defined inductively as follows:

- B₁ = ltrim^{*}(ã₁), where, for a positive integer a, we define ltrim^{*}(a) = Ø if a is a power of 2, and ltrim^{*}(a) = {ltrim(a)} ∪ ltrim^{*}(ltrim(a)) otherwise.
- $B_{j+1} = \text{Itrim}^*(\{\tilde{a}_{j+1}, \tilde{a}_j + \tilde{a}_{j+1}\} \cup \{2^i + \tilde{a}_{j+1} : 0 \le i \le n-1\} \cup \{b + \tilde{a}_{j+1} : b \in B_j\})$, where we extend $\text{Itrim}^*(\cdot)$ to sets by $\text{Itrim}^*(A) := \bigcup_{a \in A} \text{Itrim}^*(a)$.

Proof It suffices to prove that if τ is as claimed in the lemma, then $\tau' = V_{bch}(\partial_j \tau)$ as in Lemma 4.3 has this form as well (moreover, we may assume that τ' is not fully dyadic). We need to consider the cases (A)–(I) in Lemma 4.3, but we can right away settle (I), where τ' is fully dyadic, as well as (B) and (D), which only permute the dyadic part. This leaves us with cases (A), (C), (E), (F), (G), and (H).

First let τ be raw, with $b = \tilde{a}_{q+1}$. In cases (A) and (C) τ' contains $\operatorname{lpow}(\tilde{a}_{q+2})$ followed by $b' := \operatorname{ltrim}(\tilde{a}_{q+2})$, at the (q + 1)st position. If b' is a power of 2, then τ' is raw, and otherwise, we have $b' \in B_{q+1}$ and $b' < \operatorname{lpow}(\tilde{a}_{q+1})$; the latter is the required entry larger than b' in the dyadic part. Hence τ' is a processed simplex as claimed in the lemma.

In (E) and (G), we have a situation similar to the one just discussed, except that $b' = \tilde{a}_{q+1} + 2^i$ for some i < n in (E), and $b' = \tilde{a}_{q+1} + \tilde{a}_{q+2}$ in (G). Moreover, in (E), b' is at position q + 1, while in (G) it is at position q + 2. Again we find that τ' is a processed simplex of the claimed form. In cases (F) and (H), we either get τ' fully dyadic, or the breakpoint value of τ' is ltrim $(\tilde{a}_{q'+1})$ for some $q' \ge q + 1$, preceded by $|pow(\tilde{a}_{q'+1})$. Then τ' is a processed simplex as in the lemma as well, and the discussion of a raw τ is finished.

Now let τ be processed, with $b \in B_{q+1}$, $b < \max(2^{i_1}, \ldots, 2^{i_q})$. In cases (A) and (C) τ' may be raw, which is fine, or processed with breakpoint value $\operatorname{ltrim}(b)$, which lies in B_{q+1} , since B_{q+1} is closed under $\operatorname{ltrim}(\cdot)$.

Case (E) is, in a sense, the most sophisticated, and it is here where the inductive hypothesis $b < \max(2^{i_1}, \ldots, 2^{i_q})$ is crucial. In the setting of (E), 2^{i_p} is the maximum of the dyadic part of τ , and so $2^{i_p} > b$. Let $b' = b + 2^{i_{p+1}}$, where $2^{i_{p+1}} < 2^{i_p}$; by the conditions in case (E), we have $b' > 2^{i_p}$.

We claim that $\operatorname{ltrim}(b') \in \operatorname{ltrim}^*(b)$ (this will show that $b' \in B_{q+1}$ and thus τ' is as required). To check this, let us write, for brevity, $u = i_p$ and $v = i_{p+1}$, and let $\beta_{u-1}\beta_{u-2}\cdots\beta_0$ be the binary notation for b, i.e., $b = \sum_{i=0}^{u-1}\beta_i 2^i$, $\beta_i \in \{0, 1\}$. Since $2^u - 2^v < b < 2^u$, we have $\beta_{u-1} = \cdots = \beta_v = 1$. Then b' in binary is $1000\cdots0\beta_{v-1}\beta_{v-2}\cdots\beta_0$, and so $\operatorname{ltrim}(b')$ can be obtained from b by iterating $\operatorname{ltrim}(\cdot)$. Thus, $b' \in B_{q+1}$ indeed.

The consideration in cases (F) and (H) is the same as the one for $\tilde{\tau}$ raw.

The last case to consider is (G). Here the dyadic part of τ' is longer than that of τ . By induction, we have $b \in B_{q+1}$, and so $\operatorname{ltrim}(b + \tilde{a}_{q+2}) \in B_{q+2}$ by the definition of B_{q+2} (or it is a power of 2, in which case τ' is raw). As in the previous case, the entry $\operatorname{lpow}(b + \tilde{a}_{q+2})$ supplies the power of 2 greater than $\operatorname{ltrim}(b + \tilde{a}_{q+2})$, as required for the induction. The lemma is proved.

لے Springer

Corollary 4.6 For $\tilde{\tau}$ as in Lemma 4.5, we have $|t\text{-reach}(\tilde{\tau})| = O(n^{2k})$, with the implicit constant depending on k.

Proof For each $\tau \in \text{t-reach}(\tilde{\tau})$, we have at most n^k choices for the dyadic part (which includes fixing q, the length of the dyadic part). A raw τ is already determined by $\tilde{\tau}$ and by the dyadic part, while for τ processed, we also need to specify b.

The definition of B_j gives $|B_1| \le n$ and $|B_{j+1}| \le 3n + n^2 + n|B_j|$, which yields $|B_j| = O(n^j)$, and the corollary follows.

Remark A more careful (and more complicated) analysis should probably give $O(n^k)$ instead of $O(n^{2k})$ in Corollary 4.6. However, as we will now indicate, our vector field is not much better; there can indeed be about n^k reachable simplices in t-reach($\tilde{\tau}$).

To see this, let us take *n* that is an integer multiple of k^2 , i.e., $n = k^2 \ell$, and let us consider a source simplex $\tilde{\sigma} = [\tilde{a}_1|\cdots|\tilde{a}_k]$, where $\tilde{a}_i := (2^{\ell} - 1)2^{(i-1)\ell}$, i = 1, 2, ..., k. Put differently, if we think of the binary encoding of each \tilde{a}_i as consisting of *k* blocks of ℓ bits each (thus, \tilde{a}_i has at most n/k bits and size $(\tilde{\sigma}) \le n$), then \tilde{a}_i has 1's in the *i*th block and 0's elsewhere. It can be shown that each simplex $\sigma = [a_1|\cdots|a_k]$, where a_i has exactly one 1 in the *i*th block and 0's everywhere else, belongs to reach $(\tilde{\sigma})$. Since for each *i*, the position of the single 1 in a_i can be chosen in ℓ ways, we have $|\text{reach}(\tilde{\sigma})| > \ell^k = (n/k^2)^k$.

It would be interesting to see if one could reach a significantly better bound with a different vector field, or if there is perhaps a good lower bound valid for every vector field.

Acknowledgements We would like to thank Martin Čadek, Lukáš Vokřínek, and Uli Wagner for useful discussions and ongoing collaboration. Moreover, we thank Uli Wagner and Martin Čadek for insightful comments on a preliminary version of the manuscript. The research by J. M. and M. K. was supported by the Institute for Theoretical Computer Science (ITI), Charles University, Prague (project 1M0545 of the Czech Ministry of Education) and by the ERC Advanced Grant No. 267165. The research by M. K. was also supported by the project GAUK 49209.

References

- D.J. Anick, The computation of rational homotopy groups is #β-hard, in *Computers in Geometry and Topology*, Proc. Conf. Chicago/III., 1986. Lect. Notes Pure Appl. Math., vol. 114 (1989), pp. 1–56.
- M. Čadek, M. Krčál, J. Matoušek, F. Sergeraert, L. Vokřínek, U. Wagner, Computing all maps into a sphere. Preprint arXiv:1105.6257 (2011). Extended abstract in *Proc. ACM–SIAM Symposium on Discrete Algorithms* (SODA 2012).
- M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, U. Wagner, Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. Preprint arXiv:1211.3093 (2012).
- M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, U. Wagner, Extendability of continuous maps is undecidable. Preprint arXiv:1302.2370 (2013).
- H. Cartan, Algèbres d'Eilenberg–MacLane et homotopie. Exposés 2 à 16, Séminaire Henri Cartan (École Normale Supérieure, Paris, 1956).
- 6. A. Clément, Integral cohomology of finite Postnikov towers. Doctoral Thesis, Univ. de Lausanne, 2002.
- 7. E.B. Curtis, Simplicial homotopy theory, Adv. Math. 6, 107-209 (1971).
- 8. S. Eilenberg, S. Mac Lane, On the groups of *H*(*Π*, *n*). I, Ann. Math. 58, 55–106 (1953).
- 9. R. Forman, Morse theory for cell complexes, Adv. Math. 134(1), 90-145 (1998).
- R. Forman, A user's guide to discrete Morse theory, Séminaire Lotharingien de Combinatoire 48 (2002). Article B48c.

- P. Franek, S. Ratschan, P. Zgliczynski, Satisfiability of systems of equations of real analytic functions is quasi-decidable, in *Proc. 36th International Symposium on Mathematical Foundations of Computer Science (MFCS)*. LNCS, vol. 6907 (Springer, Berlin, 2011), pp. 315–326.
- G. Friedman, An elementary illustrated introduction to simplicial sets, *Rocky Mt. J. Math.* 42(2), 353–423 (2012).
- 13. P.G. Goerss, J.F. Jardine, Simplicial Homotopy Theory (Birkhäuser, Basel, 1999).
- J. Heras, V. Pascual, J. Rubio, F. Sergeraert, fKenzo: a user interface for computations in algebraic topology, J. Symb. Comput. 46(6), 685–698 (2011).
- 15. J.P. May, *Simplicial Objects in Algebraic Topology* (Chicago University Press, Chicago, 1992). Reprint of the 1967 original; the page numbers do not quite agree with the 1967 edition.
- A. Romero, J. Rubio, Computing the homology of groups: the geometric way. Preprint arXiv: 1107.3396, http://arxiv.org/pdf/1107.3396v1 (2011).
- 17. A. Romero, F. Sergeraert, Discrete vector fields and fundamental algebraic topology. Preprint (2011) an updated version at http://www-fourier.ujf-grenoble.fr/~sergerar/Papers/. arXiv:1005.5685
- A. Romero, J. Rubio, F. Sergeraert, Computing spectral sequences, J. Symb. Comput. 41(10), 1059– 1079 (2006).
- 19. J. Rubio, F. Sergeraert, Constructive algebraic topology, Bull. Sci. Math. 126(5), 389-412 (2002).
- J. Rubio, F. Sergeraert, Constructive homological algebra and applications. Preprint arXiv:1208.3816 (2012). Written in 2006 for a MAP Summer School at the University of Genova.
- 21. F. Sergeraert, The computability problem in algebraic topology, Adv. Math. 104(1), 1-29 (1994).
- F. Sergeraert, Introduction to combinatorial homotopy theory. Available at http://www-fourier.ujfgrenoble.fr/~sergerar/Papers/. (2008).
- J.-P. Serre, Cohomologie modulo 2 des complexes d'Eilenberg-MacLane, Comment. Math. Helv. 27, 198–232 (1953).