FINDING BOUNDED SIMPLICIAL SETS WITH FINITE HOMOLOGY

SPUR FINAL PAPER, SUMMER 2021

PRESTON CRANFORD AND PETER ROWLEY

Mentor: Robert Burklund
Project suggested by Robert Burklund
August 4, 2021

ABSTRACT. The central problem in computational algebraic topology is the computation of the homotopy groups of a given space, represented as a simplicial set. Algorithms have been found which achieve this, but the running times depend on the size of the input simplicial set. In order to reduce this dependence on the simplicial set chosen, we describe in this paper a procedure which, given a prime p and a finite, simply-connected simplicial set with finite integral homology, finds a p-locally equivalent simplicial set with size upper bounded by a function of dimension and homology. Using this in conjunction with the above algorithm, the p-local homology can be calculated such that the running time dependence on the size of the initial simplicial set is contained in a separate preprocessing step.

1. INTRODUCTION

An important area of algebraic topology is the study of maps into and out of a given space X. In the process of understanding such maps, one of the most fundamental questions that can be asked surrounds the structure of all maps from spheres into X, or the study of the homotopy groups of X. However, despite being so fundamental, and despite their ties to homology through results such as the Hurewicz theorem, the homotopy groups of a space have proven to be extremely difficult to calculate. Because of this, the calculation of homotopy groups has become one of the central problems of computational algebraic topology.

Despite the difficulty, there do exist algorithms which compute the homotopy groups of a simply connected space which represents a finite homotopy type. When algorithmically computing homotopy groups, a topological space is represented in the computer as a simplicial set. There are various algorithms for doing this, but in large part they are all variations of an algorithm due to Brown [Bro57]. His algorithm calculates the homotopy groups of a simply connected space by inductively computing its Postnikov tower.

However, his algorithm is by no means practical in terms of its running time. In fact, the problem of computing $\pi_k(X)$ for a space X has been shown to be W[1]-hard [Mat14] when k is regarded as an input parameter, which implies that a fast general algorithm almost certainly does not exist when varying k. However, by regarding k as fixed, Čadek et. al. [ČKM$^+$14] found an algorithm which computes $\pi_k(X)$ in time polynomial in the size of the input simplicial set X. Their algorithm uses Brown's as a backbone, but they associate additional data with each intermediate simplicial set, called "polynomial time homology," which allows each of the steps to be performed in polynomial time.

Though this result is impressive, their polynomial time guarantee is slightly unsatisfactory due to its dependence on the size of the input simplicial set. In particular, this quantity has

no relation to the space represented by the simplicial set, and in fact could be arbitrarily large for even the simplest underlying spaces. As such, it would be desirable to find an algorithm which, given any simplicial set, outputs another homotopy equivalent simplicial set which has size that is bounded in terms of quantities related to the space represented by the simplicial set.

Our main result achieves this goal p-locally for any prime p with the restriction that the input space X must be finite-dimensional with finite integral homology groups. The bound is given in terms of various quantities that are derived from the dimension and homology groups of the p-localization of X. To start, we use the dimension d of X. We will also use the following quantities.

Definition 1.1. Given X as above and a prime p, we denote
$$h_p := \sum_{k=1}^{d} \operatorname{rank}_p H_k(X; \mathbb{Z}_{(p)}).$$

This quantity, the total rank of the $\mathbb{Z}_{(p)}$-homology of X, is used as a bound on the overall size of the p-local homology. In addition, let
$$h := \max_p h_p.$$

Next, recall that an abelian group A has p-torsion exponent b if b is the smallest integer such that p^b annihilates all p-torsion elements of A. Let m_p be the sum of the p-torsion exponents of $H_k(X)$ over the range $0 \le k \le d$. Intuitively, this measures how much p-torsion is in the homology of X. Then, analogously to the definition of h, let
$$m := \max_p m_p.$$

Finally, let
$$N := \prod_{p:h_p \neq 0} p.$$

We use this as a measure of the diversity and size of the primes in the homology of X.

Remark. Though all of these quantities are functions of X, this dependence is left implicit, as X is typically clear from the context.

Given these, the main result is as follows.

Theorem 1.2 (Theorem 5.4)**.** *Let X be a finite, simply-connected simplicial set of dimension d with finite integral homology, and let p be a prime. Then there exists an algorithm giving a simplicial set which is p-locally homotopy equivalent to X with at most*
$$\exp\left(mh\log(N)\exp\left(\log(2h)d + O\left(\log(d)^3\right)\right)\right)$$
non-degenerate simplices.

Remark. Though p-local equivalence is not as strong as full homotopy equivalence, we note that if X has finite homology groups, it is is still sufficient to find the homotopy groups of the space. The procedure simply involves localizing at one prime at a time and using the above result for each. Since X has finite dimension and finite homology groups, this only involves localizing at finitely many primes.

Our main theorem implies the following result for a simply-connected, finite homotopy type with finite integral homology.

Corollary 1.3 (Corollary 5.5). *Let X be a simply-connected, finite homotopy type with finite integral homology whose highest non-trivial homology is in degree d. Then there exists a finite simplicial set S representing a space p-locally equivalent to X with at most*

$$\exp\left(mh\log(N)\exp\left(\log(2h)(d+1) + O\left(\log(d)^3\right)\right)\right).$$

non-degenerate simplices.

Our algorithm largely depends on part of Brown's algorithm; we follow the presentation given in [ČKM$^+$14]. In particular, we construct the $(d+1)^{\text{st}}$ stage of a Postnikov tower of the input space X, where d is the dimension of X. With the finite homology condition as stated, this gives us a simplicial set which, after some slight modifications, is p-locally equivalent to X.

Paper Outline. In section 2, we establish notation and recall definitions. In particular, we review the definitions of Eilenberg-Maclane spaces, Postnikov systems, and twisted products, and present the standard simplicial set representative for an Eilenberg-Maclane space.

Next, in section 3 we bound the size of the homotopy groups of a simply-connected, finite-dimensional space X with finite integral homology; the bound is in terms of the size of the homology groups. This is accomplished by combining two results from the literature: one that bounds the rank of the homotopy groups and another that bounds the order of elements of the homotopy groups.

In section 4 we use the Postnikov system algorithm of Brown to create a simplicial set that will be shown to satisfy the requirements of Theorem 1.2. This is accomplished by finding a skeleton of a Postnikov stage of X via Brown's algorithm, then removing simplices so that the Postnikov map induces a homotopy equivalence.

In section 5 we bound the size of the resulting simplicial set by bounding the size of a Postnikov stage of X. This is done by analyzing the size of the standard simplicial set model of an Eilenberg-Maclane space and then expressing sizes of Postnikov stages in terms of these spaces.

Finally, in section 6 we consider potential ways to strengthen our main result in the future, such as improving the bound in Theorem 1.2 and extending it to include spaces with infinite homology groups. We also discuss potential directions that could be considered in pursuit of either of these goals.

Conventions. Let S be any simplicial set. Its set of k-simplices is denoted S_k, the face operators $S_k \to X_{S-1}$ are denoted $d_0, \ldots, d_k$, and the degeneracy operators $S_k \to S_{k+1}$ are denoted $s_0, \ldots, s_k$. Given any simplicial set Y, we denote by $sk_k(Y)$ the k-skeleton of Y. We use $C_*(Y)$ to denote the normalized chain complex of Y, which is generated by the set of nondegenerate simplices in each degree. The chain differential map $C_n \to C_{n-1}$ will be denoted by ∂_n, or just ∂ if n is clear from the context.

In this paper X will denote a fixed finite, simply-connected simplicial set of dimension d with finite integral homology unless otherwise indicated. We will denote by d the dimension of X. Though this overlaps with the notation for face maps, it will always be clear from context which one is meant.

We take all exponent bases and logarithm bases to be e unless otherwise indicated.

2. PRELIMINARIES

In this section we establish notation and recall definitions. In particular, we review the definitions of Eilenberg-Maclane spaces, Postnikov systems, and twisted products, and present the standard simplicial set representative of an Eilenberg-Maclane space.

2.1. **Eilenberg-Maclane Spaces and Postnikov Systems.** We begin by defining Eilenberg-Maclane spaces and presenting the standard simplicial models for them. We will then use these concepts to define a Postnikov system of a space.

Definition 2.1. Let π be an Abelian group and let $n \geq 1$. An *Eilenberg-Maclane Space* $K(\pi, n)$ is a space such that $\pi_n(K(\pi, n)) = \pi$ and $\pi_k(K(\pi, n)) = 0$ for all $k > 0, k \neq n$.

It is a standard fact of algebraic topology that there exists a $K(\pi, n)$ and it is unique up to homotopy equivalence among spaces homotopy equivalent to CW complexes.

We now present the standard simplicial model for $K(\pi, n)$. From here on in this paper, the notation $K(\pi, n)$ will refer to this simplicial set, defined as follows.

Definition 2.2. Let Δ^ℓ denote the simplicial set representing the standard ℓ-simplex. Then we define $K(\pi, n)_\ell$ as the set of cocycles $Z^n(\Delta^\ell; \pi)$. The face operator d_i is defined in the following way. Each simplex $\sigma \in K(\pi, n)$ is represented by a cocyle in $Z^n(\Delta^l, \pi)$, so let $d_i\sigma$ be the restriction of σ to the i^{th} $(l - 1)$-face of Δ^l. The degeneracy operators are defined similarly.

This definition rise to a similar, yet slightly larger, simplicial set $E(\pi, n)$. This is defined exactly as above, except that $E(\pi, n)_\ell$ is the set of cochains $C^n(\Delta^\ell; \pi)$, rather than the cocycles. Since $E(\pi, n)$ is defined as a set of cochains and $K(\pi, n)$ is defined as a set of cocycles, there is a simplicial map $\delta_* : E(\pi, n) \to K(\pi, n + 1)$ which is induced by the coboundary map $\delta : C^n(\Delta^\ell; \pi) \to Z^{n+1}(\Delta^\ell; \pi)$. We can now define a Postnikov system of a space.

Definition 2.3. Let X be a simplicial set. A *Postnikov system* of X is a collection of simplicial sets $\{P_k\}_{k=0}^{\infty}$ and maps $\varphi_k : X \to P_k$ and $p_k : P_k \to P_{k-1}$ which fit in a commutative diagram as in Figure 1, such that the p_k are fibrations with fiber $K(\pi_k(X), k)$. We call P_k the k^{th} *Postnikov stage*, and call p_k the k^{th} *projection map*.

The maps φ_k are such that the induced maps $\pi_n(\varphi_k) : \pi_n(X) \to \pi_n(P_k)$ are isomorphisms for $n \leq k$. In addition, for each $k \geq 1$ there is a map $\mathbf{k}_{k-1} : P_{k-1} \to K(\pi_k(X), k + 1)$ called the $(k - 1)^{\text{st}}$ *Postnikov class*; based on this, P_k is obtained via the pullback in Figure 2.

2.2. **Twisted Products.** A useful construction in the analysis of Postnikov systems is the twisted Cartesian product. According to [ČKM+14], this construction is intended to mimic a fiber bundle in simplicial sets; they do not give details, but see their paper for slightly more of the intuition.
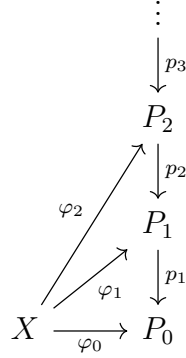
$$\vdots$$

$$\downarrow p_3$$

$$P_2$$

$$\varphi_2 \qquad \nearrow \qquad \downarrow p_2$$

$$P_1$$

$$\nearrow$$

$$\varphi_1 \qquad \downarrow p_1$$

$$X \xrightarrow{\quad \varphi_0 \quad} P_0$$

FIGURE 1. Postnikov Tower

$$
\begin{array}{ccc}
P_k & \longrightarrow & E(\pi_k(X), k) \\
\downarrow & & \downarrow {\scriptstyle \delta_*} \\
P_{k-1} & \xrightarrow[\mathbf{k}_{k-1}]{} & K(\pi_k(X), k+1)
\end{array}
$$

FIGURE 2. Pullback square for P_k

Definition 2.4. Let F, B be a simplicial sets, and let $F \times B$ denote their Cartesian product. Let G be a simplicial group acting on F on the right. A *twisting operator* τ is a collection of maps $(\tau_k)_{k=1}^{\infty}$ with $\tau_k : B_k \to G_{k-1}$ such that the following conditions are satisfied:

- $d_0(\tau(\beta)) = \tau(d_1\beta)\tau(d_0\beta)^{-1}$
- $d_i(\tau(\beta)) = \tau(d_{i+1}\beta)$ for $i \geq 1$
- $s_i(\tau(\beta)) = \tau(s_{i+1}\beta)$ for all i
- $\tau(s_0\beta) = e_k$ for all $\beta \in B_k$, where e_k is the identity of G_k.

Definition 2.5. Let F, B, and G be as above, and let τ be a twisting operator. Then the *twisted Cartesian product* (or *twisted product*) $F \times_\tau B$ is a simplicial set with $(F \times_\tau B)_k = F_k \times B_k$; its face and degeneracy maps coincide with those of the Cartesian product, with the exception of d_0, which satisfies

$$d_0(f, b) = (d_0(f)\tau(b), d_0(b)).$$

The map $\delta_* : E(\pi, k) \to K(\pi, k+1)$ described above is a fibration with fiber $K(\pi, k+1)$. In keeping with the intuition that twisted products represent fiber bundles, therefore, it turns out that $E(\pi, k)$ can be described as a twisted product $K(\pi, k) \times_\tau K(\pi, k+1)$. The twisting operator τ is described as follows. Let $z \in Z^{k+1}(\Delta^l; \pi)$, and let a k-face of Δ^l be denoted by a $(k+1)$-tuple $(i_0, \ldots, i_k)$ with $0 \leq i_0 < \cdots < i_k \leq \ell$. Then we define

$$(\tau(z))(i_0, \ldots, i_k) := z(0, i_0 + 1, i_1 + 1, \ldots, i_k + 1) - z(1, i_0 + 1, i_1 + 1, \ldots, i_k + 1)$$

It is a consequence of this fact that the pullback shown in Figure 2 gives rise to an isomorphism as follows, where $\tau^* := \tau \circ \mathbf{k}_{k-1}$:

$$P_k \cong P_{k-1} \times_{\tau^*} K(\pi_k(X), k+1).$$

We refer the reader to [May67, Proposition 18.7] for more details.

3. Bounding Homotopy Groups

In this section we bound the size of the homotopy groups $\pi_n(X)$ of a space X in terms of it's homology invariants h, n, M. As in the main theorem, we will only consider the case that X has finitely many nonzero homology groups, each of which is of finite order. It is well known by [Ser53] that this implies that the homotopy groups are finite as well. In order to bound $|\pi_n(X)|$ based on homology information in this case, we need two separate bounds: a bound on the rank of $\pi_n(X)$, and a bound on the orders of the individual elements.

Let $X_{(p)}$ denote the p-localization of X. Since $\pi_n(X)$ is finite for all n, we have $X \cong \prod_p X_{(p)}$, so $\pi_n(X) \cong \prod_p \pi_n(X_{(p)})$; as such, we will focus on bounding the homotopy groups of $X_{(p)}$ for each p. Note the following bound for the rank of $\pi_n(X_{(p)})$:

Proposition 3.1 ([BS])**.** *For X simply-connected and for each prime p we have that*

$$rank(\pi_n(X_{(p)})) \le f(n)r^n,$$

where $f(n) \in \exp\left(O\left(\log(n)^3\right)\right)$ and $r \le h_p$.

We also have the following bound on the torsion exponent of $\pi_n(X_{(p)})$.

Proposition 3.2 ([Bar60], pg. 125)**.** *Let X be simply connected. Then for each prime p we have*

$$p^{cn}\pi_n(X_{(p)}) = 0,$$

where $c \le 2m_p$.

Putting these two bounds together, we get the following:

Proposition 3.3. *Let X be simply connected, and let p be any prime. Then*

$$|\pi_n(X_{(p)})| \le \exp(m_p \log(p) \exp(n \log(h_p) + O\left(\log(n)^3\right)).$$

Proof. The above bounds are combined as follows.

$$
\begin{aligned}
|\pi_n(X_{(p)})| &\le \left(p^{cn}\right)^{\mathrm{rank}(\pi_n(X_{(p)}))} \\
&\le p^{cnr^n f(n)} \\
&\le p^{2m_p nh_p^n f(n)} \\
&\le \exp\left(\log(p)2m_p nh_p^n \exp\left(O\left(\log(n)^3\right)\right)\right) \\
&\le \exp\left(m_p \log(p) \exp\left(\log(2nh_p^n) + O\left(\log(n)^3\right)\right)\right) \\
&\le \exp\left(m_p \log(p) \exp\left(\log(2n) + n\log(h_p) + O\left(\log(n)^3\right)\right)\right) \\
&\le \exp(m_p \log(p) \exp\left(n\log(h_p) + O\left(\log(n)^3\right)\right).
\end{aligned}
$$

$\square$

We can now easily combine the above bound at each value of p to get the following overall bound.

Proposition 3.4. *Let X be simply connected, and let p be any prime. Then*

$$|\pi_n(X)| \le \exp\left(m \log(N) \exp(n \log(h) + O\left(\log(n)^3\right)\right)$$

Proof. As noted above, we know that $\pi_n(X) \cong \prod_p \pi_n(X_{(p)})$. Therefore, to bound $\pi_n(X)$ multiply the bounds given by Proposition 3.3 at each prime p which shows up in the homology of X. We first replace each m_p with m, because $m_p \leq m$ for all p by definition. Translating this product into a sum in the exponent replaces the $\log(p)$ with $\sum_p \log(p)$, and then combining the logs allows us to replace this with $\log(N)$. Recall also that by definition, $h_p \leq h$ for all p, so we can replace each $\log h_p$ with $\log h$, giving us the above expression. $\square$

4. Simplicial Set Algorithm

In this section we describe the algorithm used to construct the simplicial set required for the main theorem. This procedure relies heavily on the algorithm of Brown [Bro57] for calculating the homotopy groups of an arbitrary simply-connected simplicial set, as presented in [ČKM+14]. We will describe this algorithm in minimal detail.

4.1. **Postnikov Tower Algorithm.** We begin by summarizing at a high level the algorithm of Brown [Bro57] for constructing a Postnikov system for the input space X. As always, X is a finite, simply-connected simplicial set; their algorithm does not require the finiteness assumption on the homology of X.

The algorithm proceeds inductively, starting with the trivial zeroth stage P_0, which is a point, as well as φ_0 the constant map. Given P_{k-1} and φ_{k-1}, the authors simultaneously find the group $\pi_k(X)$ and a simplicial set representative of the Postnikov class $\mathbf{k}_{k-1} : P_{k-1} \to K(\pi_k(X), k+1)$ using an algebraic construction. The k^{th} Postnikov stage is then constructed as the simplicial set pullback shown in Figure 2.

The construction of φ_k is not important to the rest of the algorithm, so we refer the reader to [ČKM+14].

Notice, in particular, that P_k is constructed as a pullback in a way that automatically gives us the following relation:

$$P_k \subset P_{k-1} \times E(\pi_k(X), k)$$

Inductively, this implies that it is true that

(1) $$P_k \subset E(\pi_2(X), 2) \times \cdots \times E(\pi_k(X), k).$$

Note that $E(\pi_1(X), 1)$ is omitted because X is simply connected. Equation (1) is the most important takeaway from this algorithm as we move forward.

4.2. **Constructing an Equivalent Simplicial Set.** We now describe how the above algorithm is used in order to find a size-bounded simplicial set which is p-locally homotopy equivalent to X, as in the main theorem. In the following section we analyze the size of the simplicial set that results.

Our procedure is motivated by the following property of p-localization.

Theorem 4.1. [Bou75] *Let X be as in the main theorem, let Y be a finite simplicial set , and let $f : X \to Y$ be a map which induces isomorphisms on $\mathbb{Z}_{(p)}$-homology in every degree. Then f induces a homotopy equivalence on p-localizations.*

The construction of a map f and a space Y as above uses the Postnikov tower algorithm as follows. Let d be the dimension of X. Starting with X, we run the algorithm to get P_{d+1}, the $(d+1)^{\text{st}}$ stage of the Postnikov tower. The algorithm also gives us a map $\varphi_{d+1} : X \to P_{d+1}$ which induces isomorphisms in π_n for $n \leq d+1$. This gives us the following.

Lemma 4.2. *The map $\varphi_{d+1} : X \to P_{d+1}$ induces an isomorphism in homology through degree $d+1$.*

Proof. Since φ_{d+1} induces homotopy isomorphisms in degrees $\leq d+1$, it follows by the homotopy long exact sequence that after replacing P_{d+1} with the mapping cylinder of φ_{d+1}, $\pi_n(P_{d+1}, X) = 0$ for $n \leq d+1$. Therefore, the relative Hurewicz theorem tells us that $H_n(P_{d+1}, X) = 0$ for $n \leq d+1$, which implies that φ_{n+1} induces integer homology isomorphisms in degrees $\leq d+1$ as well. $\square$

The construction of P_{d+1} also gives us a crucial finiteness condition.

Lemma 4.3. *The simplicial set P_{d+1} is levelwise finite.*

Proof. Recall from (1) that

$$P_{d+1} \subset E(\pi_2(X), 2) \times \cdots \times E(\pi_{d+1}(X), d+1).$$

Since the n-simplices of $E(\pi_k(X), k)$ are the k-cochains $C_k(\Delta^n, \pi_k(X))$ and $\pi_k(X)$ is a finite group by assumption, it follows that there are finitely many n-simplices in $E(\pi_k(X), k)$ for each n and k. Therefore, P_{d+1} is a levelwise-finite simplicial set. $\square$

Based on these results, we can get a finite simplicial set which satisfies almost all of our requirements by considering $sk_{d+2}(P_{d+1})$, the $(d+2)$-skeleton of P_{d+1}. It is clear that this preserves the homology through degree $d+1$, so Lemma 4.2 still holds. We will continue to use the notation φ_{d+1} for the map $X \to sk_{d+2}(P_{d+1})$, because the map itself has not changed. Note, in particular, that φ_{d+1} induces homology isomorphisms in every degree other than $d+2$.

What remains, therefore, is to remove some of the $(d+2)$-simplices of $sk_{d+2}(P_{d+1})$ to create a sub-simplicial-set Y with $H_{d+1}(Y; \mathbb{Z}_{(p)}) = H_{d+2}(Y; \mathbb{Z}_{(p)}) = 0$. The map $X \to Y$ induced by φ_{d+1} will therefore induce $\mathbb{Z}_{(p)}$-homology isomorphisms in every degree, so this is sufficient. To do this, we wish to remove enough $(d+2)$-simplices from $sk_{d+2}(P_{d+1})$ such that there are no more $(d+2)$-cycles, but so that all $(d+1)$-boundaries remain boundaries. This is done in the following lemma.

Lemma 4.4. *We can find a subset $\{t_1, \ldots, t_\ell\}$ of the $(d+2)$-simplices of $sk_{d+2}(P_{d+1})$ such that, when the t_i are regarded as elements of $C_{d+2}(sk_{d+2}(P_{d+1}); \mathbb{Z}_{(p)})$, the set $\{\partial t_1, \ldots, \partial t_\ell\}$ is a basis of $\mathrm{im}\partial_{d+2} \subset C_{d+1}(sk_{d+2}(P_{d+1}); \mathbb{Z}_{(p)})$.*

Proof. Note that since $H_{d+1}(sk_{d+2}(P_{d+1}); \mathbb{Z}_{(p)}) = 0$, it follows that the map ∂_{d+2} surjects onto $Z_{d+1}(sk_{d+2}(P_{d+1}); \mathbb{Z}_{(p)})$. Let $\mathbf{B}_Z$ be a basis of of $Z_{d+1}(sk_{d+2}(P_{d+1}); \mathbb{Z}_{(p)})$. We construct the $\mathbb{Z}_{(p)}$-matrix representation M of the chain differential ∂_{d+2}, using the standard basis of C_{d+2} and the basis $\mathbf{B}_Z$ of the target space. It is sufficient to find a subset of the columns of M that form a basis of the $\mathbb{Z}_{(p)}$-span of the columns.

We will find this subset by induction on the size of $\mathbf{B}_Z$, which is the number of rows of M. For the base case of size 1, note that in a $1 \times n$ matrix representing a surjective map, at least one of the entries must not be divisible by p. Then, since we have $\mathbb{Z}_{(p)}$-coefficients, that entry is invertible, so it generates the span of all of the entries of the matrix. Thus, our desired simplex t_1 is the simplex corresponding to the column of this entry.

Now assume that M has ℓ rows. Since it represents a surjective map, at least one of the columns is not divisible by p. It follows that, using $\mathbb{Z}_{(p)}$ coefficients, we can scale such a column to be not divisible by any prime.

Let t_1 be a $(d+2)$-simplex corresponding to a column which is not divisible by p. It is clear that, since t_1's column can be scaled to be indivisible by any prime, $\operatorname{im}(\partial)/\partial t_1$ has dimension $\ell - 1$. In addition, the induced map $\partial^1 : C_{d+2}(sk_{d+2}(P_{d+1}); \mathbb{Z}_{(p)})/t_1 \to Z_{d+1}(sk_{d+2}(P_{d+1}); \mathbb{Z}_{(p)})/\partial t_1$ is clearly still surjective. Thus, we can find $t_2, \ldots, t_\ell$ using the inductive hypothesis, and we are done. $\qquad\square$

We can now put together the above results to prove the following lemma, which provides the simplicial set required by the main theorem; the remainder of the proof will consist of analyzing the size of the Postnikov stage, which takes place in the next section.

Lemma 4.5. *Let X be a finite, simply-connected simplicial set with finite integral homology, and let p be a prime. Then there exists an algorithm giving a finite simplicial set $Y \subset sk_{d+2}(P_{d+1})$ where Y is p-locally homotopy equivalent to X.*

Proof. As above, we can use Brown's algorithm [Bro57] to construct $sk_{d+2}(P_{d+1})$. There exists a finite subset $T \subset (P_{d+1})_{d+2}$ of non-degenerate simplices such that $\{dt | t \in T\}$ is a basis of $Z_{d+1}(P_{d+1}; \mathbb{Z}_{(p)})$. Let Y be a simplicial set such that $Y_n = (P_{d+1})_n$ for $n \le d+1$ and $Y_{d+2} = T$. Then the map $\varphi_{d+1} : X \to Y$ is easily seen to induce $\mathbb{Z}_{(p)}$-homology isomorphisms in all degrees other than $d+1$ and $d+2$. In addition, since $B_{d+1}(Y; \mathbb{Z}_{(p)}) = B_{d+1}(P_{d+1}; \mathbb{Z}_{(p)}) = Z_{d+1}(P_{d+1}; \mathbb{Z}_{(p)})$, we have $H_{d+1}(Y; \mathbb{Z}_{(p)}) = 0$, so φ_{d+1} induces an isomorphism in degree $d+1$. Finally, T is linearly independent, so we have that $Z_{d+2}(Y; \mathbb{Z}_{(p)}) = H_{d+2}(Y; \mathbb{Z}_{(p)}) = 0$, meaning φ_{d+1} induces an isomorphism in degree $d+2$ as well, which completes the proof by Lemma 4.2. $\qquad\square$

5. Bounding Postnikov Stage Size

In this section we bound the the cardinality of the simplicial set formed in the algorithm for the main theorem, as described in Section 4.

Since the Postnikov stages P_k are described in terms of the simplicial set models of the Eilenberg-Maclane spaces $K(\pi, k)$, we first seek to understand $|K(\pi, k)_n|$ for a given finite group π.

Proposition 5.1. *In the standard simplicial set model of $K(\pi, k)$ we have*

$$|K(\pi, k)_n| = |\pi|^{\binom{n}{k}}.$$

Proof. Recall that elements of $K(\pi, k)_n = Z^k(\Delta^n; \pi)$ are labellings of the k-faces of Δ^n with elements of π in such a way that satisfies the cocycle condition. Representing Δ^n as the set $[n] = \{0, 1, \ldots, n\}$, this is equivalent to labelling all of the $(k+1)$-element subsets of this set with elements of π. For convenience we will always label subsets with the elements in increasing order: $\{a_0, a_1, \ldots, a_k\}$ implicitly has $a_0 < a_1 < \cdots < a_k$.

We claim that assigning labels to each subset of the form $\{0, a_1, \ldots, a_k\}$ uniquely determines a valid k-cocycle on the n-simplex. We will begin by showing uniqueness, then check that it is a valid cocycle. Let our cocycle be called z, and consider the face $\{a_0, a_1, \ldots, a_k\}$ with $a_0 \ne 0$. The cocycle condition tells us that the values of z on the boundary of the $(k+1)$-face $\{0, a_0, \ldots, a_k\}$ add to zero in the alternating sum of the coboundary map. Rearranging this condition, we get

(2)
$$z(\{a_0, \ldots, a_k\}) = \sum_{i=0}^{k} (-1)^i z(\{0, \ldots, \hat{a}_i, \ldots\}).$$

The expression in the summation represents omitting only a_i from $\{0, a_0, \ldots, a_k\}$. The right-hand-side of (2) is clearly a well-defined value, so we have uniqueness.

Now we just have to show that the resulting cochain is indeed a cocycle. This amounts to checking that every $(k+1)$-face not containing vertex zero is labelled with zero under the coboundary map. To this end, consider the face $\{a_0, \ldots, a_{k+1}\}$. We must show that

$$(3) \qquad \sum_{i=0}^{k+1} (-1)^i z(\{\ldots, \hat{a}_i, \ldots\}) = 0.$$

We will substitute each term using equation (2). In the resulting summation, it is clear that all terms are multiples of $z(\{0, \ldots, \hat{a}_i, \ldots, \hat{a}_j, \ldots\})$ for $0 \le i < j \le k+1$. Fix some i and j, and consider the coefficient of this term. It is obtained in two ways: omitting a_j in (3) and then omitting a_i in (2), or vice versa. For the first way, we omit the j^{th} term and then the i^{th} term, giving a coefficient of $(-1)^j(-1)^i = (-1)^{i+j}$. For the second way, we omit the i^{th} term and then, since $j > i$, we omit the $(j-1)^{\text{st}}$ term, for a coefficient of $(-1)^i(-1)^{j-1} = (-1)^{i+j-1}$. Therefore, the overall coefficient of this term is

$$(-1)^{i+j} + (-1)^{i+j-1} = 0.$$

Since this term was chosen arbitrarily, the overall expression (3) is zero, as desired.

We now know that, as claimed, labelling each subset of $[n]$ containing 0 gives a unique k-cocycle on the n-simplex. Since there are $\binom{n}{k}$ total k-simplices of the form $\{0, a_1, \ldots, a_k\}$, it follows that

$$|K(\pi, k)_n| = |\pi|^{\binom{n}{k}}.$$

$\square$

Now we are prepared to bound the size of $sk_n(P_k)$ using our understanding of the sizes of the simplicial set models of the $K(\pi, k)$s.

Proposition 5.2. *Let P_k be the k^{th} Postnikov stage of a simplicial set X. Then*

$$|(sk_n(P_k))| \le \sum_{\ell=0}^{n} \prod_{j=2}^{k} |\pi_j(X)|^{\binom{\ell}{j}}.$$

Proof. We proceed by induction on k. Recall the pullback in Figure 2 used to construct P_k from P_{k-1}.

As mentioned in [Bro57], the fibration $\delta_* : E(\pi_k(X), k) \to K(\pi_k(X), k+1)$ is a fiber bundle with fiber $K(\pi_k(X), k)$. Therefore, since it is surjective, for each $p \in P_{k-1}$ the subset of P_k corresponding to it in the pullback is isomorphic to $K(\pi_k(X), k)$. It follows that

$$|(P_k)_\ell| \le |(P_{k-1})_\ell||(K(\pi_k(X), k)_\ell)|.$$

Since $|K(\pi, k)_\ell| = |\pi|^{\binom{\ell}{k}}$, induction on k gives us

$$|((P_k)_\ell| \le \prod_{j=2}^{k} |\pi_j(Y)|^{\binom{\ell}{j}}$$

Adding together over all ℓ from 0 to n gives us the desired expression. $\square$

Now we can bound the size of $sk_{d+2}(P_{d+1})$, which was used to construct Y in Lemma 4.5.

Lemma 5.3.

$$|(sk_{d+2}(P_{d+1}))| \leq \exp\left(mh\log(N)\exp\left(d\log(2h) + O\left(\log(d)^3\right)\right)\right)$$

Proof. By plugging the formula for a upper-bound for $\pi_j(X)$ from Proposition 3.4 into the formula for a upper bound for $|(sk_{d+2}(P_{d+1}))|$ from Proposition 5.2 we get that

$$|(sk_{d+2}(P_{d+1}))| \leq \sum_{\ell=0}^{d+2}\prod_{j=2}^{d+1}|\pi_j(X)|^{\binom{\ell}{j}}$$

$$\leq \sum_{\ell=0}^{d+2}\prod_{j=2}^{d+1}\left(\exp(m\log(N)\exp(j\log(h) + O\left(\log(j)^3\right)))\right)^{\binom{\ell}{j}}$$

We upper-bound the sum using $\sum_{\ell=0}^{d+2} a_\ell \leq (d+3)\max(a_\ell)$, noting that $\max(a_\ell)$ is achieved here at $\ell = d + 2$.

$$\sum_{\ell=0}^{d+2}\prod_{j=2}^{d+1}\left(\exp(m\log(N)\exp(j\log(h) + O\left(\log(j)^3\right)))\right)^{\binom{\ell}{j}}$$

$$\leq (d+3)\prod_{j=2}^{d+1}\left(\exp(m\log(N)\exp(j\log(h) + O\left(\log(j)^3\right)))\right)^{\binom{d+2}{j}}$$

We express the product of exponentials as a exponential of a sum.

$$(d+3)\prod_{j=2}^{d+1}\left(\exp(m\log(N)\exp(j\log(h) + O\left(\log(j)^3\right)))\right)^{\binom{d+2}{j}}$$

$$\leq (d+3)\left(\exp\left(m\log(N)\sum_{j=2}^{d+1}\left(\binom{d+2}{j}\exp\left(j\log(h) + O\left(\log(j)^3\right)\right)\right)\right)\right)$$

We upper-bound the sum using $\sum_{j=2}^{d+1} a_j b_j c_j \leq d\max(a_j)\max(b_j)\max(c_j)$.

$$(d+3)\left(\exp\left(m\log(N)\sum_{j=2}^{d+1}\left(\binom{d+2}{j}\exp\left(j\log(h) + O\left(\log(j)^3\right)\right)\right)\right)\right)$$

$$\leq (d+3)\left(\exp\left(m\log(N)d\left(\exp(\log(2)d)\exp\left((d+1)\log(h) + O\left(\log(d+1)^3\right)\right)\right)\right)\right)$$

We rewrite $d+3$ as $\exp\log(d+3)$ on the outside, rewrite d as $\exp\log(d)$ on the inside, write $\exp(a)\exp(b)$ as $\exp(a+b)$, and use $O\left(\log(d+1)^3\right) = O\left(\log(d)^3\right)$.

$$(d+3)\left(\exp\left(m\log(N)d\left(\exp(\log(2)d)\exp\left((d+1)\log(h) + O\left(\log(d+1)^3\right)\right)\right)\right)\right)$$

$$\leq \exp\left(\log(d+3) + m\log(N)\exp(\log(d))\left(\exp\left(\log(2)d + d\log(h) + \log(h) + O\left(\log(d)^3\right)\right)\right)\right)$$

We use $a + b \leq ab$ for sufficiently large a, b and write $\exp(a)\exp(b)$ as $\exp(a + b)$.

$$\leq \exp\left(\log(d + 3) + m\log(N)\exp(\log(d))\left(\exp\left(\log(2)d + d\log(h) + \log(h) + O\left(\log(d)^3\right)\right)\right)\right)$$
$$\leq \exp\left(\log(d + 3)m\log(N)\left(\exp\left(\log(d) + (\log(2h))d + \log(h) + O\left(\log(d)^3\right)\right)\right)\right)$$

We rewrite $\log(d + 3)$ as $\exp(\log\log(d + 3))$ and use $\exp(a)\exp(b) = \exp(a + b)$.

$$\exp\left(\log(d + 3)m\log(N)\left(\exp\left(\log(d) + (\log(2h))d + \log(h^2) + O\left(\log(d)^3\right)\right)\right)\right)$$
$$\leq \exp\left(m\log(N)\left(\exp\left(\log\log(d + 3) + \log(d) + (\log(2h))d + \log(h^2) + O\left(\log(d)^3\right)\right)\right)\right)$$

We use $\log\log(d + 3) + \log(d) = O(\log(d)^3)$.

$$\exp\left(m\log(N)\left(\exp\left(\log\log(d + 3) + \log(d) + (\log(2h))d + \log(h) + O\left(\log(d)^3\right)\right)\right)\right)$$
$$\leq \exp\left(m\log(N)\left(\exp\left((\log(2h))d + \log(h) + O\left(\log(d)^3\right)\right)\right)\right)$$

We write $\exp(\log(h))$ as h.

$$\exp\left(m\log(N)\left(\exp\left((\log(2h))d + \log(h) + O\left(\log(d)^3\right)\right)\right)\right)$$
$$\leq \exp\left(mh\log(N)\exp\left(\log(2h)d + O\left(\log(d)^3\right)\right)\right).$$

$\square$

Using this, we can easily prove our main theorem.

Theorem 5.4. *Let X be a finite, simply-connected simplicial set of dimension d with finite integral homology, and let p be a prime. Then there exists an algorithm giving a simplicial set which is p-locally homotopy equivalent to X with at most*

$$\exp\left(mh\log(N)\exp\left(\log(2h)d + O\left(\log(d)^3\right)\right)\right)$$

non-degenerate simplices.

Proof. From Lemma 4.5 we know there exists an algorithm giving a finite simplicial set $Y \subset sk_{d+2}(P_{d+1})$ which is p-locally homotopy equivalent to X. Since $Y \subset sk_{d+2}(P_{d+1})$, we know that the size of Y is upper-bounded by that of $sk_{d+2}(P_{d+1})$, which Lemma 5.3 bounds by the desired expression. $\square$

Note that our main theorem implies the following result about homotopy types in general.

Corollary 5.5. *Let X be a simply-connected, finite homotopy type with finite integral homology whose highest non-trivial homology is in degree d. Then there exists a finite simplicial set S representing a space p-locally equivalent to X with at most*

$$\exp\left(mh\log(N)\exp\left(\log(2h)(d + 1) + O\left(\log(d)^3\right)\right)\right).$$

non-degenerate simplices.

Proof. Since X is a finite homotopy type, it has a finite simplicial set representative to which we can apply the main theorem. However, we also wish to remove the dependence on dimension. To this end, note that X is simply-connected and its homology is finitely generated, so there exists a finite CW complex homotopy equivalent to X with dimension $d+1$ [Hat02, Proposition 4C.1]. It then follows by simplicial approximation that there exists a finite simplicial set homotopy equivalent to this CW complex with the same dimension. Applying the main theorem to this simplicial set gives the desired result. $\square$

6. Future Directions

In this section we consider potential future directions of our work on this problem. There are two main extensions that we wish to consider. For one, we would like to reduce the size of the simplicial set generated in Theorem 5.4. We would also like to generalize the algorithm to work when X has infinite homology groups, ideally finding a simplicial set which is homotopy equivalent fully rather than just p-locally. The latter direction is useful because it allows for the algorithm to be used in the calculation of the homotopy groups of any simply connected, finite simplicial set.

For both of the goals mentioned above, a potential strategy to explore centers on finding smaller simplicial set representatives for the Eilenberg-Maclane spaces $K(\pi, n)$. Since our size analysis centers on a twisted product of such spaces, this has a good chance of allowing us to reduce our size bound. In addition, the main obstacle to using our current strategy in the general case is that when a homotopy group has $\mathbb{Z}$ summands, the standard simplicial model of $K(\pi, n)$ is no longer levelwise finite, meaning that our current strategy yields an infinite simplicial set. Thus, finding better simplicial set representatives for these spaces would be very helpful for improving the results of this paper.

One method to consider in pursuit of this is inspired by the method used by Krčál et. al. in [KMS13] to equip $K(\mathbb{Z}, 1)$ with polynomial time homology (defined in [ČKM+14]; the precise definition is not important for us). They use structures called discrete vector fields.

Definition 6.1. Let X be a simplicial set, and let τ be a nondegenerate simplex of X. A *regular face* of τ is a simplex σ such that $\sigma = d_i\tau$ for exactly one value of i.

Definition 6.2. Let X be a simplicial set. A *discrete vector field* on X is a set of ordered pairs (σ, τ) such that σ is a regular face of τ and such that no face appears in more than one ordered pair. The σ faces are called *sources*, the τ faces are called *targets*, and the remaining faces are called *critical faces*.

The reason that these structures may yield results is due to their apparent relation to Kan completion. When the Kan completion functor is applied to a simplicial set, it fills in the horns that exist in the set; discrete vector fields seem to identify a way to "un-fill-in" horns that had been filled in. This is promising because Kan completion produces a homotopy equivalent space, so going in the reverse direction may also yield a homotopy equivalent space with a smaller simplicial set representative. The concept of discrete vector fields may have to be refined further in order to achieve this, however.

References

[Bar60] M. G. Barratt, *Spaces of finite characteristic*, Quart. J. Math. Oxford Ser. (2) **11** (1960), 124–136. MR120647

[Bou75] A. K. Bousfield, *The localization of spaces with respect to homology*, Topology **14** (1975), 133–150. MR380779

[Bro57] Edgar H. Brown Jr., *Finite computability of Postnikov complexes*, Ann. of Math. (2) **65** (1957), 1–20. MR83733

[BS] Robert Burklund and Andy Senger, *Appendix to: How big are the stable homotopy groups of spheres?*. In preparation.

[Hat02] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002. MR1867354

[KMS13] Marek Krčál, Jiří Matoušek, and Francis Sergeraert, *Polynomial-time homology for simplicial Eilenberg-MacLane spaces*, Found. Comput. Math. **13** (2013), no. 6, 935–963. MR3124946

[Mat14] J. Matoušek, *Computing higher homotopy groups is w[1]-hard*, 2014.

[May67] J. Peter May, *Simplicial objects in algebraic topology*, Van Nostrand Mathematical Studies, No. 11, D. Van Nostrand Co., Inc., Princeton, N.J.-Toronto, Ont.-London, 1967. MR0222892

[Ser53] Jean-Pierre Serre, *Groupes d'homotopie et classes de groupes abéliens*, Ann. of Math. (2) **58** (1953), 258–294. MR59548

[ČKM⁺14] Martin Čadek, Marek Krčál, Jiří Matoušek, Lukáš Vokřínek, and Uli Wagner, *Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension*, SIAM J. Comput. **43** (2014), no. 5, 1728–1780. MR3268623

Massachusetts Institute of Technology, Cambridge, MA 02139, USA
Email address: `prestonc@mit.edu`

Massachusetts Institute of Technology, Cambridge, MA 02139, USA
Email address: `rowley@mit.edu`