The codimension-one cohomology of $\mathrm{SL}_n \mathbb{Z}$

Thomas Church and Andrew Putman[*]

August 2, 2015

Abstract

We prove that $\mathrm{H}^{\binom{n}{2}-1}(\mathrm{SL}_n \mathbb{Z}; \mathbb{Q}) = 0$, where $\binom{n}{2}$ is the cohomological dimension of $\mathrm{SL}_n \mathbb{Z}$, and similarly for $\mathrm{GL}_n \mathbb{Z}$. We also prove analogous vanishing theorems for cohomology with coefficients in a rational representation of the algebraic group GL_n. These theorems are derived from a presentation of the Steinberg module for $\mathrm{SL}_n \mathbb{Z}$ whose generators are integral apartment classes, generalizing Manin's presentation for the Steinberg module of $\mathrm{SL}_2 \mathbb{Z}$. This presentation was originally constructed by Bykovskii. We give a new topological proof of it.

1 Introduction

The cohomology of $\mathrm{SL}_n \mathbb{Z}$ plays a fundamental role in many areas of mathematics. The Borel Stability Theorem [Bo] determines $\mathrm{H}^k(\mathrm{SL}_n \mathbb{Z}; \mathbb{Q})$ when k is sufficiently small (conjecturally, for $k < n-1$). However, little is known outside this stable range. Recall that if Γ is a virtually torsion-free group, the virtual cohomological dimension of Γ is

$$\mathrm{vcd}(\Gamma) := \max\{k \mid \mathrm{H}^k(\Gamma; V \otimes \mathbb{Q}) \neq 0 \text{ for some } \Gamma\text{-module } V\}.$$

Borel–Serre [BoSe] proved that

$$\mathrm{vcd}(\mathrm{SL}_n \mathbb{Z}) = \mathrm{vcd}(\mathrm{GL}_n \mathbb{Z}) = \binom{n}{2}.$$

The cohomology of $\mathrm{SL}_n \mathbb{Z}$ in degrees near $\binom{n}{2}$ is thus the "most unstable" cohomology. In 1976, Lee–Szczarba [LSz, Theorem 1.3] proved that $\mathrm{H}^{\binom{n}{2}}(\mathrm{SL}_n \mathbb{Z}; \mathbb{Q}) = 0$. This vanishing was recently extended to $\mathrm{H}^{\binom{n}{2}}(\mathrm{SL}_n \mathbb{Z}; V_\lambda) = 0$ for rational representations V_λ of the algebraic group GL_n by Church–Farb–Putman [CFP2].

Our main theorem concerns the cohomology in codimension 1. For $\lambda = (\lambda_1, \ldots, \lambda_n) \in \mathbb{Z}^n$ with $\lambda_1 \geq \cdots \geq \lambda_n$, let V_λ be the rational representation of $\mathrm{GL}_n \mathbb{Q}$ with highest weight λ. Define $\|\lambda\| = \sum_{i=1}^n (\lambda_i - \lambda_n)$.

Theorem A (**Codimension-one Vanishing Theorem**). *For any rational representation V_λ of $\mathrm{GL}_n \mathbb{Q}$, we have*

$$\mathrm{H}^{\binom{n}{2}-1}(\mathrm{SL}_n \mathbb{Z}; V_\lambda) = \mathrm{H}^{\binom{n}{2}-1}(\mathrm{GL}_n \mathbb{Z}; V_\lambda) = 0 \qquad \text{for all } n \geq 3 + \|\lambda\|.$$

In particular,

$$\mathrm{H}^{\binom{n}{2}-1}(\mathrm{SL}_n \mathbb{Z}; \mathbb{Q}) = \mathrm{H}^{\binom{n}{2}-1}(\mathrm{GL}_n \mathbb{Z}; \mathbb{Q}) = 0 \qquad \text{for all } n \geq 3.$$

Remark 1.1. Theorem A proves the $k = 1$ case of a conjecture of Church–Farb–Putman [CFP1] which asserts that for all $k \geq 0$, we have $\mathrm{H}^{\binom{n}{2}-k}(\mathrm{SL}_n \mathbb{Z}; \mathbb{Q}) = 0$ for $n > k + 1$.

Steinberg module. To compute $\mathrm{H}^{\binom{n}{2}-k}(\mathrm{SL}_n \mathbb{Z}; \mathbb{Q})$ for small k, the crucial object to understand is the *Steinberg module*, which we now discuss. The Tits building $\mathcal{T}_n$ for $\mathrm{GL}_n \mathbb{Q}$ is the simplicial complex whose p-simplices are flags of subspaces

$$0 \subsetneq V_0 \subsetneq \cdots \subsetneq V_p \subsetneq \mathbb{Q}^n.$$

By the Solomon-Tits theorem, $\mathcal{T}_n$ is homotopy equivalent to an infinite wedge of $(n-2)$-dimensional spheres. The group $\mathrm{GL}_n \mathbb{Q}$ naturally acts on $\mathcal{T}_n$ by simplicial automorphisms, and the *Steinberg module* for $\mathrm{GL}_n \mathbb{Q}$ is the $\mathrm{GL}_n \mathbb{Q}$-module

$$\mathrm{St}_n := \widetilde{\mathrm{H}}_{n-2}(\mathcal{T}_n; \mathbb{Z}).$$

Borel–Serre duality. While $\mathrm{SL}_n \mathbb{Z}$ does not satisfy Poincaré duality, Borel–Serre [BoSe] proved that it does satisfy *virtual Bieri–Eckmann duality* with rational dualizing module St_n. This means that for any $\mathbb{Q} \, \mathrm{SL}_n \mathbb{Z}$-module V and any $k \geq 0$, we have

$$\mathrm{H}^{\binom{n}{2}-k}\left(\mathrm{SL}_n \mathbb{Z}; V\right) \cong \mathrm{H}_k\left(\mathrm{SL}_n \mathbb{Z}; \mathrm{St}_n \otimes V\right). \tag{1}$$

A similar result holds for $\mathrm{GL}_n \mathbb{Z}$ (with the same dualizing module St_n).

Presentation of St_n. To prove Theorem A, we compute the right-hand side of (1) using a presentation of the $\mathrm{GL}_n \mathbb{Z}$-module St_n. The following theorem was originally proved by Bykovskii [By] and generalizes the presentation of St_2 given by Manin in 1972 [Man, Theorem 1.9]. The second purpose of this paper is to offer a new proof of it.

Theorem B (Presentation of St_n). *For $n \geq 2$, the Steinberg module St_n is the abelian group with generators $[v_1, \ldots, v_n]$, one for each ordered basis $\{v_1, \ldots, v_n\}$ of $\mathbb{Z}^n$, subject to the following three families of relations.*
  R1. $[v_1, v_2, v_3, \ldots, v_n] = [v_1, v_1 + v_2, v_3, \ldots, v_n] + [v_1 + v_2, v_2, v_3, \ldots, v_n].$
  R2. $[\pm v_1, \pm v_2, \ldots, \pm v_n] = [v_1, v_2, \ldots, v_n]$ *for any choices of signs.*
  R3. $[v_{\sigma(1)}, v_{\sigma(2)}, \ldots, v_{\sigma(n)}] = (-1)^\sigma \cdot [v_1, v_2, \ldots, v_n]$ *for any $\sigma \in S_n$.*
The $\mathrm{GL}_n \mathbb{Z}$-action on St_n is defined by $M \cdot [v_1, \ldots, v_n] = [M \cdot v_1, \ldots, M \cdot v_n]$.

The relation R2 can be replaced by the single relation $[-v_1, v_2, \ldots, v_n] = [v_1, v_2, \ldots, v_n]$. Relations R1 and R2 then involve *only the first two vectors*; these relations are the "stabilization" of Manin's relations for St_2 from $\mathrm{GL}_2 \mathbb{Z}$ to $\mathrm{GL}_n \mathbb{Z}$. In this light, what Theorem B says is that *no additional relations* are needed to present St_n, once the permutations S_n are taken into account.

Apartment classes and Bykovskii's proof of Theorem B. The general theory of spherical buildings automatically provides a generating set for St_n. Namely, every rational basis $\{w_1, \ldots, w_n\}$ for $\mathbb{Q}^n$ determines a spherical apartment in $\mathcal{T}_n$ homeomorphic to S^{n-2} whose fundamental class determines an apartment class $[w_1, \ldots, w_n] \in \widetilde{\mathrm{H}}_{n-2}(\mathcal{T}_n; \mathbb{Z}) = \mathrm{St}_n$, and the general theory implies that St_n is generated by these rational apartment classes. However, the generating set in Theorem B is much smaller: it consists only of the *integral* apartment classes $[v_1, \ldots, v_n]$, i.e. those for which $\{v_1, \ldots, v_n\}$ is an integral basis for $\mathbb{Z}^n$.

That St_n is generated by these integral apartment classes was proved by Ash–Rudolph [AR] in 1979. To do this, they gave an algorithm for expressing an arbitrary rational apartment class as a sum of integral apartment classes.

Bykovskii proved Theorem B by carefully examining Ash–Rudolph's algorithm, which requires making many arbitrary choices, and showing that the only ambiguity in its output comes from the relations in Theorem B. We remark that from this perspective, Theorem B appears as the integral analogue of Lee–Szczarba's presentation of St_n as a $\mathrm{GL}_n\,\mathbb{Q}$-module [LSz].

Our proof of Theorem B. Our proof of Theorem B is quite different. It is inspired by our alternate proof of Ash–Rudolph's theorem in Church–Farb–Putman [CFP2] and by Manin's original proof of Theorem B for St_2. We use topology to show directly that the homology of $\mathcal{T}_n$ is generated by integral apartment classes; non-integral apartment classes never show up in our proof. The key is the *complex of partial augmented frames* for $\mathbb{Z}^n$ defined below, which provides an "integral model" for the Tits building $\mathcal{T}_n$. We begin first with the more familiar complex of partial frames.

Definition 1.2. Let V be a finite-rank free abelian group.
- A *line* in V is a 2-element set $\{v, -v\}$ of primitive vectors in V; we denote it by $v^{\pm}$.
- A *frame* for V is a set $\{v_1^{\pm}, \ldots, v_n^{\pm}\}$ of lines such that $\{v_1, \ldots, v_n\}$ is a basis for V.
- A *partial frame* for V is a frame for a direct summand of V, or equivalently a set of lines in V that can be completed to a frame for V.

The *complex of partial frames* for $\mathbb{Z}^n$, denoted $\mathcal{B}_n$, is the simplicial complex whose p-simplices are partial frames for $\mathbb{Z}^n$ of cardinality $(p+1)$.

The complex $\mathcal{B}_n$ is $(n-1)$-dimensional, and Maazen [Maa] proved that $\mathcal{B}_n$ is $(n-2)$-connected. This connectivity is what we used in [CFP2] to prove Ash–Rudolph's theorem on generators for St_n. However, to obtain a *presentation* for St_n this is not enough; we need to attach higher-dimensional cells to $\mathcal{B}_n$ to improve its connectivity.

Improving connectivity: the complex of partial augmented frames. To motivate the cells we add, we recall how Manin found his presentation for St_2. The simplicial complex $\mathcal{B}_2$ is the following graph.
- The vertices are lines $(a, b)^{\pm}$, where $(a, b) \in \mathbb{Z}^2$ is a primitive vector.
- Vertices $(a, b)^{\pm}$ and $(c, d)^{\pm}$ are joined by an edge exactly when $\{(a, b), (c, d)\}$ is a basis for $\mathbb{Z}^2$, or equivalently when $ad - bc = \pm 1$.

This is exactly the classical Farey graph; see Figure 1. This graph is connected, but not simply-connected. The reduced 0-chains $\widetilde{C}_0(\mathcal{B}_2; \mathbb{Z})$ can be identified (essentially by definition) with $\mathrm{St}_2 = \widetilde{\mathrm{H}}_0(\mathcal{T}_2; \mathbb{Z})$, and it is not hard to see that the 1-chains $C_1(\mathcal{B}_2; \mathbb{Z})$ have a presentation just like that in Theorem B, but where we impose only the relations R2 and R3. What this means is that the first homology $\mathrm{H}_1(\mathcal{B}_2; \mathbb{Z})$ measures exactly the *additional* relations beyond R2 and R3 needed to present St_2.

As is clear from Figure 1, the group $\mathrm{H}_1(\mathcal{B}_2; \mathbb{Z})$ is spanned by the boundaries of the evident triangles in the Farey graph. Under our identification of $\mathcal{B}_2$ with the Farey graph, these triangles are formed by triples of vertices $\{v_1^{\pm}, v_2^{\pm}, v_3^{\pm}\}$ such that $\{v_1, v_2\}$ is a basis for $\mathbb{Z}^2$ and $\pm v_1 \pm v_2 \pm v_3 = 0$ for some choice of signs (in which case $\{v_2, v_3\}$ and $\{v_1, v_3\}$ are also bases for $\mathbb{Z}^2$). Reordering and changing the signs of the v_i, we can assume that $v_3 = v_1 + v_2$. The relation in St_2 corresponding to the boundary of the triangle $\{v_1^{\pm}, v_2^{\pm}, (v_1 + v_2)^{\pm}\}$ is

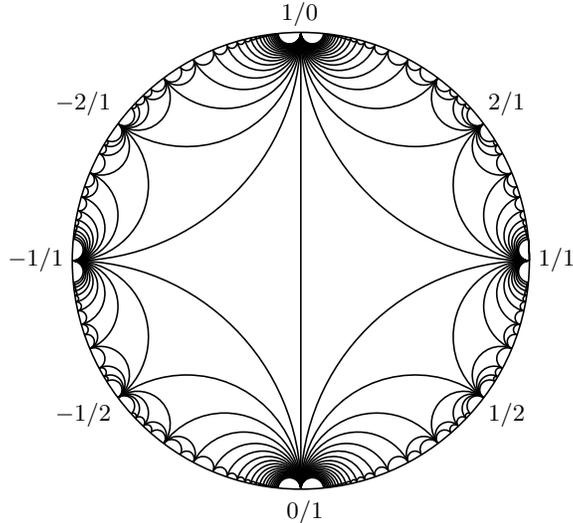$$[v_1, v_1 + v_2] - [v_1, v_2] + [v_1 + v_2, v_2] = 0,$$

3

Figure 1: *The complex $\mathcal{B}_2$ is isomorphic to the Farey graph under the identification that takes the line $(a, b)^{\pm} \in \mathbb{Z}^2$ to $\frac{a}{b} \in \mathbb{Q} \cup \{\infty\}$.*

which is precisely the relation R1.

Manin's theorem that the relations R1 together with R2 and R3 suffice to present St_2 thus follows from the fact that attaching the above triangles to the Farey graph yields a simply-connected simplicial complex. This motivates the following definition.

Definition 1.3. Let V be a finite-rank free abelian group.
- An *augmented frame* for V is a collection $\{v_0^{\pm}, v_1^{\pm}, \ldots, v_n^{\pm}\}$ of lines in V such that $\{v_1^{\pm}, \ldots, v_n^{\pm}\}$ is a frame for V and $\pm v_0 \pm v_1 \pm v_2 = 0$ for some choice of signs.
- A *partial augmented frame* for V is a set of lines in V that is either a frame or an augmented frame for a direct summand of V; equivalently, a set of lines is a partial augmented frame for V if it can be completed to an augmented frame for V.

The *complex of partial augmented frames* for $\mathbb{Z}^n$, denoted $\mathcal{BA}_n$, is the simplicial complex whose p-simplices are the partial augmented frames for $\mathbb{Z}^n$ of cardinality $(p + 1)$.

The final main theorem of this paper is as follows; the definition of a Cohen–Macaulay complex is recalled in §2.1 below. We remark that this theorem plays a fundamental role in the second author's recent work with Day on the second homology group of the Torelli subgroup of $\mathrm{Aut}(F_n)$; see [DP].

Theorem C ($\mathcal{BA}_n$ is Cohen–Macaulay). *For all $n \geq 2$, the complex $\mathcal{BA}_n$ is Cohen–Macaulay of dimension n. In particular, $\mathcal{BA}_n$ is $(n-1)$-connected.*

Remark 1.4. Since $\mathcal{BA}_n$ is n-dimensional, the connectivity in Theorem C cannot be improved unless $\mathcal{BA}_n$ is contractible. Since $\mathcal{BA}_2$ is the complex obtained by filling in the triangles in the Farey graph, the complex $\mathcal{BA}_2$ is contractible. However, it seems unlikely that $\mathcal{BA}_n$ would be contractible for any $n \geq 3$.

Outline. The logical relation between our three main theorems is that Theorem C $\implies$ Theorem B $\implies$ Theorem A. However, the proof of Theorem C occupies more than half of the paper, so we defer the proof of Theorem C until §4. We prove Theorem B in §2 and prove Theorem A in §3, both assuming Theorem C.

Acknowledgments. We are very grateful to Benson Farb, who was closely involved during the development of these results, but declined to be listed as a coauthor.

4

2 Generators and relations for the Steinberg module

In this section, we derive Theorem B from Theorem C, which will be proved in §4. We begin in §2.1 with some basic results about posets and then prove Theorem B in §2.2.

2.1 The topology of posets

Recall that a d-dimensional complex is d-*spherical* if it is $(d-1)$-connected, in which case it is homotopy equivalent to a wedge of d-spheres. A simplicial complex X is *Cohen–Macaulay* (abbreviated CM) of dimension d if the following hold.
- X is d-spherical.
- For every $(k-1)$-simplex σ^{k-1} of X, the link $\mathrm{Link}_X(\sigma^{k-1})$ is $(d-k)$-spherical.

Remark 2.1. This should be compared with the definition of a combinatorial d-manifold, which is a d-dimensional simplicial complex M such that for every $(k-1)$-simplex σ^{k-1}, the link $\mathrm{Link}_M(\sigma^{k-1})$ is a combinatorial $(d-k)$-sphere.

Let A be a poset. Recall that the *geometric realization* of A is the simplicial complex $|A|$ whose k-simplices are chains $a_0 \lneq a_1 \lneq \cdots \lneq a_k$ in A. Whenever we say that A has some topological property, we mean that $|A|$ has that topological property. In particular, we define $\mathrm{H}_*(A)$ to equal $\mathrm{H}_*(|A|)$. The following is a key example.

Example 2.2. Let X be a simplicial complex. Define $\mathcal{P}(X)$ to be the poset of simplices of X under inclusion. Then $|\mathcal{P}(X)|$ is isomorphic to the barycentric subdivision of X. In particular, there is a canonical isomorphism $\mathrm{H}_*(\mathcal{P}(X)) \cong \mathrm{H}_*(X)$.

For $a \in A$, the *height* of a, denoted $\mathrm{ht}(a)$, is the maximal k such that there exists a chain $a_0 \lneq a_1 \lneq \cdots \lneq a_k = a$ in A. If B is another poset and $F\colon A \to B$ is a poset map, for $b \in B$ we define $F_{\leq b}$ to be the subposet $\{a \in A \mid F(a) \leq b\}$ of A. With these definitions, we have the following proposition, which slightly generalizes a result of Quillen.

Proposition 2.3. *Fix $m \geq 0$ and let $F\colon A \to B$ be a map of posets. Assume that B is CM of dimension d and that for all $b \in B$, the fiber $F_{\leq b}$ is $(\mathrm{ht}(b) + m)$-spherical (or more generally, that $\widetilde{\mathrm{H}}_q(F_{\leq b}) = 0$ for $q \neq \mathrm{ht}(b) + m$). Then $F\colon A \to B$ is $(d+m)$-acyclic. In particular, $F_*\colon \widetilde{\mathrm{H}}_d(A) \to \widetilde{\mathrm{H}}_d(B)$ is an isomorphism if $m \geq 1$.*

Proof. This can be proved exactly like [Q, Theorem 9.1]. The necessary conditions on B are satisfied since it is CM. Quillen's hypothesis that $F_{\leq b}$ is $\mathrm{ht}(b)$-spherical is used only to conclude that $\widetilde{\mathrm{H}}_q(F_{\leq b}) = 0$ for $q \neq \mathrm{ht}(b)$, so we can replace this with the hypothesis that $\widetilde{\mathrm{H}}_q(F_{\leq b}) = 0$ for $q \neq \mathrm{ht}(b) + m$. We conclude that the spectral sequence $E_{pq}^2 \implies \mathrm{H}_{p+q}(A)$ of [Q, (9.3)] vanishes in the range $p + q < d + m$, except for $E_{d0}^2 = \mathrm{H}_d(B)$. Therefore $F\colon A \to B$ is $(d+m)$-acyclic, as claimed. $\qquad\square$

2.2 The proof of Theorem B

We now prove Theorem B. To do this we will use Theorem C, whose proof is postponed until §4, in two different places.

The subcomplex $\mathcal{BA}_n'$. We begin by defining a subcomplex $\mathcal{BA}_n'$ of $\mathcal{BA}_n$. Consider a simplex $\sigma = \{v_1^{\pm}, \ldots, v_k^{\pm}\}$ of $\mathcal{BA}_n$. By definition, the submodule $\mathrm{span}_{\mathbb{Z}}(v_1, \ldots, v_k)$ of $\mathbb{Z}^n$ is a direct summand (of rank k if σ is a partial frame and of rank $(k-1)$ otherwise). Let $\mathcal{BA}_n'$ be the subcomplex of $\mathcal{BA}_n$ consisting of simplices $\sigma = \{v_1^{\pm}, \ldots, v_k^{\pm}\}$ of $\mathcal{BA}_n$ such that $\mathrm{span}_{\mathbb{Z}}(v_1, \ldots, v_k)$ is a *proper* direct summand of $\mathbb{Z}^n$. The only simplices that are omitted from $\mathcal{BA}_n'$ are the frames of $\mathbb{Z}^n$ (which are $(n-1)$-simplices) and the augmented frames of $\mathbb{Z}^n$ (which are n-simplices).

The proof now has three main steps.

Step 1. *The abelian group described by the presentation in Theorem B coincides with the relative homology $\mathrm{H}_{n-1}(\mathcal{BA}_n, \mathcal{BA}_n'; \mathbb{Z})$.*

To prove this, we will compute the relative homology via the relative simplicial chain complex $C_*(\mathcal{BA}_n, \mathcal{BA}_n')$. Since the $(n-2)$-skeleton of $\mathcal{BA}_n$ is contained in $\mathcal{BA}_n'$, we have $C_k(\mathcal{BA}_n, \mathcal{BA}_n') = 0$ for $k \leq n-2$. It follows that

$$\mathrm{H}_{n-1}(\mathcal{BA}_n, \mathcal{BA}_n'; \mathbb{Z}) = \mathrm{coker}(C_n(\mathcal{BA}_n, \mathcal{BA}_n') \xrightarrow{\partial} C_{n-1}(\mathcal{BA}_n, \mathcal{BA}_n')).$$

Define $\mathcal{I}_0 = C_{n-1}(\mathcal{BA}_n, \mathcal{BA}_n')$ and $\mathcal{I}_1 = C_n(\mathcal{BA}_n, \mathcal{BA}_n')$. Our goal is to describe $\mathcal{I}_0$ and $\mathcal{I}_1$ and the differential ∂.

The simplices that contribute to $\mathcal{I}_0$ are the $(n-1)$-simplices of $\mathcal{BA}_n$ that do not lie in $\mathcal{BA}_n'$, i.e. those corresponding to frames $\{v_1^{\pm}, \ldots, v_n^{\pm}\}$ for $\mathbb{Z}^n$. To specify such a frame, it is enough to give the vectors $v_1, \ldots, v_n$. The only ambiguity is that multiplying the vectors v_i by ± 1 does not change the frame, nor does permuting the vectors; however, permuting the vectors does change the orientation of the corresponding simplex. We deduce that $\mathcal{I}_0$ is the abelian group with generators the set of formal symbols $\langle\!\langle v_1 \ldots, v_n \rangle\!\rangle$ for bases $\{v_1, \ldots, v_n\}$ of $\mathbb{Z}^n$ subject to the following relations.

   S2. $\langle\!\langle \pm v_1, \pm v_2, \ldots, \pm v_n \rangle\!\rangle = \langle\!\langle v_1, v_2, \ldots, v_n \rangle\!\rangle$ for any choices of signs.

   S3. $\langle\!\langle v_{\sigma(1)}, v_{\sigma(2)}, \ldots, v_{\sigma(n)} \rangle\!\rangle = (-1)^{\sigma} \cdot \langle\!\langle v_1, v_2, \ldots, v_n \rangle\!\rangle$ for any $\sigma \in S_n$.

The simplices that contribute to $\mathcal{I}_1$ are the n-simplices of $\mathcal{BA}_n$ that do not lie in $\mathcal{BA}_n'$ (a vacuous condition since $\mathcal{BA}_n'$ does not contain any n-simplices). These correspond to augmented frames $\{v_0^{\pm}, \ldots, v_n^{\pm}\}$ for $\mathbb{Z}^n$. By definition, $\{v_1, \ldots, v_n\}$ is a basis for $\mathbb{Z}^n$ and $\pm v_0 \pm v_1 \pm v_2 = 0$ for some choice of signs. Multiplying v_0 and v_1 and v_2 by appropriate choices of signs, we can arrange for $v_0 = v_1 + v_2$. For $3 \leq i \leq n$, the set $\{v_0, \ldots, \widehat{v}_i, \ldots, v_n\}$ spans a proper direct summand of $\mathbb{Z}^n$, so this term of the boundary vanishes in $\mathcal{I}_0$. This implies that under the boundary map ∂, the generator of $\mathcal{I}_1$ corresponding to the augmented frame $\{(v_1 + v_2)^{\pm}, v_1^{\pm} \ldots, v_n^{\pm}\}$ has image in $\mathcal{I}_0$ equal to

$$\langle\!\langle v_1, \ldots, v_n \rangle\!\rangle - \langle\!\langle v_1 + v_2, v_2, \ldots, v_n \rangle\!\rangle + \langle\!\langle v_1 + v_2, v_1, v_3, \ldots, v_n \rangle\!\rangle.$$

Applying the relation S2 and rearranging, we see that $\mathrm{H}_{n-1}(\mathcal{BA}_n, \mathcal{BA}_n'; \mathbb{Z}) = \mathrm{coker}(\partial)$ is the quotient of $\mathcal{I}_0$ by the set of relations

   S1. $\langle\!\langle v_1, v_2, v_3, \ldots, v_n \rangle\!\rangle = \langle\!\langle v_1, v_1 + v_2, v_3, \ldots, v_n \rangle\!\rangle + \langle\!\langle v_1 + v_2, v_2, v_3, \ldots, v_n \rangle\!\rangle.$

The relations S1, S2, and S3 correspond exactly to the relations R1, R2, and R3 in Theorem B, yielding the identity claimed in Step 1.

Step 2. *We have $\mathrm{H}_{n-1}(\mathcal{BA}_n, \mathcal{BA}_n'; \mathbb{Z}) \cong \widetilde{\mathrm{H}}_{n-2}(\mathcal{BA}_n'; \mathbb{Z})$.*

This is our first invocation of Theorem C, which states that $\mathcal{BA}_n$ is CM of dimension n. In particular,

$$\mathrm{H}_{n-1}(\mathcal{BA}_n; \mathbb{Z}) = \mathrm{H}_{n-2}(\mathcal{BA}_n; \mathbb{Z}) = 0.$$

From the long exact sequence for relative homology, we obtain the desired isomorphism.

Step 3. *We have* $\widetilde{\mathrm{H}}_{n-2}(\mathcal{BA}'_n; \mathbb{Z}) \cong \mathrm{St}_n$.

Let $\mathbb{T}_n$ denote the poset of proper nonzero direct summands of $\mathbb{Z}^n$ under inclusion. A direct summand of $\mathbb{Z}^n$ is uniquely determined by the subspace of $\mathbb{Q}^n$ it spans and vice versa. This means that $\mathbb{T}_n$ is isomorphic to the poset of proper nonzero $\mathbb{Q}$-subspaces of $\mathbb{Z}^n$, so its geometric realization $|\mathbb{T}_n|$ can be identified with the Tits building $\mathcal{T}_n$ from the introduction, whose homology is the Steinberg module. In other words,

$$\widetilde{\mathrm{H}}_{n-2}(\mathbb{T}_n; \mathbb{Z}) \cong \widetilde{\mathrm{H}}_{n-2}(\mathcal{T}_n; \mathbb{Z}) = \mathrm{St}_n \,.$$

Recall that $\mathcal{P}(\mathcal{BA}'_n)$ is the poset of simplices of $\mathcal{BA}'_n$. There is a poset map

$$F \colon \mathcal{P}(\mathcal{BA}'_n) \to \mathbb{T}_n$$

defined by

$$F(\{v_1^\pm, \ldots, v_k^\pm\}) = \mathrm{span}_{\mathbb{Z}}(v_1^\pm, \ldots, v_k^\pm).$$

We remark that F can only be defined on $\mathcal{BA}'_n$ and not on $\mathcal{BA}_n$ since $\mathbb{T}_n$ consists of *proper* direct summands. To prove the isomorphism claimed in Step 3, we will use Proposition 2.3 (with $d = n - 2$ and $m = 1$) to prove that F induces an isomorphism $F_* \colon \widetilde{\mathrm{H}}_{n-2}(\mathcal{BA}'_n) \xrightarrow{\cong} \widetilde{\mathrm{H}}_{n-2}(\mathbb{T}_n; \mathbb{Z})$.

We need to verify that F satisfies the conditions of Proposition 2.3. The Solomon-Tits theorem states that $\mathbb{T}_n$ is CM of dimension $n - 2$ (see e.g. [Br2, Remark IV.4.3] or [Q, Example 8.2]). It remains to verify the second condition of Proposition 2.3 for each direct summand $V \in \mathbb{T}_n$.

If $\mathrm{rank}(V) = 1$, the fiber $F_{\leq V}$ is easy to describe: a rank-1 direct summand V contains only one line, so $F_{\leq V}$ is a single point. In particular, the hypothesis $\widetilde{\mathrm{H}}_q(\mathbb{F}_{\leq V}) = 0$ holds for *all* q in this case.

Now consider a direct summand V with $\mathrm{rank}(V) = \ell \geq 2$. A partial augmented frame $\{v_1^\pm, \ldots, v_k^\pm\}$ lies in $F_{\leq V}$ if and only if $\mathrm{span}_{\mathbb{Z}}(v_1^\pm, \ldots, v_k^\pm)$ is contained in V. Let $U = \mathrm{span}_{\mathbb{Z}}(v_1^\pm, \ldots, v_k^\pm)$ be the direct summand of $\mathbb{Z}^n$ spanned by these lines. We have the following well-known but important lemma (which holds over any Dedekind domain): if U and V are direct summands of $\mathbb{Z}^n$ and $U \subset V$, then U is a direct summand of V. What this means is that $F_{\leq V}$ consists of those collections of lines that form a partial augmented frame for V. Thus by choosing an isomorphism $V \cong \mathbb{Z}^\ell$, we obtain an identification

$$F_{\leq V} \cong \mathcal{BA}_\ell.$$

We now invoke Theorem C for the second time: it states that $\mathcal{BA}_\ell$ is ℓ-spherical, so $F_{\leq V}$ is $\mathrm{rank}(V)$-spherical. Since $\mathrm{ht}(V) = \mathrm{rank}(V) - 1$, this verifies the desired hypothesis for $F_{\leq V}$.

3 The vanishing theorem

In this section, we use Theorem B to prove Theorem A. The actual proof is contained in §3.2. This is preceded by §3.1, which contains some preliminary lemmas.

3.1 Ingredients of the vanishing theorem

This section contains two ingredients needed for the proof of Theorem A. The first is as follows.

Lemma 3.1. *Let G be a group and let M and N be G-modules. Assume that N is a vector space over a field of characteristic 0. Also, let*

$$\cdots \longrightarrow F_2 \longrightarrow F_1 \longrightarrow F_0 \longrightarrow M \longrightarrow 0$$

be a resolution of M by flat G-modules. Then the homology of the chain complex

$$\cdots \longrightarrow F_2 \otimes_G N \longrightarrow F_1 \otimes_G N \longrightarrow F_0 \otimes_G N \longrightarrow 0$$

equals $\mathrm{H}_*(G; M \otimes N)$, where G acts diagonally on $M \otimes N$.

Proof. This combines the statements of [Br1, Proposition III.2.1] and [Br1, Proposition III.2.2]. $\square$

To make Lemma 3.1 useful, we need a simple way of recognizing flat G-modules. Our second lemma is such a criterion.

Lemma 3.2. *Let G be a group, let X be a simplicial complex on which G acts simplicially, and let Y be a subcomplex of X which is preserved by the G-action. For some $n \geq 0$, assume that the setwise stabilizer subgroup G_σ is finite for every n-simplex σ of X that is not contained in Y. Then the G-module $C_n(X, Y; \mathbb{Q})$ of relative simplicial n-chains is flat.*

Proof. For an oriented n-simplex σ of X that is not contained in Y, let $[\sigma]$ be the associated basis element of $C_n(X, Y; \mathbb{Q})$. Define $M_\sigma \subset C_n(X, Y; \mathbb{Q})$ to be the span of $\{[g(\sigma)] \mid g \in G\}$, so M_σ is a G-submodule of $C_n(X, Y; \mathbb{Q})$. As in the statement of the lemma, G_σ will denote the setwise stabilizer subgroup of σ. This subgroup may reverse the orientation of σ. Let $\mathbb{Q}_\sigma$ be the G_σ-module whose underlying vector space is $\mathbb{Q}$ but where an element of G_σ acts by ± 1 depending on whether or not it reverses the orientation of σ. We then have that

$$M_\sigma \cong \mathrm{Ind}_{G_\sigma}^G \mathbb{Q}_\sigma.$$

Since $\mathbb{Q}_\sigma$ is an irreducible representation of the finite group G_σ, it is a direct summand of $\mathbb{Q}[G_\sigma]$. It follows that M_σ is a direct summand of

$$\mathrm{Ind}_{G_\sigma}^G \mathbb{Q}[G_\sigma] \cong \mathbb{Q}[G].$$

Since $\mathbb{Q}[G]$ is a localization of the free G-module $\mathbb{Z}[G]$, it is a flat G-module. We deduce that M_σ is a flat G-module. Choosing representatives for the G-orbits of n-simplices of X not lying in Y determines an isomorphism

$$C_n(X, Y; \mathbb{Q}) \cong \bigoplus_{\sigma \in (X^{(n)} - Y^{(n)})/G} M_\sigma,$$

so $C_n(X, Y; \mathbb{Q})$ is a flat G-module, as desired. $\square$

3.2 The proof of Theorem A

We now prove Theorem A. We begin by recalling its statement. Fixing some $\lambda \in \mathbb{Z}^n$ and some $n \geq 3 + \|\lambda\|$, this theorem asserts that

$$\mathrm{H}^{\binom{n}{2}-1}(\mathrm{SL}_n \mathbb{Z}; V_\lambda) = \mathrm{H}^{\binom{n}{2}-1}(\mathrm{GL}_n \mathbb{Z}; V_\lambda) = 0.$$

Since V_λ is a vector space over a field of characteristic 0, the basic properties of the transfer map (see [Br1, Chapter III.9]) show that the vector space $\mathrm{H}^{\binom{n}{2}-1}(\mathrm{GL}_n\,\mathbb{Z}; V_\lambda)$ is a subspace of $\mathrm{H}^{\binom{n}{2}-1}(\mathrm{SL}_n\,\mathbb{Z}; V_\lambda)$, so it is enough to deal with $\mathrm{SL}_n\,\mathbb{Z}$. As we discussed in the introduction, Borel–Serre [BoSe, Eq. (1)] proved that there is an isomorphism

$$\mathrm{H}^{\binom{n}{2}-k}(\mathrm{SL}_n\,\mathbb{Z}; V_\lambda) \cong \mathrm{H}_k(\mathrm{SL}_n\,\mathbb{Z}; \mathrm{St}_n \otimes V_\lambda) \qquad (k \geq 0)$$

Let $\mathcal{BA}'_n$ be the subcomplex of $\mathcal{BA}_n$ introduced in §2.2. Define $\mathcal{I}_0^{\mathbb{Q}} = C_{n-1}(\mathcal{BA}_n, \mathcal{BA}'_n; \mathbb{Q})$ and $\mathcal{I}_1^{\mathbb{Q}} = C_n(\mathcal{BA}_n, \mathcal{BA}'_n; \mathbb{Q})$, and set $\mathrm{St}_n^{\mathbb{Q}} = \mathrm{St}_n \otimes \mathbb{Q}$. Since $\mathrm{St}_n \otimes V_\lambda = \mathrm{St}_n^{\mathbb{Q}} \otimes V_\lambda$, our goal is to show that

$$\mathrm{H}_1(\mathrm{SL}_n\,\mathbb{Z}; \mathrm{St}_n^{\mathbb{Q}} \otimes V_\lambda) = 0. \tag{2}$$

Since $\mathbb{Q}$ is a flat $\mathbb{Z}$-module, it follows from the proof of Theorem B in §2.2 that there is an exact sequence

$$\mathcal{I}_1^{\mathbb{Q}} \longrightarrow \mathcal{I}_0^{\mathbb{Q}} \longrightarrow \mathrm{St}_n^{\mathbb{Q}} \longrightarrow 0.$$

The $(n-1)$-simplices of $\mathcal{BA}_n$ that do not lie in $\mathcal{BA}'_n$ are the frames $\{v_1^\pm, \ldots, v_n^\pm\}$ for $\mathbb{Z}^n$. The $\mathrm{GL}_n\,\mathbb{Z}$-stabilizer of such a frame is a finite group isomorphic to $S_n^\pm$, the $2^n \cdot n!$-element group of signed permutation matrices. The $\mathrm{SL}_n\,\mathbb{Z}$-stabilizer of each frame is thus finite, so Lemma 3.2 shows that $\mathcal{I}_0^{\mathbb{Q}}$ is a flat $\mathrm{SL}_n\,\mathbb{Z}$-module. Similarly, the n-simplices of $\mathcal{BA}_n$ that do not lie in $\mathcal{BA}'_n$ are the augmented frames $\{v_0^\pm, \ldots, v_n^\pm\}$ for $\mathbb{Z}^n$. The $\mathrm{GL}_n\,\mathbb{Z}$-stabilizer of an augmented frame is isomorphic to $D_6 \times S_{n-2}^\pm$, where D_6 is the dihedral group of order 12, so the $\mathrm{SL}_n\,\mathbb{Z}$-stabilizer of an augmented frame is finite. By Lemma 3.2, $\mathcal{I}_1^{\mathbb{Q}}$ is a flat $\mathrm{SL}_n\,\mathbb{Z}$-module as well.

We may therefore extend this exact sequence to a flat resolution of the $\mathrm{SL}_n\,\mathbb{Z}$-module $\mathrm{St}_n^{\mathbb{Q}}$:

$$\cdots \longrightarrow F_3 \longrightarrow F_2 \longrightarrow \mathcal{I}_1^{\mathbb{Q}} \longrightarrow \mathcal{I}_0^{\mathbb{Q}} \longrightarrow \mathrm{St}_n^{\mathbb{Q}} \longrightarrow 0. \tag{3}$$

Lemma 3.1 says that $\mathrm{H}_*(\mathrm{SL}_n\,\mathbb{Z}; \mathrm{St}_n^{\mathbb{Q}} \otimes V_\lambda)$ is computed by the homology of the chain complex

$$\cdots \longrightarrow F_3 \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V_\lambda \longrightarrow F_2 \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V_\lambda \longrightarrow \mathcal{I}_1^{\mathbb{Q}} \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V_\lambda \longrightarrow \mathcal{I}_0^{\mathbb{Q}} \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V_\lambda \longrightarrow 0.$$

To prove (2), it is therefore enough to show that $\mathcal{I}_1^{\mathbb{Q}} \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V_\lambda = 0$ under our assumption that $n \geq 3 + \|\lambda\|$. We remark that in our earlier paper [CFP2] with Benson Farb, we used a similar argument to show that $\mathcal{I}_0^{\mathbb{Q}} \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V_\lambda = 0$ for $n \geq 2 + \|\lambda\|$; see [CFP2, Theorem C], which shows the vanishing of $\mathrm{H}^{\binom{n}{2}}(\mathrm{SL}_n\,\mathbb{Z}; V_\lambda)$ and applies also to $\mathrm{SL}_n\,\mathcal{O}_K$ for many number rings $\mathcal{O}_K$.

Define the partition $\lambda' = (\lambda'_1, \ldots, \lambda'_{n-1}, 0)$ via the formula $\lambda'_i = \lambda_i - \lambda_n$; observe that $\|\lambda'\| = \|\lambda\|$. As $\mathrm{GL}_n\,\mathbb{Q}$-representations, we have $V_\lambda \cong V_{\lambda'} \otimes \det^{\otimes \lambda_n}$, so as a representation of $\mathrm{SL}_n\,\mathbb{Q}$ or $\mathrm{SL}_n\,\mathbb{Z}$, the representation V_λ is isomorphic to $V_{\lambda'}$. Let $V := V_{(1)}$ denote the standard $\mathrm{SL}_n\,\mathbb{Z}$-representation on $\mathbb{Q}^n$. Using Schur–Weyl duality, we can embed $V_{\lambda'}$ as a direct summand of $V^{\otimes k}$ with $k = \|\lambda'\| = \|\lambda\|$. It thus suffices to show that $\mathcal{I}_1^{\mathbb{Q}} \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V^{\otimes k} = 0$ when $n \geq 3 + k$.

Fix an augmented frame $\sigma = \{v_0^\pm, \ldots, v_n^\pm\}$ for $\mathbb{Z}^n$, and choose representatives so that $\{v_1, \ldots, v_n\}$ is a basis for $\mathbb{Z}^n$ and $v_0 = v_1 + v_2$. Orienting σ using the ordering on the v_i determines a generator $[\sigma]$ of $\mathcal{I}_1^{\mathbb{Q}} \cong C_n(\mathcal{BA}_n, \mathcal{BA}'_n; \mathbb{Q})$. Moreover, fix arbitrary indices $i_1, \ldots, i_k \in \{1, \ldots, n\}$ and consider the element $w = v_{i_1} \otimes \cdots \otimes v_{i_k} \in V^{\otimes k}$.

We now prove that the image of $[\sigma] \otimes w$ in $\mathcal{I}_1^{\mathbb{Q}} \otimes_{\mathrm{SL}_n\,\mathbb{Z}} V^{\otimes k}$ is 0. We do this by constructing an element φ in the stabilizer of the augmented frame σ that satisfies $\varphi([\sigma]) = -[\sigma]$ and $\varphi(w) = w$.

Since $n \geq 3 + k$, we can find some $3 \leq j \leq n$ such that $j \notin \{i_1, \ldots, i_k\}$. We consider two cases separately.

- First, if we can find a second index $3 \leq j' \leq n$ with $j \neq j'$ and $j' \notin \{i_1, \ldots, i_k\}$, we define $\varphi \in \mathrm{SL}_n \mathbb{Z}$ by

$$
\varphi(v_i) = \begin{cases} -v_{j'} & \text{if } i = j, \\ v_j & \text{if } i = j', \\ v_i & \text{otherwise} \end{cases} \qquad (1 \leq i \leq n).
$$

  Note that $\varphi(v_0) = \varphi(v_1 + v_2) = v_1 + v_2 = v_0$, so the element φ preserves the augmented frame σ. Since φ exchanges the lines $v_j^{\pm}$ and $v_{j'}^{\pm}$, it reverses the orientation of the corresponding simplex, so $\varphi([\sigma]) = -[\sigma]$. By construction φ fixes all the vectors $v_{i_1}, \ldots, v_{i_k}$, so $\varphi(w) = w$.

- The second case is that no such j' exists. Neither 1 nor 2 can belong to $\{i_1, \ldots, i_k\}$ since $n \geq 3 + k$. In this case we define $\varphi \in \mathrm{SL}_n \mathbb{Z}$ via the formula

$$
\varphi(v_i) = \begin{cases} v_2 & \text{if } i = 1, \\ v_1 & \text{if } i = 2, \\ -v_j & \text{if } i = j, \\ v_i & \text{otherwise} \end{cases} \qquad (1 \leq i \leq n).
$$

  Note that $\varphi(v_0) = \varphi(v_1 + v_2) = v_2 + v_1 = v_0$, so again φ preserves the augmented frame σ. Since φ exchanges the lines $v_1^{\pm}$ and $v_2^{\pm}$, we have $\varphi([\sigma]) = -[\sigma]$. Since neither 1 nor 2 nor j belongs to $\{i_1, \ldots, i_k\}$, the vectors $v_{i_1}, \ldots, v_{i_k}$ are all fixed by φ, so $\varphi(w) = w$.

In both cases, our chosen element satisfies $\varphi([\sigma]) = -[\sigma]$ and $\varphi(w) = w$. It follows that the images of $[\sigma] \otimes w$ and $-[\sigma] \otimes w = \varphi([\sigma] \otimes w)$ coincide in $\mathcal{I}_1^{\mathbb{Q}} \otimes_{\mathrm{SL}_n \mathbb{Z}} V^{\otimes k}$, and thus must be 0.

Since $\{v_1, \ldots, v_k\}$ is a basis for $\mathbb{Z}^n$, the tensors $v_{i_1} \otimes \cdots \otimes v_{i_k}$ constitute a basis for $V^{\otimes k}$. In other words, w was an arbitrary basis element of $V^{\otimes k}$, so the vanishing of $[\sigma] \otimes w$ implies that the image of $[\sigma] \otimes V^{\otimes k}$ in $\mathcal{I}_1^{\mathbb{Q}} \otimes_{\mathrm{SL}_n \mathbb{Z}} V^{\otimes k}$ vanishes. Since $[\sigma]$ was an arbitrary generator of $\mathcal{I}_1^{\mathbb{Q}}$, this implies that $\mathcal{I}_1^{\mathbb{Q}} \otimes_{\mathrm{SL}_n \mathbb{Z}} V^{\otimes k} = 0$, as desired. This completes the proof of Theorem A.

4 The complex of partial augmented frames is CM

The remainder of the paper is occupied with the proof of Theorem C, which asserts that the n-dimensional complex $\mathcal{BA}_n$ is Cohen–Macaulay (CM) of dimension n.

4.1 Warmup: The complex of partial frames is CM

Recall from the introduction that $\mathcal{B}_n$ is the complex of partial frames of $\mathbb{Z}^n$. In this section, we will prove that $\mathcal{B}_n$ is CM of dimension $(n-1)$. This theorem is similar to a result of Maazen [Maa] and we could deduce it from his work, but we include a proof since it provides a simpler venue to preview the ideas that we will use in our proof of Theorem C. Moreover, we will use both this result and the details of its proof in multiple places during the proof of Theorem C.

During our proof, we will need to understand the links of various simplices of $\mathcal{B}_n$, so we make the following definition. Throughout this section, $\{e_1, \ldots, e_p\}$ will denote the standard basis for $\mathbb{Z}^p$; the context will indicate what value of p we are using at any particular point.

Definition 4.1. For $n \geq 0$ and $m \geq 0$, let $\mathcal{B}_n^m$ be the subcomplex $\mathrm{Link}_{\mathcal{B}_{m+n}}(\{e_1^{\pm}, \ldots, e_m^{\pm}\})$ of $\mathcal{B}_{m+n}$.

The main result of this section is the following theorem, which as we said above is closely related to a theorem of Maazen [Maa]. Of course, $\mathcal{B}_n^0$ is equal to $\mathcal{B}_n$, so this theorem proves that $\mathcal{B}_n$ is CM of dimension $n - 1$, as claimed.

Theorem 4.2. *For all $n \geq 0$ and $m \geq 0$, the complex $\mathcal{B}_n^m$ is CM of dimension $n - 1$.*

We preface the proof of Theorem 4.2 with two lemmas. Analogues of these two lemmas will be at the heart of our proof of the more difficult Theorem C (and the second lemma here will also be used directly during that proof).

Lemma 4.3. *Consider $n \geq 0$ and $m \geq 0$. For some $1 \leq k \leq n$, let σ be a $(k-1)$-simplex of $\mathcal{B}_n^m$. Then the complex $\mathrm{Link}_{\mathcal{B}_n^m}(\sigma)$ is isomorphic to $\mathcal{B}_{n-k}^{m+k}$.*

Proof. Write $\sigma = \{v_1^{\pm}, \ldots, v_k^{\pm}\}$, so $\{e_1, \ldots, e_m, v_1, \ldots, v_k\}$ is a basis for a direct summand of $\mathbb{Z}^{m+n}$. Extend this to a basis $\{e_1, \ldots, e_m, v_1, \ldots, v_n\}$ for $\mathbb{Z}^{m+n}$. Define $\varphi \in \mathrm{GL}_{m+n}(\mathbb{Z})$ by the formulas $\varphi(e_i) = e_i$ for $1 \leq i \leq m$ and $\varphi(v_j) = e_{m+j}$ for $1 \leq j \leq n$. Then φ induces an automorphism of $\mathcal{B}_n^m$ that takes $\mathrm{Link}_{\mathcal{B}_n^m}(\sigma)$ to $\mathcal{B}_{n-k}^{m+k} \subset \mathcal{B}_n^m$. $\square$

Definition 4.4. Consider $n \geq 0$ and $m \geq 0$. Assume that some linear map $F: \mathbb{Z}^{m+n} \to \mathbb{Z}$ has been fixed. Given a subcomplex X of $\mathcal{B}_n^m$ and given $N > 0$, we define $X^{<N}$ to be the full subcomplex of X spanned by the set of vertices $v^{\pm}$ of X satisfying $|F(v)| < N$. This condition is well-defined since $|F(v)| = |F(-v)|$.

Lemma 4.5. *Consider $n \geq 0$ and $m \geq 0$. Let $F: \mathbb{Z}^{m+n} \to \mathbb{Z}$ be a fixed linear map and let $N > 0$. Let σ be a simplex of $\mathcal{B}_n^m$ such that some vertex $w^{\pm}$ of σ satisfies $F(w) = N$. Then there exists a simplicial retraction $\pi: \mathrm{Link}_{\mathcal{B}_n^m}(\sigma) \twoheadrightarrow \mathrm{Link}_{\mathcal{B}_n^m}(\sigma)^{<N}$.*

Proof. Define $X = \mathrm{Link}_{\mathcal{B}_n^m}(\sigma)$ and write $\sigma = \{w_1^{\pm}, \ldots, w_p^{\pm}\}$ with $w_1 = w$. Our goal is to construct a simplicial retraction $X \twoheadrightarrow X^{<N}$. Say that $v \in \mathbb{Z}^{m+n}$ is *F-nonnegative* if $F(v) \geq 0$. We begin by defining a map $\widehat{\pi}: X^{(0)} \to (X^{<N})^{(0)}$ on 0-simplices as follows.

- Consider a vertex $v^{\pm}$ of X. Replacing v with $-v$ if necessary, we can assume that v is F-nonnegative. Define $q_v \in \mathbb{N}$ to be the result $\lfloor \frac{F(v)}{N} \rfloor$ of dividing $F(v)$ by N, so $0 \leq F(v) - q_v N < N$. We then set

$$\widehat{\pi}(v^{\pm}) = (v - q_v w)^{\pm}. \tag{4}$$

  This is well-defined; the only possible ambiguity occurs when $F(v) = 0$ and hence both v and $-v$ are F-nonnegative, but in that case we have $q_v = q_{-v} = 0$ so $\widehat{\pi}(v^{\pm}) = v^{\pm}$ no matter what choice we make.

By definition, $\widehat{\pi}(v^{\pm}) = v^{\pm}$ if $v^{\pm}$ is a vertex of $X^{<N}$, and similarly $\widehat{\pi}(v^{\pm}) \in X^{<N}$ for any vertex $v^{\pm}$ of X. To complete the proof, we must prove that $\widehat{\pi}$ extends over the higher-dimensional simplices of X. Consider a $(k-1)$-simplex $\{v_1^{\pm}, \ldots, v_k^{\pm}\}$ of X, so $\{e_1, \ldots, e_m, w_1, \ldots, w_p, v_1, \ldots, v_k\}$ is a basis for a rank-$(m + p + k)$ direct summand U of $\mathbb{Z}^n$. Replace the v_i by $-v_i$ if necessary to ensure that the v_i are F-nonnegative and set $v_i' = v_i - q_{v_i} w_1$, so $\widehat{\pi}(v_i^{\pm}) = (v_i')^{\pm}$. Since each v_i' is obtained by adding some multiple of w_1 to v_i, the set $\{e_1, \ldots, e_m, w_1, \ldots, w_p, v_1', \ldots, v_k'\}$ is also a basis for U. We conclude that $\{\widehat{\pi}(v_1^{\pm}), \ldots, \widehat{\pi}(v_k^{\pm})\}$ is a $(k-1)$-simplex of $X^{<N}$, as desired. $\square$

Remark 4.6. If a partial frame σ of $\mathbb{Z}^\ell$ is contained in a summand V of $\mathbb{Z}^\ell$, then in fact σ is a partial frame of V. Although this fact may seem obvious, it need not hold over other rings and its failure can lead to great difficulty. For example, over the ring $A = \mathbb{R}[x, y, z]/(x^2 + y^2 + z^2 - 1)$, the vector $v = xe_1 + ye_2 + ze_3$ is part of a basis for A^4 and has e_4-coordinate 0, but v is *not* part of any basis containing the vector e_4. Nevertheless, over $\mathbb{Z}$ the claim follows from the fact that if a summand U of $\mathbb{Z}^\ell$ is contained in another summand V of $\mathbb{Z}^\ell$, then U is a summand of V; this property holds not only for $\mathbb{Z}$ but for any Dedekind domain. For the same reason, a partial augmented frame of $\mathbb{Z}^\ell$ that is contained in a summand V of $\mathbb{Z}^\ell$ is in fact a partial augmented frame of V; this will be used in the next section in the proof of Theorem C.

We now come to the proof of Theorem 4.2. This proof could be written in the language of combinatorial Morse theory without great difficulty, but it would be much more awkward to express our later proof of Theorem C in this language (as we illustrate afterwards in Remark 4.18). Since our goal is to motivate the proof of Theorem C, we follow its structure here.

Proof of Theorem 4.2. We prove the theorem by induction on n. For the base case $n = 0$, we must prove for all $m \geq 0$ that $\mathcal{B}_0^m$ is CM of dimension -1, i.e. that the simplicial complex $\mathcal{B}_0^m$ is empty. Since $\{e_1^\pm, \ldots, e_m^\pm\}$ is already a frame for $\mathbb{Z}^m$, it is a maximal simplex of $\mathcal{B}_m$. Therefore its link $\mathcal{B}_0^m$ is empty, as desired.

Now fix $n > 0$ and $m \geq 0$ and assume that $\mathcal{B}_{n'}^{m'}$ is CM of dimension $n' - 1$ for all $n' < n$ and all $m' \geq 0$. Since every frame for $\mathbb{Z}^{m+n}$ consists of $m + n$ lines, the complex $\mathcal{B}_n^m$ is $(n - 1)$-dimensional. Lemma 4.3 and our induction hypothesis implies that for all $(k-1)$-simplices σ of $\mathcal{B}_n^m$ with $1 \leq k \leq n$, the complex $\mathrm{Link}_{\mathcal{B}_n^m}(\sigma)$ is CM of dimension $n - k$. All that remains to show is that $\mathcal{B}_n^m$ is $(n - 2)$-connected.

Fix $0 \leq p \leq n - 2$, let S^p be a combinatorial triangulation of a p-sphere, and let $\varphi \colon S^p \to \mathcal{B}_n^m$ be a simplicial map. Our goal is to show that φ can be homotoped to a constant map. Let $F \colon \mathbb{Z}^{m+n} \twoheadrightarrow \mathbb{Z}$ be the linear map taking $v \in \mathbb{Z}^{m+n}$ to the e_{m+n}-coordinate of v. For a vertex $v^\pm$ of $\mathcal{B}_n^m$, define $r(v^\pm) = |F(v)|$; this is well-defined since $|F(v)| = |F(-v)|$. We then define

$$R(\varphi) = \max\{r(\varphi(x)) \mid x \text{ a vertex of } S^p\}.$$

This will be our measure of complexity for φ.

If $R(\varphi) = 0$, then every simplex σ of $\varphi(S^p)$ is contained in the summand $\ker F$ of $\mathbb{Z}^{m+n}$. In particular, $\{e_1^\pm, \ldots, e_m^\pm\} * \sigma$ is a partial frame contained in $\ker F$; by Remark 4.6, it is in fact a partial frame for $\ker F$, so it can be extended to a partial frame for $\mathbb{Z}^{n+m}$ by adding the line $e_{m+n}^\pm$. In other words, the entire image $\varphi(S^p)$ is contained in the star (indeed, in the link) of $e_{m+n}^\pm$. We conclude that when $R(\varphi) = 0$, the desired null-homotopy is obtained by homotoping φ to the constant map at the vertex $e_{m+n}^\pm$.

We can therefore assume that $R(\varphi) = R > 0$; we want to homotope φ so as to reduce $R(\varphi)$. Consider the following condition on a simplex σ of S^p:

$$r(\varphi(x)) = R \text{ for all vertices } x \text{ of } \sigma. \tag{5}$$

Since $R(\varphi) = R$, there must be some simplex σ of S^p satisfying (5). We can therefore choose a simplex σ of S^p satisfying (5) whose dimension k is maximal among those satisfying (5). This maximality implies that φ takes $\mathrm{Link}_{S^p}(\sigma)$ to $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$.

12

Let ℓ be the dimension of the simplex $\varphi(\sigma)$; we certainly have $\ell \leq k$, but we might have $\ell < k$ if φ restricted to σ is not injective. Combining Lemma 4.3 with our induction hypothesis, we see that $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))$ is CM of dimension $(n - \ell - 2)$, and in particular is $(n - \ell - 3)$-connected. This retracts to $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$ by Lemma 4.5, so its retract $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$ is also $(n - \ell - 3)$-connected.

By the definition of a combinatorial triangulation, the link $\mathrm{Link}_{S^p}(\sigma)$ is a combinatorial $(p - k - 1)$-sphere. Since $p \leq n - 2$ and $\ell \leq k$, we have $p - k - 1 \leq n - \ell - 3$, so $\varphi|_{\mathrm{Link}_{S^p}(\sigma)}$ is null-homotopic via a homotopy inside $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$. Using Zeeman's relative simplicial approximation theorem [Z], we conclude that there exists a combinatorial $(p - k)$-ball B with $\partial B \cong \mathrm{Link}_{S^p}(\sigma)$ and a simplicial map $\psi \colon B \to \mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$ such that $\psi|_{\partial B} = \varphi|_{\mathrm{Link}_{S^p}(\sigma)}$.

The map ψ extends to the $(p + 1)$-ball $\sigma * B$ as $(\varphi|_\sigma) * \psi \colon \sigma * B \to \mathcal{B}_n^m$. The boundary of $\sigma * B$ is the union of the p-ball $\sigma * (\partial B) = \mathrm{Star}_{S^p}(\sigma)$, on which $\varphi|_\sigma * \psi = \varphi|_{\mathrm{Star}_{S^p}(\sigma)}$, and the p-ball $(\partial \sigma) * B$. We can thus homotope φ across this $(p + 1)$-ball to replace $\varphi|_{\mathrm{Star}_{S^p}(\sigma)}$ with $\varphi|_{\partial \sigma} * \psi \colon (\partial \sigma) * B \to \mathcal{B}_n^m$.

The key property of this modification is that it eliminates the simplex σ and does not add any other simplices satisfying (5). Indeed, every new simplex is the join of a simplex in $\partial \sigma$ with a nonempty simplex in B; since $\psi(B)$ is contained in $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$, such a simplex has at least one vertex with $r(\varphi(x)) < R$, so it will not satisfy (5). Repeating this process, we can homotope φ to eliminate *all* simplices satisfying (5); in other words, we can homotope φ so that $R(\varphi) < R$.

By induction, we can homotope φ so that $R(\varphi) = 0$. At this point, as explained above, φ can be directly contracted to a constant map, so this concludes the proof that $\mathcal{B}_n^m$ is $(n - 2)$-connected. $\qquad\square$

4.2 The complex $\mathcal{BA}_n^m$

We now turn to the proof of Theorem C, which asserts that the complex $\mathcal{BA}_n$ is CM of dimension n. Just as for $\mathcal{B}_n$, we will need to understand links of simplices in $\mathcal{BA}_n$. However, for technical reasons the heart of our argument will deal not with the entire link, but rather with the following subcomplex of the link. Recall that $\{e_1, \ldots, e_p\}$ denotes the standard basis for $\mathbb{Z}^p$, where $p \geq 1$ is determined by context.

Definition 4.7. For $n \geq 1$ and $m \geq 0$ with $m + n \geq 2$, define $\mathcal{BA}_n^m$ to be the full subcomplex of $\mathrm{Link}_{\mathcal{BA}_{m+n}}(\{e_1^\pm, \ldots, e_m^\pm\})$ spanned by vertices $v^\pm$ of $\mathrm{Link}_{\mathcal{BA}_{m+n}}(\{e_1^\pm, \ldots, e_m^\pm\})$ such that $v \notin \mathrm{span}_{\mathbb{Z}}(e_1, \ldots, e_m) \subset \mathbb{Z}^{m+n}$.

For example, even though $\{e_1^\pm, e_2^\pm, (e_1 + e_2)^\pm\}$ is a simplex of $\mathcal{BA}_{n+2}$, the vertex $(e_1 + e_2)^\pm$ is excluded from $\mathcal{BA}_n^2$. Our main theorem is then as follows. It reduces to Theorem C when $m = 0$.

Theorem C'. For $n \geq 1$ and $m \geq 0$ with $m + n \geq 2$, the complex $\mathcal{BA}_n^m$ is CM of dimension n.

Remark 4.8. We have intentionally refrained from defining $\mathcal{BA}_n^m$ in the case when $n = 0$ or the case when $m + n < 2$. The reason is that $\mathcal{BA}_n^m$ would be degenerate in these cases; not only would Theorem C' be false in these cases, $\mathcal{BA}_n^m$ would not even be n-dimensional.

We will prove Theorem C' in §4.5. This is preceded by §4.3, which describes the links in $\mathcal{BA}_n^m$ (or certain subcomplexes of the links) and establishes the base case for our induction,

and by §4.4, which constructs certain retractions on links in $\mathcal{BA}_n^m$. Before we start with all of this, we close this section by introducing some terminology for simplices of $\mathcal{BA}_n^m$.

Definition 4.9. Fix $n \geq 1$ and $m \geq 0$ with $m + n \geq 2$. We divide the simplices of $\mathcal{BA}_n^m$ into three mutually exclusive types.
- A *standard simplex* is a simplex $\{v_1^\pm, \ldots, v_p^\pm\}$ such that $\{e_1^\pm, \ldots, e_m^\pm, v_1^\pm, \ldots, v_p^\pm\}$ is a simplex of $\mathcal{B}_n^m$. In other words, $\{e_1, \ldots, e_m, v_1, \ldots, v_p\}$ is a basis for a direct summand of $\mathbb{Z}^{m+n}$.
- An *internally additive simplex* is a simplex that can be written as $\{v_0^\pm, \ldots, v_p^\pm\}$, where $\{v_1^\pm, \ldots, v_p^\pm\}$ is a standard simplex and $\pm v_0 \pm v_1 \pm v_2 = 0$ for some choice of signs. We will call $\{v_0^\pm, v_1^\pm, v_2^\pm\}$ the *additive core* of our simplex; this subset is well-defined since $\{v_0, v_1, v_2\}$ is the minimal linearly dependent subset of $\{v_0, \ldots, v_p\}$.
- An *externally additive simplex* is a simplex that can be written as $\{v_0^\pm, \ldots, v_p^\pm\}$, where $\{v_1^\pm, \ldots, v_p^\pm\}$ is a standard simplex and $\pm v_0 \pm v_1 \pm e_i = 0$ for some choice of signs and some $1 \leq i \leq m$. We will call $\{v_0^\pm, v_1^\pm\}$ the *additive core* of our simplex; it is well-defined just as for internally additive simplices.

An *additive simplex* is a simplex which is either internally or externally additive.

Remark 4.10. We emphasize that the classification in Definition 4.9 applies to a simplex *as a simplex of* $\mathcal{BA}_n^m$. The same collection of lines might be classified differently as a simplex of $\mathcal{BA}_{n'}^{m'}$. For example, a partial frame that forms an externally additive simplex of $\mathcal{BA}_n^m$ would be a standard simplex when considered as a simplex of $\mathcal{BA}_{n+m}^0$.

4.3 Describing links in $\mathcal{BA}_n^m$

In this section, we describe the links of simplices in $\mathcal{BA}_n^m$, as we did for $\mathcal{B}_n^m$ in Lemma 4.3. To handle the link of a standard simplex, we are forced to deal with a certain subcomplex of the link (just as $\mathcal{BA}_n^m$ is a *subcomplex* of the full link in $\mathcal{BA}_{m+n}$); the reason is that the retraction constructed in Proposition 4.17 below cannot be extended across the entire link.

Definition 4.11. Given a standard simplex $\sigma = \{v_1^\pm, \ldots, v_p^\pm\}$ of $\mathcal{BA}_n^m$, define $\widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(\sigma)$ to be the full subcomplex of $\mathrm{Link}_{\mathcal{BA}_n^m}(\sigma)$ spanned by vertices $v^\pm$ of $\mathrm{Link}_{\mathcal{BA}_n^m}(\sigma)$ such that $v \notin \mathrm{span}_\mathbb{Z}(e_1, \ldots, e_m, v_1, \ldots, v_p) \subset \mathbb{Z}^{m+n}$.

Lemma 4.12. *Consider $n \geq 1$ and $m \geq 0$ with $m + n \geq 2$. For some $1 \leq k \leq n+1$, let σ be a $(k-1)$-simplex of $\mathcal{BA}_n^m$.*
  (a) *If σ is an additive simplex, then $\mathrm{Link}_{\mathcal{BA}_n^m}(\sigma)$ is isomorphic to $\mathcal{B}_{n-k+1}^{m+k-1}$.*
  (b) *If σ is a standard simplex and $k \neq n$, then $\widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(\sigma)$ is isomorphic to $\mathcal{BA}_{n-k}^{m+k}$.*
  (c) *If σ is a standard simplex, define $X = \mathrm{Link}_{\mathcal{BA}_n^m}(\sigma)$ and $\widehat{X} = \widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(\sigma)$. Then for all vertices $v^\pm$ of X that do not lie in $\widehat{X}$, the complex $\mathrm{Link}_X(v^\pm)$ lies in $\widehat{X}$ and is isomorphic to $\mathcal{B}_{n-k}^{m+k}$.*

Proof. Parts (a) and (b) are proved exactly like Lemma 4.3 since by an appropriate automorphism of $\mathcal{BA}_n^m$, we may assume that $\sigma = \{e_{m+1}^\pm, \ldots, e_{m+k-1}^\pm, (e_{m+1} + e_{m+2})^\pm\}$ or $\sigma = \{e_{m+1}^\pm, \ldots, e_{m+k}^\pm\}$, respectively. Part (c) is a consequence of Part (a). $\square$

In the proof of the next proposition, we make use of the following lemma. It is certainly standard, but we could not find a proof in the literature.

14

Lemma 4.13. *Let X be obtained from the simplicial complex Y by coning off the subcomplex Z. If Y is CM of dimension n and Z is CM of dimension $n - 1$, then X is CM of dimension n.*

Proof. Let p be the cone point. Since Z is $(n-2)$-connected, the pair (X, Y) is $(n-1)$-connected, so X is n-spherical. There are four kinds of simplices of X. The first is $\{p\}$, whose link is $\text{Link}_X\{p\} = Z$, which is CM of dimension $(n-1)$ by assumption. The second is $\sigma * \{p\}$ for some simplex σ of Z; its link is $\text{Link}_X(\sigma * \{p\}) = \text{Link}_Z \sigma$, which is CM of the appropriate dimension since Z is CM of dimension $(n-1)$. The third is a simplex σ of Y that does not lie in Z; its link is $\text{Link}_X \sigma = \text{Link}_Y \sigma$, which is CM of the appropriate dimension since Y is CM of dimension n. The fourth is a simplex σ of Z; its link is $\text{Link}_X \sigma = \{p\} * \text{Link}_Z \sigma$, which is CM of the appropriate dimension because Z is CM of dimension $(n-1)$. $\square$

Proposition 4.14. *Fix $n \geq 1$ and $m \geq 0$ such that $m + n \geq 2$. Assume that $\mathcal{BA}_{n'}^{m'}$ is CM of dimension n' for all $1 \leq n' < n$ and $m' \geq 0$ such that $m' + n' = m + n$. Then for every $(k-1)$-simplex σ of $\mathcal{BA}_n^m$, the subcomplex $\text{Link}_{\mathcal{BA}_n^m}(\sigma)$ is CM of dimension $n - k$.*

Proof. If σ is an additive simplex, then Lemma 4.12(a) asserts that $\text{Link}_{\mathcal{BA}_n^m}(\sigma)$ is isomorphic to $\mathcal{B}_{n-k+1}^{m+k-1}$, which Theorem 4.2 says is CM of dimension $n - k$.

If σ is a standard simplex with $k = n$, then we can write $\sigma = \{v_1^\pm, \ldots, v_n^\pm\}$ with $\{e_1^\pm, \ldots, e_m^\pm, v_1^\pm, \ldots, v_n^\pm\}$ a frame for $\mathbb{Z}^{m+n}$. If $m > 0$, set $v_0 = e_1 + v_1$; otherwise, since $m + n \geq 2$ we must have $n \geq 2$ and we can set $v_0 = v_1 + v_2$. In either case, the set $\{e_1^\pm, \ldots, e_m^\pm, v_0^\pm, v_1^\pm, \ldots, v_n^\pm\}$ is an augmented frame for $\mathbb{Z}^{m+n}$, so $v_0^\pm$ is contained in $\text{Link}_{\mathcal{BA}_n^m}(\sigma)$. Therefore $\text{Link}_{\mathcal{BA}_n^m}(\sigma)$ is nonempty, i.e. CM of dimension 0.

Finally, if σ is a standard simplex with $k < n$, define $X = \text{Link}_{\mathcal{BA}_n^m}(\sigma)$ and $\widehat{X} = \widehat{\text{Link}}_{\mathcal{BA}_n^m}(\sigma)$. Since $1 \leq k < n$, Lemma 4.12(b) asserts that $\widehat{X}$ is isomorphic to $\mathcal{BA}_{n-k}^{m+k}$, which by assumption is CM of dimension $n - k$. Let $v^\pm$ be a vertex of X that does not lie in $\widehat{X}$. Lemma 4.12(c) says that $\text{Link}_X(v^\pm)$ lies in $\widehat{X}$ and is isomorphic to $\mathcal{B}_{n-k}^{m+k}$, so adding $v^\pm$ to $\widehat{X}$ has the effect of coning off the subcomplex $\text{Link}_X(v^\pm) \cong \mathcal{B}_{n-k}^{m+k}$. This subcomplex is CM of dimension $n - k - 1$ by Theorem 4.2, so Lemma 4.13 tells us that coning off this subcomplex preserves the property of being CM of dimension $n - k$. Carrying this out for each vertex of X not contained in $\widehat{X}$, we conclude that $X = \text{Link}_{\mathcal{BA}_n^m}(\sigma)$ is CM of dimension $n - k$, as desired. $\square$

4.4 The retraction maps

In this section, we construct two retractions of the links in $\mathcal{BA}_n^m$ (or parts of them), just as we did for $\mathcal{B}_n^m$ in Lemma 4.5. We begin with the following definition.

Definition 4.15. Assume that some linear map $F \colon \mathbb{Z}^{m+n} \to \mathbb{Z}$ has been fixed and let X be a subcomplex of $\mathcal{BA}_n^m$. For $N > 0$, we define $X^{<N}$ to be the full subcomplex of X spanned by the set of vertices $v^\pm$ of X satisfying $|F(v)| < N$. This is well-defined since $|F(v)| = |F(-v)|$.

Our first retraction, for the link of an additive simplex, is straightforward.

Lemma 4.16. *Consider $n \geq 2$ and $m \geq 0$. Let $F \colon \mathbb{Z}^{m+n} \to \mathbb{Z}$ be a fixed linear map. Let σ be an additive simplex of $\mathcal{BA}_n^m$ such that there exists some vertex $w^\pm$ of σ with $F(w) = N > 0$. Then there exists a simplicial retraction $\pi \colon \text{Link}_{\mathcal{BA}_n^m}(\sigma) \twoheadrightarrow \text{Link}_{\mathcal{BA}_n^m}(\sigma)^{<N}$.*

15

Proof. The retraction π is defined on vertices by the same formula (4) as in Lemma 4.5. The fact that σ is an additive simplex ensures that for all simplices τ of $\mathrm{Link}_{\mathcal{BA}_n^m}(\sigma)$, the additive core of $\sigma * \tau$ is disjoint from τ, which implies that there is no difficulty in extending π over τ. $\qquad\square$

Our second retraction will be more difficult to construct because internally additive simplices are extremely constrained. Indeed, if two lines $v_1^\pm$ and $v_2^\pm$ are specified, there are only *two* lines $v_0^\pm$ for which $\{v_0^\pm, v_1^\pm, v_2^\pm\}$ is an internally additive simplex. As a result, if we attempt to define a retraction on the link of a standard simplex by (4) as in Lemma 4.16, the retraction will not extend across all additive simplices.

For example, consider the vectors $v_1 = e_1 + 9e_4$, $v_2 = e_2 + 9e_4$, $v_0 = v_1 + v_2 = e_1 + e_2 + 18e_4$, and $w = e_3 + 10e_4$, so $\{v_0^\pm, v_1^\pm, v_2^\pm\}$ forms an additive simplex of $\mathrm{Link}_{\mathcal{BA}_4}(w^\pm)$. However, if we take $F \colon \mathbb{Z}^n \to \mathbb{Z}$ to be the coefficient of e_4 and define $\widehat{\pi}$ as in (4), then we have $\widehat{\pi}(v_1^\pm) = v_1^\pm$ and $\widehat{\pi}(v_2^\pm) = v_2^\pm$, but $\widehat{\pi}(v_0^\pm) = (v_0 - w)^\pm = (e_1 + e_2 - e_3 + 8e_4)^\pm$. Thus $\{\widehat{\pi}(v_0^\pm), \widehat{\pi}(v_1^\pm), \widehat{\pi}(v_2^\pm)\}$ is not a simplex of $\mathrm{Link}_{\mathcal{BA}_n}(w^\pm)$ at all.

In general this problem seems insuperable. We will solve it only for the link of a single vertex and only after restricting to the subcomplex $\widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(w^\pm)$; even then, to make the retraction well-defined we are forced to subdivide the complex first. This is the content of the following proposition.

Proposition 4.17. *Consider $n \geq 2$ and $m \geq 0$. Let $F \colon \mathbb{Z}^{m+n} \to \mathbb{Z}$ be a fixed linear map such that $F(e_i) = 0$ for $1 \leq i \leq m$. Let $w^\pm$ be a vertex of $\mathcal{BA}_n^m$ such that $F(w) = N > 0$. Then there exists a topological retraction $\pi \colon \widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(w^\pm) \twoheadrightarrow \widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(w^\pm)^{<N}$.*

Proof. Define $X = \widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(w^\pm)$, so our goal is to construct a topological retraction $\pi \colon X \twoheadrightarrow X^{<R}$. We begin by defining a map $\widehat{\pi} \colon X^{(0)} \to (X^{<N})^{(0)}$ on 0-simplices by the same formula (4) as before. To recap, we say that $v \in \mathbb{Z}^{m+n}$ is *F-nonnegative* if $F(v) \geq 0$. For F-nonnegative v we define $q_v = \lfloor \frac{F(v)}{N} \rfloor \in \mathbb{N}$ and set $\widehat{\pi}(v^\pm) = (v - q_v w)^\pm$, so $0 \leq F(\widehat{\pi}(v^\pm)) < N$.

If $v^\pm$ is a vertex of X, then by the definition of $\widehat{\mathrm{Link}}_{\mathcal{BA}_n^m}(w^\pm)$ we know that $v \notin \mathrm{span}_{\mathbb{Z}}(e_1, \ldots, e_m, w)$, so $\{e_1^\pm, \ldots, e_m^\pm, w^\pm, v^\pm\}$ is a partial frame for a rank $m + 2$ summand of $\mathbb{Z}^{m+n}$. Thus $\{e_1^\pm, \ldots, e_m^\pm, w^\pm, \widehat{\pi}(v^\pm)\}$ is a partial frame for the same summand, so $\widehat{\pi}(v^\pm)$ is a vertex of $X^{<N}$; moreover $\widehat{\pi}(v^\pm) = v^\pm$ if $v^\pm$ is a vertex of $X^{<N}$. In other words, $\widehat{\pi}$ is a retraction of the vertices of X onto the vertices of $X^{<N}$.

Unfortunately, the map $\widehat{\pi}$ does not extend to a simplicial map on X, as we discussed above. What we will show instead is that there exists a subdivision Y of X such that $X^{<N}$ is still a subcomplex of Y (so no simplices of $X^{<N}$ are subdivided) and an extension of $\widehat{\pi}$ to Y.

The trouble will occur only on the internally additive simplices. Before we deal with these, we prove that $\widehat{\pi}$ extends over the other simplices of X. We distinguish the standard simplices lying in X into two types:

- A *$w^\pm$-standard simplex* is a simplex σ of X such that $\sigma * \{w^\pm\}$ is a standard simplex of $\mathcal{BA}_n^m$.
- A *$w^\pm$-additive simplex* is a simplex of X that can be written in the form $\{v_0^\pm, \ldots, v_p^\pm\}$ with $\{v_1^\pm, \ldots, v_p^\pm\}$ a standard simplex of $\mathcal{BA}_n^m$ and $\pm v_0 \pm v_1 \pm w = 0$ for some choice of signs.

Claim 1. *The map $\widehat{\pi}$ extends over the $w^\pm$-standard simplices σ of X.*

Proof of claim. This is identical to the proof of the corresponding statement in the proof of Lemma 4.5. $\qquad\square$

Claim 2. *The map $\widehat{\pi}$ extends over the externally additive simplices σ of X.*

Proof of claim. Write $\sigma = \{v_0^\pm, \ldots, v_p^\pm\}$, where each v_i is F-nonnegative, $\{v_1^\pm, \ldots, v_p^\pm\}$ is a $w^\pm$-standard simplex, and $\pm v_0 \pm v_1 \pm e_i = 0$ for some i and some choice of signs. Since $F(e_i) = 0$ and both $F(v_0)$ and $F(v_1)$ are nonnegative, the relation $\pm v_0 \pm v_1 \pm e_i = 0$ implies that $F(v_0) = F(v_1)$. Moreover, possibly replacing v_1 by $-v_1$ if $F(v_1) = 0$, we have $v_0 = v_1 + \varepsilon e_i$ for some $\varepsilon \in \{\pm 1\}$. Since $F(v_0) = F(v_1)$ we have $q_{v_0} = q_{v_1}$, so

$$\{\widehat{\pi}(v_0^\pm), \ldots, \widehat{\pi}(v_p^\pm)\} = \{(v_1 - q_{v_1}w + \varepsilon e_i)^\pm, (v_1 - q_{v_1}w)^\pm, \ldots, (v_p - q_{v_p}w)^\pm\}.$$

This is an externally additive simplex of $X^{<N}$. $\qquad\square$

Claim 3. *The map $\widehat{\pi}$ extends over the $w^\pm$-additive simplices σ of X.*

Proof of claim. Write $\sigma = \{v_0^\pm, \ldots, v_p^\pm\}$, where each v_i is F-nonnegative, $\{v_1^\pm, \ldots, v_p^\pm\}$ is a $w^\pm$-standard simplex, and $\pm v_0 \pm v_1 \pm w = 0$ for some choice of signs. Exchanging v_0 and v_1 if necessary, we can assume that $F(v_0) \geq F(v_1)$.

We first consider the case where $F(v_0) \geq N$. Since $F(w) = N$, in this case the relation $\pm v_0 \pm v_1 \pm w = 0$ implies that $v_0 = v_1 + w$ as long as $F(v_1) > 0$. When $F(v_1) = 0$, it implies only that $v_0 = \pm v_1 + w$, but replacing v_1 by $-v_1$ we can still assume that $v_0 = v_1 + w$. The relation $v_0 = v_1 + w$ implies that $q_{v_0} = q_{v_1} + 1$, so we have

$$\widehat{\pi}(v_0^\pm) = (v_0 - q_{v_0}w)^\pm = ((v_1 + w) - (q_{v_1} + 1)w)^\pm = (v_1 - q_{v_1}w)^\pm = \widehat{\pi}(v_1^\pm).$$

In other words, the $w^\pm$-additive edge $\{v_0^\pm, v_1^\pm\}$ of X is collapsed by $\widehat{\pi}$ to a single vertex of $X^{<N}$. Similarly, the $w^\pm$-additive p-simplex σ is collapsed by $\widehat{\pi}$ to

$$\{\widehat{\pi}(v_0^\pm), \ldots, \widehat{\pi}(v_p^\pm)\} = \{(v_1 - q_{v_1}w)^\pm, \ldots, (v_p - q_{v_p}w)^\pm\},$$

a $w^\pm$-standard $(p-1)$-simplex of $X^{<N}$. We remark that this is the only case in which the dimension of a simplex is decreased by $\widehat{\pi}$.

In the remaining case, we have $0 \leq F(v_1) \leq F(v_0) < N$. Since $F(w) = N$, the relation $\pm v_0 \pm v_1 \pm w = 0$ implies in this case that $w = v_0 + v_1$. Since $F(v_0) < N$ and $F(v_1) < N$, we have $q_{v_0} = q_{v_1} = 0$, so $\widehat{\pi}(v_0^\pm) = v_0^\pm$ and $\widehat{\pi}(v_1^\pm) = v_1^\pm$. Therefore

$$\{\widehat{\pi}(v_0^\pm), \widehat{\pi}(v_1^\pm), \ldots, \widehat{\pi}(v_p^\pm)\} = \{v_0^\pm, v_1^\pm, (v_2 - q_{v_2}w)^\pm, \ldots, (v_p - q_{v_p}w)^\pm\}.$$

is a $w^\pm$-additive p-simplex of $X^{<N}$. $\qquad\square$

The last remaining class of simplices are the internally additive simplices. It will turn out that certain kinds of internally additive simplices will cause trouble. Consider an internally additive simplex σ. Write $\sigma = \{v_0^\pm, \ldots, v_p^\pm\}$ where each v_i is F-nonnegative and $\pm v_0 \pm v_1 \pm v_2 = 0$ for some choice of signs. If all three signs are the same, then we must have $F(v_0) = F(v_1) = F(v_2) = 0$, so we can negate v_0 without changing the fact that v_0 is F-nonnegative. The upshot is that we can assume that the three signs are not all the same. Reordering the v_i if necessary, we can thus assume that $v_0 = v_1 + v_2$. We call σ a *carrying simplex* if

$$\left\lfloor \frac{F(v_0)}{N} \right\rfloor \neq \left\lfloor \frac{F(v_1)}{N} \right\rfloor + \left\lfloor \frac{F(v_2)}{N} \right\rfloor. \tag{6}$$

To check that this is well-defined, observe that for the inequality (6) to hold we must have $F(v_i) > 0$ for $0 \leq i \leq 2$, in which case v_0 is uniquely determined since $F(v_0)$ is the maximum value among $\{F(v_0), F(v_1), F(v_2)\}$. The following claim shows that the carrying simplices are the only possible source of trouble.

Claim 4. *The map $\widehat{\pi}$ extends over the internally additive simplices σ of X that are not carrying simplices.*

Proof of claim. Write $\sigma = \{v_0^{\pm}, \ldots, v_p^{\pm}\}$, where each v_i is F-nonnegative and $v_0 = v_1 + v_2$. Since σ is not a carrying simplex, we have

$$q_{v_0} = \left\lfloor \frac{F(v_0)}{N} \right\rfloor = \left\lfloor \frac{F(v_1)}{N} \right\rfloor + \left\lfloor \frac{F(v_2)}{N} \right\rfloor = q_{v_1} + q_{v_2}.$$

This implies that

$$\widehat{\pi}(v_0^{\pm}) = (v_0 - q_{v_0}w)^{\pm} = ((v_1 + v_2) - (q_{v_1} + q_{v_2})w)^{\pm} = ((v_1 - q_{v_1}w) + (v_2 - q_{v_2}w))^{\pm}.$$

Therefore

$$\{\widehat{\pi}(v_0^{\pm}), \widehat{\pi}(v_1^{\pm}), \widehat{\pi}(v_2^{\pm}), \ldots, \widehat{\pi}(v_p^{\pm})\}$$
$$= \{((v_1 - q_{v_1}w) + (v_2 - q_{v_2}w))^{\pm}, (v_1 - q_{v_1}w)^{\pm}, (v_2 - q_{v_2}w)^{\pm} \ldots, (v_p - q_{v_p}w)^{\pm}\}$$

is an internally additive simplex of $X^{<N}$. $\qquad\square$

It remains to deal with the carrying simplices. The key to our approach is the observation that even though the inequality (6) may hold, the two sides never differ by more than 1. Formally, this is the observation that the function ω_N on $\mathbb{Z}^2$ defined by

$$\left\lfloor \frac{n_1 + n_2}{N} \right\rfloor = \left\lfloor \frac{n_1}{N} \right\rfloor + \left\lfloor \frac{n_2}{N} \right\rfloor + \omega_N(n_1, n_2),$$

takes values *only in* $\{0, 1\}$. We remark that ω_N descends to a well-defined function $\overline{\omega}_N \colon (\mathbb{Z}/N)^2 \to \{0, 1\}$. Regarding its image as lying in $\{0, 1\} \subset \mathbb{Z}/N$, the function $\overline{\omega}_N$ is a group cocycle on $\mathbb{Z}/N$ whose cohomology class is the Euler class of the nonsplit central extension

$$0 \longrightarrow \mathbb{Z}/N \longrightarrow \mathbb{Z}/N^2 \longrightarrow \mathbb{Z}/N \longrightarrow 0.$$

It is known as the *carrying cocycle* because it records when carrying is necessary when adding modulo N; see Isaksen [I].

Let $\mathcal{C}$ be the set of 2-dimensional carrying simplices. Define the simplicial complex Y to be the result of subdividing X by adding a vertex τ_c to the center of each simplex $c \in \mathcal{C}$. No carrying simplex can be contained in $X^{<N}$ since $\omega_N(0, -) = \omega_N(-, 0) = 0$. Thus the subcomplex $X^{<N}$ is not affected by this subdivision, so we can regard $X^{<N}$ as a subcomplex of Y. Extend $\widehat{\pi} \colon X^{(0)} \to (X^{<N})^{(0)}$ to $\widehat{\pi} \colon Y^{(0)} \to (X^{<N})^{(0)}$ as follows.

- Given $c \in \mathcal{C}$, write $c = \{v_0^{\pm}, v_1^{\pm}, v_2^{\pm}\}$, where each v_i has $F(v_i) > 0$ and $v_0 = v_1 + v_2$. The ordering of the v_i is not canonical (as we mentioned above, v_0 is uniquely determined, but there is no way to distinguish v_1 and v_2), so simply make an arbitrary choice for each c. Define $\widehat{\pi}(\tau_c) = (v_1 - q_{v_1}w - w)^{\pm}$.

We first verify that $\widehat{\pi}(\tau_c)$ lies in $X^{<N}$. By definition, $0 \leq F(v_1 - q_{v_1} w) < N$. However, the fact that $\omega_N(N \cdot n_1, -) = 0$ means that a carrying simplex cannot have $F(v_1 - q_{v_1} w) = 0$ since this would imply $F(v_1) = q_{v_1} F(w) = q_{v_1} N$. Therefore $0 < F(v_1 - q_{v_1} w) < N$ and hence $-N < F(v_1 - q_{v_1} w - w) < 0$, as desired.

Claim 5. *The map $\widehat{\pi}$ extends over the images in Y of carrying simplices σ of X.*

Proof of claim. Write $\sigma = \{v_0^{\pm}, \ldots, v_p^{\pm}\}$, where each v_i is F-nonnegative and $v_0 = v_1 + v_2$ and where for $c = \{v_0^{\pm}, v_1^{\pm}, v_2^{\pm}\}$ we have $\widehat{\pi}(\tau_c) = (v_1 + q_{v_1} w - w)^{\pm}$. By the definition of a carrying simplex, we have

$$q_{v_0} = \left\lfloor \frac{F(v_1) + F(v_2)}{N} \right\rfloor = \left\lfloor \frac{F(v_1)}{N} \right\rfloor + \left\lfloor \frac{F(v_2)}{N} \right\rfloor + \omega_N(F(v_1), F(v_2)) = q_{v_1} + q_{v_2} + 1.$$

To simplify our notation, for $0 \leq i \leq p$ we define $v_i' = v_i - q_{v_i} w$, so $\widehat{\pi}(v_i^{\pm}) = (v_i')^{\pm}$. Observe that $\widehat{\pi}(\tau_c) = (v_1' - w)^{\pm}$ and $v_0' = v_1' + v_2' - w$.

The image of σ in Y consists of the three simplices

$$\alpha = \{\tau_c, v_1^{\pm}, v_2^{\pm}, v_3^{\pm}, \ldots, v_p^{\pm}\}, \quad \beta = \{v_0^{\pm}, \tau_c, v_2^{\pm}, v_3^{\pm}, \ldots, v_p^{\pm}\}, \quad \gamma = \{v_0^{\pm}, v_1^{\pm}, \tau_c, v_3^{\pm}, \ldots, v_p^{\pm}\}.$$

We verify that $\widehat{\pi}$ extends over each of these in turn. For the first simplex α and the third simplex γ, we use $\widehat{\pi}(\tau_c) = (v_1' - w)^{\pm}$ to write

$$\begin{aligned}
\pi(\alpha) &= \{\widehat{\pi}(\tau_c), \widehat{\pi}(v_1^{\pm}), \widehat{\pi}(v_2^{\pm}), \widehat{\pi}(v_3^{\pm}), \ldots, \widehat{\pi}(v_p^{\pm})\} \\
&= \{\widehat{\pi}(\tau_c), (v_1')^{\pm}, (v_2')^{\pm}, \ldots, (v_p')^{\pm}\} \\
&= \{(v_1' - w)^{\pm}, (v_1')^{\pm}, (v_2')^{\pm}, \ldots, (v_p')^{\pm}\}. \\
\pi(\gamma) &= \{\widehat{\pi}(v_0^{\pm}), \widehat{\pi}(v_1^{\pm}), \widehat{\pi}(\tau_c), \widehat{\pi}(v_3^{\pm}), \ldots, \widehat{\pi}(v_p^{\pm})\} \\
&= \{(v_0')^{\pm}, (v_1')^{\pm}, \widehat{\pi}(\tau_c), (v_3')^{\pm}, \ldots, (v_p')^{\pm}\} \\
&= \{(v_0')^{\pm}, (v_1')^{\pm}, (v_1' - w)^{\pm}, (v_3')^{\pm}, \ldots, (v_p')^{\pm}\}.
\end{aligned}$$

These are both $w^{\pm}$-additive simplices of $X^{<N}$. For the second simplex β, from $v_0' = v_1' + v_2' - w$ we deduce the alternate identity $\widehat{\pi}(\tau_c) = (v_0' - v_2')^{\pm}$, which we use to write

$$\begin{aligned}
\pi(\beta) &= \{\widehat{\pi}(v_0^{\pm}), \widehat{\pi}(\tau_c), \widehat{\pi}(v_2^{\pm}), \widehat{\pi}(v_3^{\pm}), \ldots, \widehat{\pi}(v_p^{\pm})\} \\
&= \{(v_0')^{\pm}, \widehat{\pi}(\tau_c), (v_2')^{\pm}, \ldots, (v_p')^{\pm}\} \\
&= \{(v_0')^{\pm}, (v_0' - v_2')^{\pm}, (v_2')^{\pm}, \ldots, (v_p')^{\pm}\}.
\end{aligned}$$

This is an internally additive simplex of $X^{<N}$. $\qquad\square$

Claims 1–5 demonstrate that $\widehat{\pi} \colon Y^{(0)} \to (X^{<N})^{(0)}$ extends over every simplex of Y, so it defines a simplicial retraction $\pi \colon Y \twoheadrightarrow X^{<N}$. Since Y is a subdivision of X, their realizations are homeomorphic, so this defines a topological retraction $\pi \colon X \twoheadrightarrow X^{<N}$. This completes the proof of Proposition 4.17. $\qquad\square$

4.5 The proof of Theorem C$'$

We finally prove Theorem C$'$, which asserts that $\mathcal{BA}_n^m$ is CM of dimension n for $n \geq 1$ and $m \geq 0$ with $m + n \geq 2$. The proof will be by induction on n.

Base case. We begin with the base case $n = 1$. Our goal is to prove for $m \geq 1$ that $\mathcal{BA}_1^m$ is CM of dimension 1, i.e. is a connected nonempty graph. The vertices of the 1-dimensional complex $\mathcal{BA}_1^m$ are the vertices of $\mathcal{B}_1^m$, namely the lines spanned by vectors $w \in \mathbb{Z}^{m+1}$ such that $\{e_1, \ldots, e_m, w\}$ is a basis for $\mathbb{Z}^{m+1}$. We can write such a vector as

$$w = a_1 e_1 + \cdots + a_m e_m \pm e_{m+1}$$

for some $a_i \in \mathbb{Z}$ and some sign. Replacing w with $-w$ changes the final sign, so we deduce that the vertices of $\mathcal{B}_1^m$ are in bijection with elements $\mathbf{a} \in \mathbb{Z}^m$ via the bijection that takes $\mathbf{a} = (a_1, \ldots, a_m)$ to the line $v_{\mathbf{a}}^{\pm}$ with

$$v_{\mathbf{a}} = a_1 e_1 + \cdots a_m e_m + e_{m+1}.$$

Every 1-simplex of $\mathcal{BA}_1^m$ is externally additive since an internally additive simplex has dimension at least 2. Two lines $v_{\mathbf{a}}^{\pm}$ and $v_{\mathbf{a}'}^{\pm}$ determine an externally additive 1-simplex precisely when $\varepsilon v_{\mathbf{a}'} + \varepsilon' v_{\mathbf{a}} + \varepsilon'' e_i = 0$ for some $1 \leq i \leq m$ and some $\varepsilon, \varepsilon', \varepsilon'' = \pm 1$. Examining the coefficient of e_{m+1} in this expression, we see that $\varepsilon = \varepsilon'$. This implies that $v_{\mathbf{a}}^{\pm}$ and $v_{\mathbf{a}'}^{\pm}$ determine an externally additive 1-simplex exactly when $\mathbf{a} \in \mathbb{Z}^m$ and $\mathbf{a}' \in \mathbb{Z}^m$ differ by a standard basis vector.

We conclude that $\mathcal{BA}_1^m$ is isomorphic to the Cayley graph of $\mathbb{Z}^m$ with respect to the generating set $\{e_1, \ldots, e_m\}$, and is thus connected. We remark that this is one point in the argument where working with lines and frames is essential; if we worked instead with primitive *vectors* and *bases*, we would obtain a disconnected graph consisting of two copies of this Cayley graph, one consisting of all vectors with e_{m+1} coordinate 1 and the other consisting of those with coordinate -1. This concludes the proof of the base case.

Inductive step. We now assume that $n > 1$ and that $\mathcal{BA}_{n'}^{m'}$ is CM of dimension n' for all $1 \leq n' < n$ and $m' \geq 0$ with $m' + n' \geq 2$. Under these assumptions, Proposition 4.14 states that all links in $\mathcal{BA}_n^m$ are CM of the appropriate dimension, so it is enough to prove that $\mathcal{BA}_n^m$ is $(n-1)$-connected.

Fix $0 \leq p \leq n - 1$, let S^p be a combinatorial triangulation of a p-sphere, and let $\varphi \colon S^p \to \mathcal{BA}_n^m$ be a simplicial map. Our goal is to show that φ can be homotoped to a constant map. Let $F \colon \mathbb{Z}^{m+n} \twoheadrightarrow \mathbb{Z}$ be the linear map taking $v \in \mathbb{Z}^{m+n}$ to the e_{m+n}-coordinate of v. For a vertex $v^{\pm}$ of $\mathcal{BA}_n^m$, define $r(v^{\pm}) = |F(v)|$; this is well-defined since $|F(v)| = |F(-v)|$. We then define

$$R(\varphi) = \max\{r(\varphi(x)) \mid x \text{ a vertex of } S^p\}.$$

This will be our measure of complexity for φ.

If $R(\varphi) = 0$, then every simplex σ of $\varphi(S^p)$ is contained in the summand $\ker F$ of $\mathbb{Z}^{m+n}$. In particular, $\{e_1^{\pm}, \ldots, e_m^{\pm}\} * \sigma$ is a partial augmented frame contained in $\ker F$; by Remark 4.6, it is in fact a partial augmented frame for $\ker F$, so it can be extended to a partial augmented frame for $\mathbb{Z}^{n+m}$ by adding the line $e_{m+n}^{\pm}$. In other words, the entire image $\varphi(S^p)$ is contained in the star of $e_{m+n}^{\pm}$ within $\mathcal{BA}_n^m$, so we can directly contract φ to the constant map whose image is the vertex $e_{m+n}^{\pm}$.

We can therefore assume that $R(\varphi) = R > 0$. The proof now is divided into four steps. The end product of these four steps is that we can homotope φ so as to decrease $R(\varphi)$. Repeating these steps over and over, we can eventually homotope φ so that $R(\varphi) = 0$, at which point we can contract φ directly to a constant map as above.

Step 1. *Given $\varphi\colon S^p \to \mathcal{B}\mathcal{A}_n^m$ with $R(\varphi) \leq R$, we can homotope φ so that it satisfies the following Conditions C1 and C2.*

C1. *We still have $R(\varphi) \leq R$.*

C2. *If σ is a simplex of S^p such that $\varphi(\sigma)$ is an additive simplex, then for all vertices x of σ we have $r(\varphi(x)) < R$.*

Consider the following condition on a simplex σ of S^p:

$\varphi(\sigma)$ is an additive simplex, and

some vertex $v^\pm$ of $\varphi(\sigma)$ has $r(v^\pm) = R$, and $\qquad\qquad\qquad\qquad\qquad\qquad$ (7)

every vertex $w^\pm$ of $\varphi(\sigma)$ either has $r(w^\pm) = R$ or lies in the additive core of $\varphi(\sigma)$.

If φ does not satisfy Conditions C1 and C2, then there must be some simplex σ of S^p satisfying (7). We can therefore choose a simplex σ of S^p satisfying (7) whose dimension k is maximal among those satisfying (7). This maximality implies that φ takes $\mathrm{Link}_{S^p}(\sigma)$ to $\mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(\varphi(\sigma))^{<R}$.

Let ℓ be the dimension of $\varphi(\sigma)$; we certainly have $\ell \leq k$, but we might have $\ell < k$ if the restriction of φ to σ is not injective. Lemma 4.12(a) states that $\mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(\varphi(\sigma))$ is isomorphic to $\mathcal{B}_{n-\ell}^{m+\ell}$, which by Theorem 4.2 is CM of dimension $(n-\ell-1)$, and in particular is $(n-\ell-2)$-connected. Using Lemma 4.16, we deduce that its retract $\mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(\varphi(\sigma))^{<R}$ is $(n-\ell-2)$-connected.

By the definition of a combinatorial triangulation, $\mathrm{Link}_{S^p}(\sigma)$ is a combinatorial $(p-k-1)$-sphere. Since $p \leq n-1$ and $\ell \leq k$, we have $p-k-1 \leq n-\ell-2$, so $\varphi|_{\mathrm{Link}_{S^p}(\sigma)}$ is null-homotopic within $\mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(\varphi(\sigma))^{<R}$. Using Zeeman's relative simplicial approximation theorem [Z], we deduce that there exists a combinatorial $(p-k)$-ball B with $\partial B \cong \mathrm{Link}_{S^p}(\sigma)$ and a simplicial map $\psi\colon B \to \mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(\varphi(\sigma))^{<R}$ such that $\psi|_{\partial B} = \varphi|_{\mathrm{Link}_{S^p}(\sigma)}$.

The map ψ extends to the $(p+1)$-ball $\sigma * B$ as $(\varphi|_\sigma) * \psi\colon \sigma * B \to \mathcal{B}_n^m$. The boundary of $\sigma * B$ is the union of the p-ball $\sigma * (\partial B) = \mathrm{Star}_{S^p}(\sigma)$, on which $\varphi|_\sigma * \psi = \varphi|_{\mathrm{Star}_{S^p}(\sigma)}$, and the p-ball $(\partial\sigma) * B$. We can thus homotope φ across this $(p+1)$-ball to replace $\varphi|_{\mathrm{Star}_{S^p}(\sigma)}$ with $\varphi|_{\partial\sigma} * \psi\colon (\partial\sigma) * B \to \mathcal{B}_n^m$.

The key property of this modification is that it eliminates the simplex σ and does not add any other simplices satisfying (7) or any vertices mapping to vertices with $r(v^\pm) \geq R$. Indeed, every new vertex lies in B, which maps to $\mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(\varphi(\sigma))^{<R}$ by construction; this verifies the second claim. Moreover, every new simplex τ is the join of a simplex in $\partial\sigma$ with a nonempty simplex ρ in B. Its image $\varphi(\tau)$ is contained in $\varphi(\sigma * \rho) = \varphi(\sigma) * \varphi(\rho)$. Thus $\varphi(\tau)$ is only additive if it contains the additive core of $\varphi(\sigma)$, in which case this is also the additive core of $\varphi(\tau)$. Since $\varphi(\rho)$ is disjoint from $\varphi(\sigma)$ and every vertex has $r(v^\pm) < R$ by construction, the new simplex τ cannot satisfy (7).

Repeating this modification, we can homotope φ so that no simplex of S^p satisfies (7), so φ satisfies Conditions C1 and C2, as desired.

Step 2. *Given $\varphi\colon S^p \to \mathcal{B}\mathcal{A}_n^m$ satisfying Conditions C1 and C2, we can homotope φ so that it still satisfies the same Conditions C1 and C2, and additionally satisfies the following Condition C3′.*

C1. *We still have $R(\varphi) \leq R$.*

C2. *If σ is a simplex of S^p such that $\varphi(\sigma)$ is an additive simplex, then for all vertices x of σ we have $r(\varphi(x)) < R$.*

C3′. *If x_1 and x_2 are distinct vertices of S^p such that $r(\varphi(x_1)) = r(\varphi(x_2)) = R$ and $\varphi(x_1) = \varphi(x_2)$, then x_1 and x_2 are not joined by an edge in S^p.*

Although Conditions C1 and C2 are restated multiple times in this section for convenience, we emphasize that these conditions are unchanged throughout.

Consider the following condition on a simplex σ of S^p:

$$\begin{aligned} &\varphi|_\sigma \text{ is not injective, and} \\ &\text{every vertex } v^\pm \text{ of } \varphi(\sigma) \text{ has } r(v^\pm) = R. \end{aligned} \tag{8}$$

If φ satisfies Conditions C1 and C2 but not Condition C3′, then there must be some simplex σ of S^p satisfying (8). We can therefore choose a simplex σ of S^p satisfying (8) whose dimension k is maximal among those satisfying (8). This maximality implies that φ takes $\mathrm{Link}_{S^p}(\sigma)$ to $\mathrm{Link}_{\mathcal{BA}_n^m}(\varphi(\sigma))^{<R}$.

In fact, even more is true. Namely, Condition C2 implies that if τ is a simplex of $\mathrm{Link}_{S^p}(\sigma)$, then the simplex $\varphi(\tau) * \varphi(\sigma)$ of $\mathcal{BA}_n^m$ must be a standard simplex. This implies that φ actually takes $\mathrm{Link}_{S^p}(\sigma)$ to the subcomplex $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$ of $\mathrm{Link}_{\mathcal{BA}_n^m}(\varphi(\sigma))^{<R}$.

Let ℓ be the dimension of $\varphi(\sigma)$; since $\varphi|_\sigma$ is not injective, we have $\ell \leq k-1$. Theorem 4.2 says that $\mathcal{B}_n^m$ is CM of dimension $(n-1)$, and hence $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))$ is $(n-\ell-3)$-connected. By Lemma 4.5, its retract $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$ is $(n-\ell-3)$-connected.

The complex $\mathrm{Link}_{S^p}(\sigma)$ is a combinatorial $(p-k-1)$-sphere. Since $p \leq n-1$ and $\ell \leq k-1$, we have $p-k-1 \leq n-\ell-3$, so $\varphi|_{\mathrm{Link}_{S^p}(\sigma)}$ is null-homotopic within the subcomplex $\mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$ of $\mathcal{BA}_n^m$. Therefore there exists a combinatorial $(p-k)$-ball B with $\partial B \cong \mathrm{Link}_{S^p}(\sigma)$ and a simplicial map $\psi\colon B \to \mathrm{Link}_{\mathcal{B}_n^m}(\varphi(\sigma))^{<R}$ such that $\psi|_{\partial B} = \varphi|_{\mathrm{Link}_{S^p}(\sigma)}$.

As in the previous step, we can use this ball to homotope φ so as to replace $\varphi|_{\mathrm{Star}_{S^p}(\sigma)}$ with $\varphi|_{\partial\sigma} * \psi\colon (\partial\sigma) * B \to \mathcal{BA}_n^m$. The key property of this modification is that it eliminates the simplex σ and does not add any other simplices satisfying (8), while preserving Conditions C1 and C2. Indeed, every new vertex has $r(v^\pm) < R$, so Condition C1 is preserved. Every new simplex contains a new vertex, so it cannot satisfy (8). Finally, none of the simplices involved are additive since the modifications in this step take place within the subcomplex $\mathcal{B}_n^m$, so Condition C2 is preserved. Repeating this, we can ensure that no simplices satisfy (8) while preserving Conditions C1 and C2, as desired.

Step 3. *Given $\varphi\colon S^p \to \mathcal{BA}_n^m$ satisfying Conditions C1, C2, and C3′, we can homotope φ so that it satisfies the following Condition C3, as well as Conditions C1 and C2.*

C1. *We still have $R(\varphi) \leq R$.*

C2. *If σ is a simplex of S^p such that $\varphi(\sigma)$ is an additive simplex, then for all vertices x of σ we have $r(\varphi(x)) < R$.*

C3. *If x_1 and x_2 are distinct vertices of S^p such that $\mathrm{rank}(\varphi(x_1)) = \mathrm{rank}(\varphi(x_2)) = R$, then x_1 and x_2 are not joined by an edge in S^p.*

If φ does not satisfy Condition C3, then there exists an edge $e = \{x_1, x_2\}$ of S^p with $r(\varphi(x_1)) = r(\varphi(x_2)) = R$. Choose such an edge e. We will homotope φ so as to eliminate e without disturbing Conditions C1, C2, or C3′.

Choose $v_1, v_2 \in \mathbb{Z}^{m+n}$ with $F(v_1) = F(v_2) = R$ such that $\varphi(x_1) = v_1^\pm$ and $\varphi(x_2) = v_2^\pm$. Set $v_0 = v_1 - v_2$, so $F(v_0) = 0$. Condition C3′ guarantees that $v_1^\pm \neq v_2^\pm$, so $v_0 \neq 0$. Thus $\{v_0^\pm, v_1^\pm, v_2^\pm\}$ is an internally additive simplex and $v_0^\pm$ lies in $\mathrm{Link}_{\mathcal{BA}_n^m}(\{v_1^\pm, v_2^\pm\})$.

Moreover, we claim that $\varphi(\mathrm{Link}_{S^p}(e))$ is contained in the star of $v_0^\pm$ inside the subcomplex $\mathrm{Link}_{\mathcal{BA}_n^m}(\{v_1^\pm, v_2^\pm\})$. To see this, consider an arbitrary simplex $\tau = \{w_1^\pm, \ldots, w_k^\pm\}$ in $\varphi(\mathrm{Link}_{S^p}(e))$. Condition C3′ implies that τ is disjoint from $\{v_1^\pm, v_2^\pm\}$, so τ lies in $\mathrm{Link}_{\mathcal{BA}_n^m}(\{v_1^\pm, v_2^\pm\})$. By Condition C2, all simplices in $\varphi(\mathrm{Star}_{S^p}(e))$ are standard, so in

22

fact τ is a simplex of $\mathrm{Link}_{\mathcal{B}_n^m}(\{v_1^\pm, v_2^\pm\})$. This means that τ is contained is some frame $\{e_1^\pm, \ldots, e_m^\pm, v_1^\pm, v_2^\pm, w_1^\pm, \ldots, w_{n-2}^\pm\}$ for $\mathbb{Z}^{m+n}$. Then $\{e_1^\pm, \ldots, e_m^\pm, v_0^\pm, v_1^\pm, v_2^\pm, w_1^\pm, \ldots, w_{n-2}^\pm\}$ is an augmented frame, so τ is contained in the star of $v_0^\pm$, as desired.

Let B be the cone on the combinatorial $(p-2)$-sphere $\mathrm{Link}_{S^p}(e)$. Since $\varphi(\mathrm{Link}_{S^p}(e))$ is contained in the star of $v_0^\pm$, we can extend $\varphi|_{\mathrm{Link}_{S^p}(e)}$ to $\psi\colon B \to \mathrm{Link}_{\mathcal{B}_n^m}(v^\pm)$ by sending the cone point to $v_0^\pm$. As before, this lets us homotope φ to replace $\varphi|_{\mathrm{Star}_{S^p}(e)}$ with $\varphi|_{\partial e} * \psi\colon (\partial e) * B \to \mathcal{B}\mathcal{A}_n^m$. This eliminates the edge e. The only new vertex is $v_0^\pm$, so Conditions C1 and C3′ are preserved since $F(v_0) = 0$. Moreover, every new simplex in $(\partial e) * B$ contains $v_0^\pm$ but at most one of the lines $v_1^\pm$ and $v_2^\pm$. Accordingly the new simplices are all standard, so Condition C2 is preserved by this modification. Repeating this process lets us eliminate all edges violating Condition C3, as desired.

Step 4. *Given $\varphi\colon S^p \to \mathcal{B}\mathcal{A}_n^m$ satisfying Conditions C1, C2 and C3, we can homotope φ such that $R(\varphi) < R$.*

We remark that Condition C2 is not used in this step, though it is essential during Steps 2 and 3.

If $R(\varphi) = R$, then we can choose a vertex x of S^p such that $r(\varphi(x)) = R$. Write $\varphi(x) = v^\pm$ with $F(v) = R$. Conditions C1 and C3 imply that φ takes $\mathrm{Link}_{S^p}(x)$ to $\mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)^{<R}$. We would like to apply Lemma 4.12(b) and Proposition 4.17 to conclude that this subcomplex is highly connected. However, these results apply not to $\mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)$, but to its proper subcomplex $\widehat{\mathrm{Link}}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)$ defined in Definition 4.11. Nevertheless, the vertices of the former that are excluded from the latter subcomplex are the $2m$ vertices of the form $(v + e_i)^\pm$ and $(v - e_i)^\pm$ for $i = 1, \ldots, m$. Since $F(v + e_i) = F(v - e_i) = F(v) = R$, these complexes *do* coincide when we restrict to vertices with $r(w^\pm) < R$; in other words, $\widehat{\mathrm{Link}}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)^{<R} = \mathrm{Link}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)^{<R}$.

Lemma 4.12(b) states that $\widehat{\mathrm{Link}}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)$ is isomorphic to $\mathcal{B}\mathcal{A}_{n-1}^{m+1}$. Since $n > 1$, our induction hypothesis states that this is CM of dimension $(n-1)$, and in particular is $(n-2)$-connected. Proposition 4.17 states that $\widehat{\mathrm{Link}}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)$ admits a topological retraction onto $\widehat{\mathrm{Link}}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)^{<R}$, so $\widehat{\mathrm{Link}}_{\mathcal{B}\mathcal{A}_n^m}(v^\pm)^{<R}$ is $(n-2)$-connected as well.

The complex $\mathrm{Link}_{S^p}(x)$ is a combinatorial $(p-1)$-sphere. Since $p \leq n-1$, we have $p-1 \leq n-2$, so the restriction $\varphi|_{\mathrm{Link}_{S^p}(\{x\})}$ is null-homotopic within $\widehat{\mathrm{Link}}_{\mathcal{B}\mathcal{A}_n^m}(\{\varphi(x)\})^{<R}$. Just as in previous steps, this allows us to homotope φ so as to eliminate x without introducing any new vertices mapping to vertices with $r(w^\pm) = R$. This guarantees that Conditions C1 and C3 are preserved by this modification. Repeating this process lets us eliminate all vertices mapping to vertices with $r(v^\pm) = R$, at which point $R(\varphi) < R$, as desired. This completes Step 4.

Repeating the modifications of Steps 1–4, we can homotope φ so as to reduce $R(\varphi)$ to 0, at which point φ can be contracted to the constant map at the vertex $e_{m+n}^\pm$ as discussed before Step 1. We conclude that an arbitrary map $\varphi\colon S^p \to \mathcal{B}\mathcal{A}_n^m$ for $0 \leq p \leq n-1$ is null-homotopic, demonstrating that $\mathcal{B}\mathcal{A}_n^m$ is $(n-1)$-connected. This completes the proof of Theorem C′.

Remark 4.18. The hardest part of Theorem C′ was showing that $\mathcal{B}\mathcal{A}_n^m$ is $(n-1)$-connected. We proved this connectivity by defining an $\mathbb{N}$-valued function r on the vertices of $\mathcal{B}\mathcal{A}_n^m$, for which the subcomplex where $r(x) = 0$ is contractible inside $\mathcal{B}\mathcal{A}_n^m$, and showed that any sphere in $\mathcal{B}\mathcal{A}_n^m$ of the appropriate dimension could be homotoped to lie in this subcomplex. From this outline the argument seems similar to "PL Morse theory" arguments, a common

technique when proving such connectivity results. However, the structure of our proof is rather nonstandard and departs greatly from the PL Morse theory framework. Still, given the obvious similarities it is natural to wonder if our proof can be phrased in this language.

Briefly, PL Morse theory tells us that if we can find a function $F\colon X^{(0)} \to \mathbb{N}$ on the vertices of a simplicial complex X such that

- there are no "horizontal edges", i.e. edges $\{x, y\}$ with $F(x) = F(y)$, and
- the "descending link" of each vertex x with $F(x) > 0$, i.e. the full subcomplex of the link spanned by vertices y with $F(y) < F(x)$, is $(m-1)$-connected,

then the inclusion into X of the subcomplex where $F(x) = 0$ is m-connected.

However, our function r is definitely not a PL Morse function; its descending links are not highly connected and it has many horizontal edges. In fact, we do not believe that any such PL Morse function can be defined on the vertices of $\mathcal{BA}_n^m$.

It turns out that it is possible to capture the argument in §4.5 via a PL Morse function F, but only after passing to the barycentric subdivision of $\mathcal{BA}_n^m$. Unfortunately, this function is rather unwieldy (see below). Moreover, to verify that the descending links have the appropriate connectivity requires recapitulating every step in §4.1, §4.2, §4.3, and §4.4, so ultimately this perspective would provide no benefit.

For the interested reader: after defining the function r on vertices of $\mathcal{BA}_n^m$ as in our proof, for a simplex $\sigma \in \mathcal{BA}_n^m$ one should define $F(\sigma) \in \mathbb{N} \times \mathbb{N} \times \mathbb{Z} \times -\mathbb{N}$ by

$$
F(\sigma) = \Big(\max_{x \text{ vertex of } \sigma} r(x),
$$

$$
\# \text{ of vertices of } \sigma \text{ realizing the maximum } \max_{x \in \sigma} r(x),
$$

$$
\begin{cases}
-2 & \text{if } \sigma \text{ is additive and } \max_{x \in \sigma} r(x) \text{ is realized by a unique vertex} \\
-1 & \text{if } \sigma \text{ is internally additive and two of the three vertices} \\
   & \quad \text{in its additive core realize the maximum } \max_{x \in \sigma} r(x) \\
0 & \text{if } \sigma \text{ is standard} \\
1 & \text{otherwise,}
\end{cases}
$$

$$
- \dim(\sigma) \ \Big).
$$

If $\mathbb{N} \times \mathbb{N} \times \mathbb{Z} \times -\mathbb{N}$ is given the lexicographic order, then $F\colon \mathcal{P}(\mathcal{BA}_n^m) \to \mathbb{N} \times \mathbb{N} \times \mathbb{Z} \times -\mathbb{N}$ is a PL Morse function with well-ordered image whose descending links are $(n-2)$-connected.

References

[AR]   A. Ash and L. Rudolph, The modular symbol and continued fractions in higher dimensions, *Invent. Math.* 55 (1979), no. 3, 241–250.

[Bo]   A. Borel, Stable real cohomology of arithmetic groups, *Ann. Sci. École Norm. Sup.* (4) 7 (1974), 235–272.

[BoSe] A. Borel and J.-P. Serre, Corners and arithmetic groups, *Comment. Math. Helv.* 48 (1973), 436–491.

[Br1]   K. S. Brown, *Cohomology of groups*, corrected reprint of the 1982 original, Graduate Texts in Mathematics, 87, Springer, New York, 1994.

[Br2]   K. S. Brown, *Buildings*, Springer, New York, 1989. Freely available at
        `http://www.math.cornell.edu/~kbrown/buildings/`

[By]   V.A. Bykovskii, Generating elements of the annihilating ideal for modular symbols. (Russian) *Funktsional. Anal. i Prilozhen.* 37 (2003), no. 4, 27–38, 95; translation in *Funct. Anal. Appl.* 37 (2003), no. 4, 263–272.

[CFP1]   T. Church, B. Farb, and A. Putman, A stability conjecture for the unstable cohomology of $SL_n \mathbb{Z}$, mapping class groups, and $Aut(F_n)$, in *Algebraic Topology: Applications and New Directions*, 55–70, Contemp. Math. 620, Amer. Math. Soc., Providence, RI. arXiv:1208.3216.

[CFP2]   T. Church, B. Farb, and A. Putman, Integrality in the Steinberg module and the top-dimensional cohomology of $GL_n(\mathcal{O}_K)$, arXiv:1501.01307v1.

[DP]   M. Day and A. Putman, On the second homology group of the Torelli subgroup of $Aut(F_n)$, preprint 2014, arXiv:1408.6242.

[I]   D. C. Isaksen, A cohomological viewpoint on elementary school arithmetic, *Amer. Math. Monthly* 109 (2002), no. 9, 796–805.

[LSz]   R. Lee and R. H. Szczarba, On the homology and cohomology of congruence subgroups, *Invent. Math.* 33 (1976), no. 1, 15–53.

[Maa]   H. Maazen, Homology Stability for the General Linear Group, thesis, University of Utrecht (1979).

[Man]   Ju. I. Manin, Parabolic points and zeta functions of modular curves. (Russian) *Izv. Akad. Nauk SSSR Ser. Mat.* 36 (1972), 19–66; translation in *Math. USSR Izv.* Vol. 6 (1972), No. 1, 19–64.

[Q]   D. Quillen, Homotopy properties of the poset of nontrivial p-subgroups of a group, *Adv. in Math.* 28 (1978), no. 2, 101–128.

[Z]   E. C. Zeeman, Relative simplicial approximation, *Proc. Cambridge Philos. Soc.* 60 (1964), 39–43.

Thomas Church   Andrew Putman
Department of Mathematics   Department of Mathematics
Stanford University   Rice University, MS 136
450 Serra Mall   6100 Main St.
Stanford, CA 94305   Houston, TX 77005
church@math.stanford.edu   andyp@math.rice.edu