MORSE MATCHINGS ON POLYTOPES

R.M. Green and Jacob T. Harper

Department of Mathematics
University of Colorado
Campus Box 395
Boulder, CO 80309-0395
USA
E-mail: rmg@euclid.colorado.edu,
Jacob.Harper@colorado.edu

Abstract. We show how to construct homology bases for certain CW complexes in terms of discrete Morse theory and cellular homology. We apply this technique to study certain subcomplexes of the half cube polytope studied in previous works. This involves constructing explicit complete acyclic Morse matchings on the face lattice of the half cube; this procedure may be of independent interest for other highly symmetric polytopes.

Preliminary version, draft 2

1. Introduction

As a graph, the n-dimensional hypercube is bipartite and connected. This induces a partition of its vertex set $V = V_n = \{\pm 1\}^n$ into two pieces, $V^e \cup V^o = V_n^e \cup V_n^o$, where V_n^e (respectively, V_n^o) consists of those vertices whose coordinates contain an even (respectively, odd) number of occurrences of -1. We define the *half cube*, Γ_n, to be the convex hull of the 2^{n-1} points in V_n^e. Using V_n^o in place of V_n^e in this construction gives rise to an isometric copy of the half cube.

1991 *Mathematics Subject Classification*. 52B11.

In a previous work [**10**], the first author classified the faces of the half cube and explained how they assemble naturally into a regular CW complex, C_n, which is homeomorphic to a ball (see Theorem 4.2). Furthermore, for each $3 \leq k \leq n$, there is an interesting subcomplex $C_{n,k}$ of C_n obtained by deleting the interiors of all the half cube shaped faces of dimensions $l \geq k$. We also showed in [**10**, Theorem 3.3.2] that the reduced homology of $C_{n,k}$ is free over $\mathbb{Z}$ and concentrated in degree $k-1$.

The Coxeter group $W(D_n)$ acts naturally on the $(k-1)$-st homology of $C_{n,k}$, and we computed the character of this representation (over $\mathbb{C}$) in [**11**, Theorem 4.4]. The group $W(D_n)$ has two parabolic subgroups that are isomorphic to the symmetric group S_n, so the homology representations become S_n-modules by restriction. We showed in [**11**, Theorem 4.7] that the resulting representations of S_n are equivalent to the representation of S_n on the $(k-2)$-nd homology of the complement of the k-equal real hyperplane arrangement.

If $k = n-1$, the complex $C_{n,k}$ is the boundary complex of the half cube, and is therefore shellable by a well-known result of Bruggesser–Mani [**4**]. If $k < n-1$, then the fact that $C_{n,k}$ has nonzero $(k-1)$-st homology is an obstruction to shellability, which means that we cannot use the machinery of shellability to produce a homology basis for $C_{n,k}$.

The first main result of this paper (Theorem 3.7) shows how to use cellular homology and discrete Morse theory to construct an explicit integral homology basis for certain kinds of regular CW complexes, of which the complexes $C_{n,k}$ are motivating examples. In order to apply the theorem to a CW complex Y, one starts with a complete acyclic Morse matching V of a CW complex X that contains Y as a subcomplex. Under certain mild additional hypotheses, which are satisfied by $C_{n,k}$, the theorem produces an explicit set of boundaries in X that induce a homology basis for Y.

In §4, we construct a complete acyclic matching on the face lattice of the half cube. It is not a surprise that such a matching exists, as this follows abstractly from the shellability of the boundary complex of a polytope. Our motivation is to construct an explicit matching that works for half cubes of arbitrary dimension $n \geq 4$. We prove in Theorem 5.8 that this is a complete acyclic matching on the faces of the half cube (together with the empty face).

Let $b_{n,k}$ be the $(k-1)$-st Betti number of the subcomplex $C_{n,k}$. The numbers $b_{n,k}$ appear as sequence A119258 in Sloane's online encyclopedia [**16**]. Various explicit expressions for these numbers are given in [**15**], some of which are sums of products of positive integers. One of these appears in work of Björner–Welker [**3**], who study the numbers $b_{n,k}$ in the context of hyperplane arrangements. They prove that

$$b_{n,k} = \sum_{i=k}^{n} \binom{n}{i} \binom{i-1}{k-1}. \tag{1.1}$$

There is a representation theoretic explanation for this: the terms in the sum correspond to the dimensions of the irreducible constituents of the representation of the Coxeter group $W(D_n)$ on the reduced homology of $C_{n,k}$; this was shown in [**11**].

Another formula for $b_{n,k}$ is

$$b_{n,k} = \sum_{i=1}^{n} 2^{i-k} \binom{i-1}{k-1}; \tag{1.2}$$

this is a straightforward generalization of a result of Barcelo and Smith [**1**], who study the case $k = 3$ in the context of combinatorial homotopy theory ("A-theory"). In Theorem 6.2, we construct an explicit homology basis for $C_{n,k}$ in terms of cellular homology; this basis has the property that when it is enumerated in the obvious way, we recover equation (1.2).

We believe that the quest for explicit Morse matchings on the face lattices of polytopes is an aesthetically pleasing goal in its own right, similar to the discovery of an explicit shelling order on the faces of a polytope. It would be interesting to find such matchings for other polytopes. For some, such as the hypercube and the simplex, this is a fairly easy exercise. Others, such as the hypersimplex, present about the same level of difficulty as the half cube; this is described in the second author's thesis [**12**] and will be published separately. It would be very interesting to have such a description for the permutahedron, some of whose subcomplexes are known to have important topological properties [**2**, Theorem 2.4].

2. Discrete Morse theory for CW complexes

We first recall the definition of a finite regular CW complex, following [**14**, §8].

If X and Y are topological spaces with $A \subset X$ and $B \subset Y$, we define a continuous map $g : (X, A) \longrightarrow (Y, B)$ to be a continuous map $g : X \longrightarrow Y$ such that $g(A) \subseteq B$. If, furthermore, $g|_{X-A} : X - A \longrightarrow Y - B$ is a homeomorphism, we call G a *relative homeomorphism*.

An *n-cell*, $e = e^n$ is a homeomorphic copy of the open n-disk $D^n - S^{n-1}$, where D^n is the closed unit ball in Euclidean n-space and S^{n-1} is its boundary, the unit $(n-1)$-sphere. We call e a *cell* if it is an n-cell for some n.

If a topological space X is a disjoint union of cells $X = \bigcup\{e : e \in E\}$, then for each $k \geq 0$, we define the k-skeleton $X^{(k)}$ of X by

$$X^{(k)} = \bigcup\{e \in E : \dim(e) \leq k\}.$$

The CW complexes we consider in this paper are all finite, which means that we can give the following abbreviated definition.

Definition 2.1. A *CW complex* is an ordered triple (X, E, Φ), where X is a Hausdorff space, E is a family of cells in X, and $\{\Phi_e : e \in E\}$ is a family of maps, such that

(i) $X = \bigcup\{e : e \in E\}$ is a disjoint union;
(ii) for each k-cell $e \in E$, the map $\Phi_e : (D^k, S^{k-1}) \longrightarrow (e \cup X^{(k-1)}, X^{(k-1)})$ is a relative homeomorphism.

A *subcomplex* of the CW complex (X, E, Φ) is a triple $(|E'|, E', \Phi')$, where $E' \subset E$,

$$|E'| := \bigcup\{e : e \in E'\} \subset X,$$

$\Phi' = \{\Phi_e : e \in E'\}$, and $\text{Im } \Phi_e \subset |E'|$ for every $e \in E'$.

The complexes considered here have the property that the maps Φ_e (regarded as mapping to their images) are all homeomorphisms. Such CW complexes are called *regular*.

An *oriented CW complex* is a CW complex together with a choice of orientation for each cell.

Cellular homology is a version of singular homology that is particularly convenient in the context of regular CW complexes. For our purposes, it is convenient to describe cellular homology in terms of intersection numbers as follows. If e^n_α is an n-cell and e^{n-1}_β is an $(n-1)$-cell of the same CW complex X, then the *incidence number* $[e^n_\alpha : e^{n-1}_\beta]$ is defined in [**9**, §2.5] as the degree of a certain map. It follows that the incidence number is an integer. A key property for our purposes is the following.

Proposition 2.2. *If X is an oriented regular CW complex then the intersection number $[e^n_\alpha : e^{n-1}_\beta]$ is equal to ± 1 if e^{n-1}_β is a face of e^n_α, and is equal to 0 otherwise.*

Proof. This is [**9**, Proposition 5.3.10]. □

To define the cellular homology of a CW complex X over a ring R, we introduce, for each integer $n \geq 0$, the n-chains of X. This is the free R-module with a basis indexed by all the n-cells, e^n_α; by abuse of notation, we will identify the basis elements with the cells, having fixed once and for all on an orientation for each cell. The boundary map $\partial = \partial_n : C_n(X;R) \longrightarrow C_{n-1}(X;R)$ is then defined to be the R-module homomorphism for which

$$\partial(e^n_\alpha) = \sum_\beta [e^n_\alpha : e^{n-1}_\beta] e^{n-1}_\beta.$$

It can be shown that $\partial \circ \partial = 0$. The homology of the complex $C_\bullet$ is the *cellular homology* of X over R. It is convenient for some purposes to introduce a unique -1-cell e^{-1}_α; this gives rise to *reduced* cellular homology.

Discrete Morse theory, which was introduced by Forman [**7**], is a combinatorial technique for computing the homology of CW complexes. By building on work of Chari [**6**], Forman later produced a version of discrete Morse theory based on acyclic matchings in Hasse diagrams [**8**]. This version of the theory plays a key role in computing the homology of $C_{n,k}$.

Definition 2.3. Let K be a finite regular CW complex. A *discrete vector field* on K is a collection of pairs of cells (K_1, K_2) such that

(i) K_1 is a face of K_2 of codimension 1 and
(ii) every cell of K lies in at most one such pair.

We call a cell of K *paired* if it lies in (a unique) one of the above pairs, and *unpaired* otherwise. If (K_1, K_2) is a pair of the matching as above, we say that K_1 is an *upward matching* face and that K_2 is a *downward matching* face.

If V is a discrete vector field on a regular CW complex K, we define a V-path to be a sequence of cells

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \alpha_2, \dots, \beta_r, \alpha_{r+1}$$

such that for each $i = 0, \dots, r$, (a) each of α_i and α_{i+1} is a codimension 1 face of β_i, (b) each (α_i, β_i) belongs to V and (c) $\alpha_i \neq \alpha_{i+1}$ for all $0 \leq i \leq r$. If $r \geq 0$, we call the V-path *nontrivial*, and if $\alpha_0 = \alpha_{r+1}$, we call the V-path *closed.* Note that all the faces α_i have the same dimension, p say, and all the faces β_i have dimension $p + 1$.

Let P be the set of cells of K, together with the empty cell $\emptyset$, which we consider to be a cell of dimension -1. Denote the set of k-cells of K by P^k. The set P becomes a partially ordered set under inclusion. Let H be the Hasse diagram of this partial order. We regard H as a directed graph, in which all edges point towards cells of larger dimension.

Suppose now that V is a discrete vector field on K. We define $H(V)$ to be the directed graph obtained from H by reversing the direction of an arrow if and only if it joins two cells $K_1 \subset K_2$ for which (K_1, K_2) is one of the pairs of V. If the graph $H(V)$ has no directed cycles, we call V an *acyclic (partial) matching* of the Hasse diagram of K.

Theorem 2.4 (Forman). *Let V be a discrete vector field on a regular CW complex K.*

(i) *There are no nontrivial closed V-paths if and only if V is an acyclic matching of the Hasse diagram of K.*

(ii) *Suppose that V is an acyclic partial matching of the Hasse diagram of K in which the empty set is unpaired. Let u_p denote the number of unpaired p-cells. Then K is homotopic to a CW complex with exactly u_p cells of dimension p for each $p \geq 0$.*

Proof. Part (i) is [**8**, Theorem 6.2] and part (ii) is [**8**, Theorem 6.3]. □

3. Homology bases for subcomplexes

Definition 3.1. Let X be a finite regular oriented CW complex and let V be an acyclic partial matching on K. Recall that P^i is the set of i-cells of K. For each k, let $V_k = V \cap (P^k \times P^{k+1})$. Define

$$d_{V,k} = \{K \in P^{k+1} : (K_1, K) \in V_k \text{ for some } K_1 \in P^k\}$$

and

$$e_{V,k} = \{K \in P^k : (K, K_2) \in V_k \text{ for some } K_2 \in P^{k+1}\}.$$

Let $D_{V,k}(X;R)$ (respectively, $E_{V,k}(X;R)$) be the free R-module on $d_{V,k}$ (respectively, $e_{V,k}$). Let $\partial_{V,k} : D_{V,k}(X;R) \longrightarrow E_{V,k}(X;R)$ be the R-module homomorphism defined by

$$\partial_{V,k}(e^{k+1}_\alpha) = \sum_{\beta \in e_{V,k}} [e^{k+1}_\alpha : e^k_\beta] e^k_\beta$$

for each e^{k+1}_α.

If $e, e' \in e_{V,k}$, we write $e' \prec e$ if both $(e,d) \in V_k$ and e' lies in the boundary of d. Let $\leq_{e,k}$ be the relation on $e_{V,k}$ given by the reflexive, transitive extension of $\prec$.

Lemma 3.2. *In the notation of Definition 3.1, $\leq_{e,k}$ is a partial order on $e_{V,k}$.*

Proof. It suffices to show that $\leq_{e,k}$ is antisymmetric. Suppose for a contradiction that this is not the case; this implies that there is a sequence

$$e_1 \prec e_2 \prec \cdots \prec e_k \prec e_{k+1} = e_1$$

with $k \geq 2$ and $e_i \neq e_{i+1}$ for all i. Let $d_1, \ldots, d_{k+1}$ be the unique elements of $d_{V,k}$ for which $(e_i, d_i) \in V_k$. It then follows that

$$e_1, d_1, e_2, d_2, \ldots, e_k, d_k, e_1$$

is a nontrivial closed V-path, which is the required contradiction. □

Proposition 3.3. *Maintain the notation of Definition 3.1, and define*

$$N = |d_{V,k}| = |e_{V,k}|.$$

Denote the elements of $d_{V,k}$ by $d_1, \ldots, d_N$ in an arbitrary (but fixed) order, and denote by e_i the element of $e_{V,k}$ paired with d_i. Let $\leq_{e,k}$ be the partial order on $e_{V,k}$ defined in Lemma 3.2, and let $\leq_{d,k}$ be the order on $d_{V,k}$ induced by the matching V.

(i) *The matrix of the linear transformation $\partial_{V,k}$ relative to $d_{V,k}$ and $e_{V,k}$ is triangular with respect to $\leq_{d,k}$ and $\leq_{e,k}$ with all the diagonal entries equal to ± 1. In particular, $\partial_{V,k}$ is an isomorphism of R-modules.*

(ii) *Suppose that there exists a c with $1 \leq c \leq N$ such that whenever we have $i \leq c < j$, e_j is not a face of d_i. Then if d_p appears with nonzero coefficient in some $d \in D_{V,k}$ for some $p > c$, then e_q appears with nonzero coefficient in $\partial(d)$ for some $q > c$.*

Proof. Let $(e, d) \in V_k$. It follows from the definitions of the partial orders that

$$\partial_{V,k}(d) = \sum_{\substack{\beta \in e_{V,k} \\ \beta \leq e}} \lambda_\beta \beta.$$

It follows from Proposition 2.2 that $\lambda_\beta \in \{-1, 0, +1\}$ for all β, and also that $\lambda_e \neq 0$. This completes the proof of the first assertion of (i), and the second assertion of (i) is immediate from the first.

We now turn to (ii); write

$$d = \sum_{i=1}^{N} \lambda_i d_i.$$

Let I be the set of all p with $c < p \leq N$ satisfying the hypotheses of (ii), and let l be a $\leq_{d,k}$-maximal element of I. It follows from the definitions that e_l appears with nonzero coefficient in $\partial(d_l)$. If $1 \leq i \leq c$, then the definition of c shows that e_l appears with zero coefficient in $\partial(d_i)$, whereas if $c < i \leq N$, then e_l can only appear with nonzero coefficient in $\partial(d_i)$ if $e_l \leq_{e,k} e_i$, by the definition of $\partial(d_i)$ and Proposition 2.2. The maximality hypothesis on l then shows that the only term in the expression

$$\partial(d) = \sum_{i=1}^{N} \lambda_i \partial(d_i)$$

that contributes a coefficient of e_l is the term $\partial(d_l)$. Setting $q = l$ completes the proof. $\square$

Lemma 3.4. *Maintain the notation of Definition 3.1, and suppose that there are no unpaired k-cells. Suppose also that $e = \sum_{\alpha \in P^k} \lambda_\alpha e_\alpha$ is a k-cycle; that is, $\partial(e) = 0$.*

(i) *If $e \neq 0$, then there exists $\alpha \in P^k$ with $\lambda_\alpha \neq 0$ such that e_α is an upward matching face.*

(ii) *If $e = \sum_{\alpha \in P^k} \lambda_\alpha e_\alpha$ and $e' = \sum_{\alpha \in P^k} \mu_\alpha e_\alpha$ are two k-cycles with the property that $\lambda_\alpha = \mu_\alpha$ whenever e_α is an upward matching face, then $e = e'$.*

Proof. Suppose that $e \neq 0$, but that $\lambda_\alpha = 0$ for every upward matching face $e_\alpha \in P^k$. It follows that

$$e = \sum_{\alpha \in d_{V,k-1}} \lambda_\alpha e_\alpha.$$

Since $e \neq 0$, Proposition 3.3 (i) shows that $\partial_{V,k-1}(e) \neq 0$. The definitions of ∂ and $\partial_{V,k-1}$ then show that $\partial(e) \neq 0$. This is a contradiction, and (i) follows. Part (ii) follows from (i) by considering the cycle $e - e'$. $\square$

Lemma 3.5. *Maintain the notation of Definition 3.1, and suppose that there are no unpaired k-cells. Then the set*

$$\{\partial(d) : d \in d_{V,k}\}$$

is an irredundantly described free R-basis for the k-cycles over R, $\ker(\partial_k)$*.*

Proof. Let e be a k-cycle. Since there are no unpaired k-cells, we may write

$$e = \sum_{\alpha \in d_{V,k-1}} \lambda_\alpha e_\alpha + \sum_{\beta \in e_{V,k}} \mu_\beta e_\beta.$$

By Proposition 3.3 (i), there exists a unique element

$$f = \sum_{\gamma \in d_{V,k}} \nu_\gamma e_\gamma$$

such that $\partial_{V,k}(f) = \sum_\beta \mu_\beta e_\beta$. It follows that for each $\beta \in e_{V,k}$, the coefficient of e_β in the k-cycle $\partial(f)$ is μ_β. Applying Lemma 3.4 (ii) to the cycles e and $\partial(f)$ then shows that $\partial(f) = e$, and it follows that the set in the statement is an R-spanning set for the cycles. The freeness assertion follows by another application of Proposition 3.3 (i). $\square$

Lemma 3.6. *Let B be a free abelian group on n generators, and let A be an abelian group generated by $\{a_1, a_2, \ldots, a_n\}$. If there is a surjective group homomorphism $\phi : A \longrightarrow B$, then ϕ is an isomorphism and $\{a_1, a_2, \ldots, a_n\}$ is a free basis for A.*

Proof. Since A is generated by n elements, it is naturally a homomorphic image $\psi(X)$ of a free abelian group X on n generators; it follows that $B = \phi(\psi(X))$ is also a quotient of X.

It follows (for example by using the fact that $\mathbb{Q}$ is a flat $\mathbb{Z}$-module) that if

$$0 \longrightarrow C \longrightarrow X \longrightarrow B \longrightarrow 0$$

is a short exact sequence of abelian groups, then $\text{rank}(X) = \text{rank}(B) + \text{rank}(C)$. The hypotheses force $\text{rank}(C) = 0$ in this case, but since every subgroup of X is torsion-free, we must have $C = 0$ and $\phi \circ \psi$ is an isomorphism. It follows that ψ is injective and is an isomorphism, which completes the proof. $\square$

Theorem 3.7. *Let X be a finite regular CW complex with a -1-dimensional cell, and suppose that V is an acyclic matching on X. Let Y be a CW subcomplex of X and let V_Y be the acyclic partial matching on Y obtained by discarding all pairs of the matching V that do not entirely lie within Y. Let K_Y be the set of unpaired cells in V_Y, and let K_X be the set of cells of $X \backslash Y$ that were paired with the elements of K_Y in the original matching V. Suppose that (a) V has no unpaired cells, and that (b) the topological boundary of each cell of K_X lies entirely within Y. Then*

(i) *the image in $H_k(Y;R)$ of the set*

$$\mathcal{B}_{Y,k} = \{\partial(k) : k \in K_X \cap P^{k+1}\}$$

is an R-spanning set for $H_k(Y;R)$, and

(ii) *if $H_k(Y;\mathbb{Z})$ is free over $\mathbb{Z}$ of rank $|K_X|$, then the image of $\mathcal{B}_{Y,k}$ is a free $\mathbb{Z}$-basis for $H_k(Y;\mathbb{Z})$.*

Proof. To prove (i), we need to show that the map ∂ induces a surjective map from $\mathcal{B}_{Y,k}$ to $H_k(Y;R)$. Let y be a k-cycle of Y; we may regard y as a k-cycle of X by extension. Number the elements of $e_{V,k}$ as $e_1, e_2, \ldots, e_N$ in such a way that

$$e_{V,k} \cap Y = \{e_1, e_2, \ldots, e_c\}$$

for some $1 \leq c \leq N$. Since Y is a subcomplex of X, we may choose the numbering so that

$$d_{V,k} \cap Y = \{d_{b+1}, d_{b+2}, \ldots, d_c\}$$

for some $0 \leq b \leq c$; it follows that

$$K_X = \{d_1, d_2, \ldots, d_b\}.$$

Hypothesis (b) shows that if e_j is a face of d_i for $1 \leq i \leq b$, then we must have $j \leq c$. The same is true if we have $b < i \leq c$, because Y is a subcomplex of X. It follows that if $i \leq c < j$, then e_j is not a face of d_i.

Lemma 3.5 and hypothesis (a) show that there exists $x \in D_{V,k}$ such that $\partial(x) = y$; let us write

$$x = \sum_{i=1}^{N} \mu_i d_i.$$

The previous paragraph and Proposition 3.3 (ii) show that we must have $\mu_i = 0$ whenever $i > c$; that is, we have

$$x = \sum_{i=1}^{c} \mu_i d_i.$$

If we define

$$x' = \sum_{i=1}^{b} \mu_i d_i,$$

it follows by hypothesis (b) that $\partial(x')$ is a cycle in Y. If $b < i \leq c$ then d_i lies in Y; this means that $\partial(d_i)$ is a boundary in Y and that $\partial(x)$ and $\partial(x')$ correspond to the same homology class in $H_k(Y; R)$. This proves part (i).

In the special case $R = \mathbb{Z}$, we note that the $\mathbb{Z}$-spanning set given in (i) has cardinality $|K_X|$. Part (ii) then follows from Lemma 3.6. $\square$

Remark 3.8. The hypotheses of Theorem 3.7 together with Theorem 2.4 (ii) show that X must be contractible.

4. The half cube

An n-dimensional *(Euclidean) polytope* Π_n is a closed, bounded, convex subset of $\mathbb{R}^n$ obtained by intersecting finitely many closed half-spaces associated to hyperplanes. We will assume that the set of hyperplanes is taken to be minimal. The part of Π_n that lies in one of the hyperplanes is called a *facet*, and each facet is an $(n-1)$-dimensional polytope. A polytope is homeomorphic to an n-ball (which follows, for example, from [**13**, Lemma 1.1]), and the boundary of the polytope, which is equal to the union of its facets, is identified with the $(n-1)$-sphere by this homeomorphism.

Iterating this construction gives rise to a set of k-dimensional polytopes Π_k (called *k-faces*) for each $0 \leq k \leq n$. The elements of Π_0 are called *vertices* and the elements of Π_1 are called *edges*. It is not hard to show that a polytope is the convex hull of its set of vertices, and that the boundary of a polytope is precisely the union of its k-faces for $0 \leq k < n$. What is less obvious, but still true [**17**, Theorem 1.1], is that the convex hull of an arbitrary finite subset of $\mathbb{R}^n$ is a polytope in the above sense. It follows that a polytope is determined by its vertex set, and we write $\Pi(V)$ for the polytope whose vertex set is V. Recalling the vertex sets V_n, V_n^e and V_n^o from the Introduction, we see that $\Pi(V_n)$ is an n-dimensional hypercube, and the half cube Γ_n is (by definition) $\Pi(V_n^e)$.

The *dimension* of a face is the dimension of its affine hull. An *automorphism* of a polytope is an isometry of its affine hull that fixes the polytope setwise. The *interior* of a face refers to its interior with respect to the induced topology on its affine hull.

The Coxeter group $W(D_n)$ is a subgroup of the group of geometric automorphisms of the half cube Γ_n, and is the full automorphism group if $n > 4$. It is generated by a set of n involutions $\{s_1, s_{1'}, s_2, s_3, \ldots, s_{n-1}\}$ which act on the set V_n: the $(n-1)$ generators s_i act by permuting the coordinates by the transposition $(i, i+1)$, and the generator $s_{1'}$ acts by the transposition $(1, 2)$ followed by a sign change on the first and second coordinates. The group $W(D_n)$ has order $2^{n-1}n!$, and acts on V_n by the subgroup of signed permutations that effect an even number of sign changes. This induces an action on $W(D_n)$ on $\mathbb{R}^n$ by orthogonal transformations fixing the half cube setwise.

Definition 4.1. Let $n \geq 4$ be an integer, and let $\mathbf{n} = \{1, 2, \ldots, n\}$.

Let $v' \in V_n^o$ and $S \subseteq \mathbf{n}$. We define the subset $K(v', S)$ of V_n^e by the condition that $v \in K(v', S)$ if and only if there exists $i \in S$ such that v and v' differ only in the i-th coordinate.

Let $v \in V_n^e$ and let $S \subseteq \mathbf{n}$. We define the subset $L(v, S)$ of V_n^e by the condition that $v' \in L(v, S)$ if and only if for all $i \notin S$, v and v' agree in the i-th coordinate. The set S is characterized as the set of coordinates at which not all points of $L(v, S)$ agree.

We call the set S in a face of the form $\Pi(K(v',S))$ or $\Pi(L(v,S))$ the *mask* of the face.

The k-faces of the half cube were classified in [**10**].

Theorem 4.2 [**10**]. *The k-faces of Γ_n for $k \leq n$ are as follows:*

(i) 2^{n-1} *0-faces (vertices) given by the elements of V_n^e;*

(ii) $2^{n-2}\binom{n}{2}$ *1-faces $\Pi(K(v',S))$, where $v' \in V_n^o$ and $|S| = 2$;*

(iii) $2^{n-1}\binom{n}{3}$ *simplex shaped 2-faces $\Pi(K(v',S))$, where $v' \in V_n^o$ and $|S| = 3$;*

(iv) $2^{n-1}\binom{n}{k+1}$ *simplex shaped k-faces $\Pi(K(v',S))$, where $v' \in V_n^o$ and $|S| = k+1$ for $3 \leq k < n$;*

(v) $2^{n-k}\binom{n}{k}$ *half cube shaped k-faces $\Pi(L(v,S))$, where $v \in V_n^e$ and $|S| = k$ for $3 \leq k \leq n$.*

Furthermore, two faces are conjugate under the action of $W(D_n)$ if and only if they have the same dimension and the same shape.

Proof. The classification of the k-faces is given in [**10**, Theorem 2.3.6], and the classification of the orbits under the action of $W(D_n)$ is given in [**10**, Theorem 4.2.3 (ii)]. □

The unique n-face in (v) above corresponds to the interior of the polytope. The k-faces assemble naturally into a regular CW complex, C_n.

Remark 4.3. The proof of Theorem 4.2 given in [**10**] is not optimal. A shorter proof of this result can be obtained by using Casselman's theorem [**5**, Theorem 3.1], which Casselman attributes to Satake and Borel–Tits. The latter result gives an explicit set of $W(D_n)$-orbit representatives of the k-faces of the half cube for each k.

In order to describe a complete acyclic matching on the faces of the half cube, it will be helpful to parametrize the faces in terms of certain sequences.

Definition 4.4. We denote a coordinate of $+1$ by the digit 0, and a coordinate of -1 by the digit 1. A face of type $\Pi(K(v',S))$ is denoted by replacing the digits in v' corresponding to coordinates in S by underlined symbols. A face of type $\Pi(L(v, S))$ is denoted by replacing the digits in v corresponding to S by asterisks. This notation is ambiguous for the 1-dimensional faces; we consider them to be faces of type $\Pi(K(v', S))$ in which the symbol associated to the rightmost coordinate in S is a $\underline{0}$.

Example 4.5.

(i) The vertex $(-1, -1, -1, +1, -1, +1, +1)$ corresponds to the sequence 1110100.

(ii) The simplex shaped face $\Pi(K(v', S))$ with

$$v' = (+1, -1, -1, +1, -1, +1, +1)$$

and $S = \{1, 3, 6, 7\}$ is denoted by the sequence $\underline{0}1\underline{1}01\underline{00}$. By toggling each coordinate in S in turn, we find the set of vertices of this face; these vertices correspond to the sequences 1110100, 0100100, 0110110, 0110101.

(iii) The half cube shaped face 010$**$1$*$010 is the convex hull of the 2^3 points obtained by filling in the asterisks with 0s and 1s in such a way that the total number of 1s is even, i.e., 0100011010, 0100110010, 0101010010, 0101111010. This face is equal to $\Pi(L(v, S))$, where v is any of the four points corresponding to the sequences listed, and $S = \{4, 5, 7\}$.

(iv) The convex hull of the pair of vertices 1110100 and 0100100 is a 1-dimensional face. This could potentially be denoted either by $\underline{1}1\underline{0}0100$ or by $\underline{0}1\underline{1}0100$, but only the first of these is correct according to Definition 4.4.

Similarly, the convex hull of the pair of vertices 0110110 and 0110101 is denoted by $01101\underline{00}$, rather than $01101\underline{11}$.

We may now describe an explicit matching on the faces of the half cube, together with the empty face, $\emptyset$. Let F be one of these faces, let S be its mask, and let $\mathbf{x}$ be the sequence associated to F by Definition 4.4. We denote the face matched with F by F', and the sequence associated with F' by $\mathbf{y}$.

(1) If F is a half cube shaped face with $\dim(F) \geq 3$, and $\mathbf{x}$ contains a 1 to the right of S, then $\mathbf{y}$ is obtained by replacing the rightmost 1 in $\mathbf{x}$ with a $*$.

(2) If F is a half cube shaped face with $\dim(F) \geq 4$, and there is no 1 in $\mathbf{x}$ to the right of S, then $\mathbf{y}$ is obtained by replacing the rightmost $*$ in S with a 1.

(3) If F is a simplex shaped face with $\dim(F) \geq 2$, and the rightmost 1 in $\mathbf{x}$ is not underlined, then $\mathbf{y}$ is obtained by underlining the rightmost 1 in $\mathbf{x}$.

(4) If F is a simplex shaped face with $\dim(F) \geq 3$, and the rightmost 1 in $\mathbf{x}$ is underlined, then $\mathbf{y}$ is obtained by replacing the rightmost $\underline{1}$ with a 1.

(5) If F is a triangle (a simplex shaped face with $\dim(F) = 2$), and the rightmost 1 in $\mathbf{x}$ is underlined, and the entries in S (reading left to right) are $\underline{011}$ or $\underline{111}$, then $\mathbf{y}$ is obtained by replacing these three entries in by $*$.

(6) If F is a half cube shaped face with $\dim(F) = 3$, and there is no 1 to the right of S in $\mathbf{x}$, then $\mathbf{y}$ is obtained from $\mathbf{x}$ by replacing the rightmost two $*$ in $\mathbf{x}$ by $\underline{1}$, and replacing the leftmost $*$ in $\mathbf{x}$ by $\underline{0}$ or by $\underline{1}$, in such a way that the total number of 1s in $\mathbf{y}$ is odd.

(7) If F is an edge and the rightmost 1 in F is not underlined, then $\mathbf{y}$ is obtained from $\mathbf{x}$ by underlining the rightmost 1.

(8) If F is a triangle and the rightmost 1 in $\mathbf{x}$ is underlined, and it is not the case that the rightmost two entries in S are equal to $\underline{11}$, then $\mathbf{y}$ is obtained from $\mathbf{x}$ by replacing the rightmost $\underline{1}$ with a 1.

(9) If F is a vertex and $\mathbf{x}$ contains at least two 1s, then $\mathbf{y}$ is obtained from $\mathbf{x}$ by replacing the rightmost 1 in $\mathbf{x}$ by $\underline{0}$ and the second rightmost 1 in $\mathbf{x}$ by $\underline{1}$.

(10) If F is an edge and the rightmost 1 in F is underlined, then $\mathbf{y}$ is obtained from $\mathbf{x}$ by replacing both underlined entries by non-underlined 1s.

(11) The empty face $\emptyset$ is matched with the vertex $00 \cdots 0$.

Example 4.6.

(i) The faces $0{*}{*}1{*}10$ and $0{*}{*}1{*}{*}0$ are matched by rules (1) and (2).

(ii) The faces $0\underline{0}1\underline{1}10\underline{0}$ and $0\underline{0}1\underline{11}0\underline{0}$ are matched by rules (3) and (4).

(iii) The faces $0\underline{1}1\underline{1}10\underline{1}$ and $0{*}1{*}10{*}$ are matched by rules (5) and (6).

(iv) The faces $01\underline{1}01\underline{0}0$ and $01\underline{1}0\underline{10}0$ are matched by rules (7) and (8).

(v) The faces 1110010 and $11\underline{1}00\underline{0}0$ are matched by rules (9) and (10).

Lemma 4.7. *Every face (including the empty face) of the half cube Γ_n is matched to another face by one, and only one, of rules (1)–(11) above. Furthermore, the smaller face in each pair of matched faces is a codimension 1 face of the larger of the pair.*

Proof. This is a case analysis based on the classification of Theorem 4.2. Let F be a (possibly empty) face of Γ_n, and let $\mathbf{x}$ and S be the corresponding sequence and mask, respectively.

The sequences corresponding to vertices all contain an even number of 1s. If this number is zero, then the vertex is matched to $\emptyset$ by (11); otherwise, the vertex is matched to an edge by (9). Notice that if rule (9) is applied, the resulting sequence satisfies the conditions of Definition 4.4.

The sequences corresponding to edges all contain an odd total number of 1s, and have rightmost underlined entry equal to $\underline{0}$ by Definition 4.4. In particular, there must be at least one 1 (underlined or otherwise) in the sequence. If the rightmost 1 is not underlined, then F is matched to a triangle by (7); otherwise, F is matched to a vertex by (10). Notice that rule (10) in this case will produce a vertex with an even number of 1s, as required.

If F is a triangle, then $\mathbf{x}$ contains an odd (and thus nonzero) number of 1s, some of which may be underlined. If the rightmost 1 is not underlined, then F is matched to a 3-simplex by (3). If the rightmost 1 is underlined and the two rightmost underlined symbols are both 1s, then F is matched to a 3-half cube by (5); otherwise, F is matched to an edge by (8). If rule (8) is applied, the rightmost underlined symbol in the resulting edge cannot be a 1, or rule (5) would have applied instead; this satisfies the requirements of Definition 4.4.

If F is a simplex of dimension at least 3, then $\mathbf{x}$ contains an odd number of 1s; in particular, there is at least one occurrence of 1. If the rightmost such occurrence is not underlined, then F is matched to a simplex of dimension one larger by (3); if the rightmost such occurrence is not underlined, then F is matched to a simplex of dimension one less by (4).

If F is a half cube, and there is a 1 in $\mathbf{x}$ to the right of S, then F is matched to a higher dimensional half cube by (1). Suppose there is no such 1. If F has dimension at least 4 (respectively, dimension equal to 3), then F is matched to a half cube of dimension one lower by (2) (respectively, (6)). □

Lemma 4.8. *Let F_1 and F_2 be faces of Γ_n. Rule (1) (respectively, (3), (5), (7), (9)) sends face F_1 to the face F_2 if and only if rule (2) (respectively, (4), (6), (8), (10)) sends F_2 to F_1.*

Proof. It is immediate from the definitions that rules (1) and (2) are inverses of each other, restricted to the appropriate domain and codomain. The same is true for rules (3) and (4).

Observe that if F is a 3-dimensional half cube shaped face, then F contains four triangular faces. These are obtained by replacing the $*$ in the mask of S by occurrences of $\underline{1}$ or $\underline{0}$ in such a way that the total number of 1s and $\underline{1}$s in the resulting sequence is odd. Precisely one of these four triangular faces has a mask of the form $\underline{011}$ or $\underline{111}$. These observations imply that rules (5) and (6) are inverses of each other.

Note that if F is an edge, then Definition 4.4 requires the rightmost underlined symbol in F to be a $\underline{0}$. If rule (7) is applicable to F, and S is the mask of the resulting triangle, then the rightmost two entries in S will be $\underline{01}$. On the other hand, if rule (8) is applicable to a triangle F' with mask S', then the rightmost two entries in S will be $\underline{01}$, and the rightmost underlined entry of the resulting edge will be $\underline{0}$. These observations show that rules (7) and (8) are inverses of each other.

Observe that if F is an edge and the rightmost 1 in F is underlined, then the restrictions of Definition 4.4 mean that this rightmost 1 is the leftmost of the two underlined symbols, and furthermore, that the rightmost underlined symbol is a $\underline{0}$. Since F contains an odd number of occurrences of 1 or $\underline{1}$, replacing both these entries with occurrences of 1 as in rule (10) will produce an even total number of 1s. Conversely, any vertex not equal to $00\cdots0$ contains an even number of 1s; in particular, it contains at least two occurrences. These observations show that rules (9) and (10) are inverses to each other, which completes the proof. □

The following result is an immediate consequence of lemmas 4.7 and 4.8.

Proposition 4.9. *Rules (1)–(11) define a complete matching on the set of faces of the half cube* Γ_n*, including the empty face.* □

5. Proof that the matching is acyclic

In §5, we will show that the complete matching of Proposition 4.9 is acyclic in the sense of §2. In order to do this, it is convenient to introduce a certain statistic on the faces of types (i)–(iv) in the classification of Theorem 4.2; we will call such faces *faces of type* K.

Definition 5.1. Let F be a face of type K, and let $\mathbf{s}$ be the sequence associated to F by Definition 4.4. Let $S' \subseteq \mathbf{n}$ be the (possibly empty) set of indices at which $\mathbf{s}$ has occurrences of 1 or $\underline{1}$. We define the *total* of F to be

$$t(F) = t(\mathbf{s}) = \sum_{i \in S'} i.$$

We define the sequence $u(\mathbf{s}) = u(F)$ from $\mathbf{s}$ by replacing all occurrences of $\underline{0}$ (respectively, $\underline{1}$) by 0 (respectively, 1).

It is immediate $u(F) = u(F')$ implies that $t(F) = t(F')$.

Example 5.2. If F is the face with sequence $0\underline{1}11\underline{0}101$, then we have $t(F) = 2 + 3 + 4 + 6 + 8 = 23$ and $u(F) = 01110101$.

Remark 5.3.

(i) In Definition 4.4, the sequence chosen to represent a given edge is the one with the lower of the two possible totals.

(ii) In the context of rule (6) of the matching, there are four triangular faces of the 3-dimensional half cube; the one paired with the half cube is the one with the highest total.

The following result is a immediate from the classification of Theorem 4.2; it will often be used in the sequel.

Lemma 5.4.

(i) *If* $\Pi(K(v', S))$ *is a simplex shaped face of* Γ_n*, then every face of* $\Pi(K(v', S))$ *can be expressed in the form* $\Pi(K(v', S'))$ *for the same* v'*, where* $S' \subset S$ *and* $S \backslash S'$ *is a singleton.*

(ii) *If* $\Pi(L(v', S))$ *is a half cube shaped face of* Γ_n*, then every half cube shaped face of* $\Pi(L(v', S))$ *can be expressed in the form* $\Pi(L(v', S'))$ *for the same* v'*, where* $S' \subset S$ *and* $S \backslash S'$ *is a singleton.* □

Remark 5.5. Some care must be taken in using Lemma 5.4 (i) for faces of the form $\Pi(K(v', T))$ when $|T| = 2$, because in this case, there are two possible choices for v'.

Lemma 5.6. *If*

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \alpha_2, \dots, \beta_r, \alpha_{r+1} = \alpha_0$$

is a nontrivial closed V*-path of faces of* Γ_n *in which the faces* α_i *have dimension* 2*, then all the faces* α_i *and* β_i *are of the form* $\Pi(K(v', S))$ *for the same* v'*. In particular, none of the* β_i *is half cube shaped, and the sequences* $u(\alpha_i)$ *and* $u(\beta_i)$ *all coincide.*

Proof. If a face β_i is a 3-dimensional half cube, it follows from Remark 5.3 (ii) that $t(\alpha_{i+1}) < t(\alpha_i)$. In contrast, if β_i is a 3-dimensional simplex, Lemma 5.4 (i) shows that $u(\alpha_i) = u(\beta_i) = u(\alpha_{i+1})$, which in turn implies that $t(\alpha_{i+1}) = t(\alpha_i)$. The conclusions (i) and (ii) now follow from the requirement that $t(\alpha_{r+1}) = t(\alpha_0)$. □

Lemma 5.7. *If*

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \alpha_2, \dots, \beta_r, \alpha_{r+1} = \alpha_0$$

is a nontrivial closed V*-path of faces of* Γ_n *in which the faces* α_i *have dimension* 1*, then all the sequences* $u(\alpha_i)$ *and* $u(\beta_i)$ *all coincide.*

Proof. It follows from Remark 5.5 and Remark 5.3 (i) that we have $t(\alpha_{i+1}) \leq t(\alpha_i)$ for $0 \leq i \leq r$, with equality holding if and only if $u(\alpha_i) = u(\beta_i) = u(\alpha_{i+1})$. The requirement that $\alpha_{r+1} = \alpha_0$ forces equality to hold at every step, and the conclusion follows from this. □

Theorem 5.8. *The matching described in §4 is a complete acyclic matching on the faces of* Γ_n *(together with the empty face).*

Proof. By Proposition 4.9, it is enough to show that the matching is acyclic. By Theorem 2.4 (i), this reduces to showing that there are no nontrivial closed V-paths. Suppose for a contradiction that

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \alpha_2, \dots, \beta_r, \alpha_{r+1} = \alpha_0$$

is such a path. We will proceed by a case analysis based on $\dim(\alpha_0)$.

The fact that there is a unique face of dimension -1 rules out the possibility of $\dim(\alpha_0) = -1$.

Suppose that $\dim(\alpha_0) = 0$. Each edge β_i has exactly two vertices contained in it, and they both appear in the path. It follows from Remark 5.3 (i) that for all $0 \leq i \leq r$, $t(\alpha_i) < t(\alpha_{i+1})$. This is incompatible with the condition $\alpha_{r+1} = \alpha_0$, which completes the proof in this case.

Suppose that $\dim(\alpha_0) \geq 3$ and that α_0 is simplex shaped. It follows by rule (3) of the matching and Lemma 5.4 (i) that all the other faces in the path are simplex shaped, with all the matched pairs being matched by rules (3) and (4). In particular, β_0 is obtained by underlining the rightmost 1 in the sequence for α_0, and α_1 is obtained from β_0 by removing the underline from one of the other symbols. This means that the rightmost 1 in the sequence of α_1 is still underlined, and α_1 is not a candidate for input to rule (3). This is a contradiction.

If α_0 is a triangle, Lemma 5.6 shows that none of the β_i is half cube shaped. In particular, rules (5) and (6) do not play a role in the path, and we can apply the argument of the above paragraph to obtain a contradiction.

If α_0 is an edge, then all the matched pairs in the path are matched by rules (7) and (8). By Lemma 5.7, the sequences $u(\alpha_i)$ and $u(\beta_i)$ all coincide. We can then copy the argument used above to deal with the case where α_i is a simplex to obtain a contradiction.

It remains to deal with the case where α_0 is a half cube shaped face with $\dim(\alpha_0) \geq 3$. If at least one of the α_i is a simplex shaped face, then we may rotate the closed path so that α_i plays the role of α_0. This has already been dealt with above, so we may assume that all the α_i are half cube shaped, and that all faces in the path are matched by rules (1) and (2).

It follows from rule (1) that β_0 is obtained by replacing the rightmost 1 by a $*$ in the sequence for α_0, and α_1 is obtained from β_0 via Lemma 5.4 (ii) by replacing one of the other occurrences of $*$ by 0 or 1. This means that α_1 has no 1 to the right of the rightmost $*$, and is not a candidate for input to rule (1). This is a contradiction and completes the proof. □

6. Homology bases for polytopal subcomplexes

In this section, we combine Theorem 5.8 with Theorem 3.7 to obtain an explicit homology basis for $C_{n,k}$.

Lemma 6.1. *Let $n \geq 4$ and let $3 \leq k < n$. Let X be the CW complex corresponding to the faces of Γ_n, including the empty face, let V be the complete acyclic matching on X given in Theorem 5.8, and let Y be the subcomplex of X corresponding to $C_{n,k}$. Then:*

(i) *X and Y satisfy the hypotheses of Theorem 3.7;*

(ii) *the unmatched faces of Y are the $(k-1)$-dimensional faces that are inputs to rules (1) or (5) of the matching; these are paired with the k-dimensional half cube shaped faces of X that are inputs to rules (2) or (6).*

Proof. We need to identify the faces of Y that are paired by V with faces in $X \backslash Y$. An inspection of the matching rules in §4 shows that these faces are the faces of Y of dimension $k-1$ satisfying the input conditions of rule (1) if $k > 3$, or rule (5) if $k = 3$. The faces of $X \backslash Y$ that are paired with these faces are k-dimensional half cubes that satisfy the input conditions of rule (2) if $k > 3$, or rule (6) if $k = 3$; this proves (ii).

The faces of a k-dimensional half cube shaped face are $(k-1)$-dimensional, and are either simplices or half cubes. All such faces are contained in Y. This shows that condition (b) of Theorem 3.7 holds, and condition (a) holds by the completeness of the matching V, completing the proof of (i). □

Theorem 6.2. *Let $n \geq 4$ and $3 < k < n$, and let B be the set of k-dimensional half cube shaped faces of Γ_n whose sequences have no 1 to the right of the rightmost occurrence of $*$.*

(i) *A basis for the $(k-1)$-st homology of $C_{n,k}$ is given by the images under the boundary map of the faces in B.*

(ii) *The $(k-1)$-st Betti number of $C_{n,k}$ is given by*

$$\sum_{i=1}^{n} 2^{i-k} \binom{i-1}{k-1}. \tag{1.2}$$

Proof. The hypotheses of Theorem 3.7 are satisfied by Lemma 6.1 (i). By Lemma 6.1 (ii), the set $\mathcal{B}_{Y,k}$ in Theorem 3.7 consists of the k-dimensional half cube shaped faces that are inputs to rules (2) or (6), and the latter coincides with the set B by the definition of the matching.

The $(k-1)$-st reduced homology of $C_{n,k}$ is free over $\mathbb{Z}$ by [**10**, Theorem 3.3.2]. Part (i) now follows from Theorem 3.7 (ii).

For part (ii), notice that the faces of B all have sequences with precisely k occurrences of $*$, and furthermore, they all end in $*00\cdots0$. Let i denote the number of symbols including and to the left of the rightmost $*$, so that the number of symbols in the sequence $*00\cdots0$ just mentioned is $n-i+1$. To form the set of such sequences for a fixed i, the leftmost $i-1$ symbols must contain $k-1$ occurrences of $*$; the remaining $i-k$ symbols can be independently chosen to be 0 or 1. (This number will be zero unless $i \geq k$.) This gives a total of $2^{i-k}\binom{i-1}{k-1}$ choices, and summing over all possible i gives the result. □

Remark 6.3. The basis of Theorem 6.2 can be used for explicit computations involving the action of $W(D_n)$ on the integral homology of $C_{n,k}$. In this case, the incidence numbers may be computed using the combinatorics of Coxeter groups.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant Number 0905768.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[**1**] H. Barcelo and S. Smith, *The discrete fundamental group of the order complex of B_n*, J. Algebraic Combin. **27** (2008), 399–421.
[**2**] A. Björner, *Random walks, arrangements, cell complexes, greedoids, and self-organizing libraries*, Building Bridges, vol. 19, Springer-Verlag, Berlin, 2008, pp. 165–203.
[**3**] A. Björner and V. Welker, *The homology of "k-equal" manifolds and related partition lattices*, Adv. Math. **110** (1995), 277–313.
[**4**] H. Bruggesser and P. Mani, *Shellable decompositions of cells and spheres*, Math. Scand. **29** (1971), 197–205.
[**5**] W. Casselman, *Geometric rationality of Satake compactifications*, Algebraic groups and Lie groups, Cambridge University Press, Cambridge, UK, 1997, pp. 81–103.
[**6**] M. Chari, *On discrete Morse functions and combinatorial decompositions*, Discrete Math. **217** (2000), 101–113.
[**7**] R. Forman, *Morse Theory for cell complexes*, Adv. Math. **134** (1998), 90–145.
[**8**] R. Forman, *A user's guide to discrete Morse theory*, Séminaire Lotharingien de Combinatoire **48** (2002).
[**9**] R. Geoghegan, *Topological Methods in Group Theory*, Springer, New York, 2008.
[**10**] R.M. Green, *Homology representations arising from the half cube*, Adv. Math. **222** (2009), 216–239.
[**11**] R.M. Green, *Homology representations arising from the half cube, II*, J. Combin. Theory Ser. A **117** (2010), 1037–1048.
[**12**] J.T. Harper, *Homology representations arising from a hypersimplex*, Ph.D. thesis, University of Colorado Boulder, 2011.
[**13**] J.R. Munkres, *Elements of algebraic topology*, Addison-Wesley, Menlo Park, CA, 1984.
[**14**] J.J. Rotman, *An introduction to algebraic topology*, Springer-Verlag, New York, 1988.
[**15**] M. Shattuck and T. Waldhauser, *Proofs of some binomial identities using the method of last squares* (preprint; `arXiv:1107.1063`).

[**16**] N.J.A. Sloane, *The on-line encyclopedia of integer sequences* (2011) (available online at `http://www.research.att.com/ njas/sequences/`).

[**17**] G.M. Ziegler, *Lectures on polytopes*, Springer-Verlag, New York, 1995.