Computing homology groups: Computing simplicial homology groups for a finite Δ-complex is, in the end, a matter of straightforward linear algebra. First the punchline! Start with the chain complex

$$\cdots \to C_{n+1}(X) \xrightarrow{\partial_{n+1}} C_n(X) \xrightarrow{\partial_n} C_{n-1}(X) \to \cdots$$

for the simplicial homology of a finite Δ-complex X, with matrices Δ_n representing the boundary maps ∂_n (in the standard bases given by the $(n+1)$- and n-simplices). We can, by row and column operations over $\mathbb{Z}$ (that is, we add integer multiples of a row or column to another, permute rows or columns, and multiply a row or column by -1) put the matrix $\Delta_n = (d_{ij})$ into *Smith normal form* $A_n = (a_{ij})$, which means that its entries are 0 except along the "diagonal" $a_{ii} = b_i$; and further, $b_i | b_{i+1}$ for all i (where everything is assumed to divide 0 and 0 divides only 0). We sketch the proof below.

Then if we let $M_n =$ the number of columns of zeros of A_n and $m_n =$ the number of *non-zero* rows of A_{n+1} [note the change of subscript!], and $b_1, \dots b_k$ the non-zero diagonal entries of A_{n+1} [note that $m_n = k$], then

$$H_n(X) \cong \mathbb{Z}^{M_n - m_n} \oplus Z_{b_1} \oplus \cdots \oplus Z_{b_k}$$

The point is that the matrices A_n are representatives of the boundary maps ∂_n, with respect to some very carefully chosen (and compatible) bases for $C_n(X)$. Essentially, $Z_n(X) \cong Z^{M_n}$ with basis the last M_n elements of our basis for $C_n(X)$, and $B_n(X) \cong b_1\mathbb{Z} \oplus \cdots \oplus b_k\mathbb{Z}$, mapping into the last k coordinates.

Specifically, we show that we can decompose each $C_n(X) = U_n \oplus V_n \oplus W_n$, where $Z_n(X) = \{0\} \oplus V_n \oplus W_n$ (so ∂_n is injective on U_n), and $\partial_n(U_n) \subseteq W_{n-1}$ (with finite index). Furthermore, we may choose bases so that the matrix of $\partial_{n+1} : U_{n+1} \to W_n$ is the diagonal matrix $\mathrm{diag}(b_1, \ldots , b_k)$. Note that, choosing bases for the V_n to fill out a basis for $C_n(X)$ (using the bases for U_n and W_n implied by the above statement), the matrix for ∂_{n+1} has the block form consisting of all 0's, except for the "U_{n+1} to W_n" part, which is the diagonal matrix. (This is a slightly permuted version of Smith normal form.) Then

$$B_n(X) = \mathrm{im}(\partial_{n+1}) = b_1\mathbb{Z} \oplus \cdots \oplus b_k\mathbb{Z} \subseteq \{0\} \oplus W_n \subseteq V_n \oplus W_n$$

so $H_n(X) = Z_n(X)/B_n(X) \cong (V_n \oplus W_n)/(\{0\} \oplus b_1\mathbb{Z} \oplus \cdots \oplus b_k\mathbb{Z}) \cong V_n \oplus Z_{b_1} \oplus \cdots \oplus Z_{b_k}$. The dimension of V_n can be determined from the normal forms of ∂_{n+1} and ∂_n as $\dim(V_n) = \dim(V_n \oplus W_n) - \dim(W_n) = \dim(Z_n(X)) - \dim(U_{n+1}) =$ number of non-pivot columns of ∂_n $-$ number of pivot rows of ∂_{n+1} $=$ number of zero columns of A_n $-$ number of non-zero rows of A_{n+1}, as desired.

It remains only to show that the matrices Δ_n can be put into Smith normal form, and that this *implies* the existence of the decompositions of $C_n(X)$ as described. The basic idea is that row and column operations on a matrix can really be interpreted as choosing different (ordered) bases for the domain or codomain, to describe a linear transformation $L : V \to W$. Working with integer matrices and operations over $\mathbb{Z}$ means that we can work with bases of $\mathbb{Z}^n$. Since $Lu_i = \sum_j a_{ij}w_j$ for bases $\{u_i\}$ and $\{w_j\}$, interchanging rows or columns corresponds to interchanging basis elements in W or V, repsectively (to make the actual vectors described by the summation remain the same). Multiplication of a row or column by -1 multiplies the corresponding basis element by -1. And adding a multiple m of one row, k, to another, ℓ amounts to replacing the basis vector u_ℓ with $u_\ell + mu_k$ (since the new row describes, in terms of the w_j, the image of that vector in V); adding a multiple m of one column k to another, ℓ, is reflected in the bases as

$$a_{i1}w_1 + \cdots + a_{im}w_m = a_{i1} + \cdots a_{ik}(w_k - mw_\ell) + \cdots + (a_{i\ell} + ma_{ik})w_\ell + \cdots + a_{im}w_m$$

so it can be interpreted as a change of basis in the codomain.

Using these operations to arrive in Smith normal form is a fairly straighforward matter of induction. Given an integer matrix $A = (a_{ij})$, Let $N(A)$ be the minimum, among all non-zero entries of A, of their absolute value. If $|a_{ij}| = N(A)$ does not divide some entry $a_{kl} = b$ of A, we can by row and column operations decrease $N(A)$. If $i = k$ or $j = \ell$, then adding the right multiple of the column or row containing a_{ij} to the one containing b will yield the remainder of b upon division by a_{ij} in b's place, lowering $N(A)$. Otherwise, we may assume that both $a_{i\ell} = \alpha a_{ij}$ and $a_{kj} = \beta a_{ij}$, and then adding $1 - \alpha$ times the i-th row to the k-th yields a row with a_{ij} in the j-th spot and $a_{k\ell} + (1 - \alpha)\beta a_{ij}$ in the ℓ-th, which a_{ij} still doesn't divide. Then we apply the previous operation to reduce $N(A)$.

Since this cannot continue indefinitely $[N(A) \in \mathbb{Z}^+]$, eventually $N(A) = |a_{ij}|$ *does* divide every entry of A. By swapping rows and columns, we may assume that $N(A) = |a_{11}|$, and then by row and column operations we can zero out every other entry in the first row and column. Striking out this row and column, we have a minor matrix B satisfying $N(A)|N(B)$. Note that then this will remain true under any subsequent row or column operation on B. By induction, there are row and column operations for B (which we interpret as operations on A not using the first row or column) putting B into Smith normal form; essentially, we just continually use the process above on smaller and smaller matrices. Since $N(A)|N(B)$ at the start, this remains true throughout the row and column operations, so each diagonal entry divides the ones that come later. All other entries are eventually zeroed out.

After we have put our matrices representing the boundary maps ∂_n into Smith normal form, this provides us with bases $\{u_i^n\}, \{w_j^{n-1}\}$ for $C_n(X)$ and $C_{n-1}(X)$ so that $\partial_n u_i^n = b_{i,n} w_i^{n-1}$ for $i \leq k_n$ and $\partial_n u_i^n = 0$ for $i > k_n$. In order to make these bases "compatible", as we desire, we need to use the fact that $\partial_n \circ \partial_{n+1} = 0$. We may assume that $b_{i,n} \neq 0$ for all $i \leq k_n$ (otherwise we just shift k_n).
The point is that $\partial_{n+1} u_i^{n+1} = b_{i,n+1} w_i^n$ implies that $0 = \partial_n \partial_{n+1} u_i^{n+1} = b_{i,n+1} \partial_n w_i^n$, so $\partial_n w_i^n = 0$ for every $i \leq k_n$ (since $b_{i,n+1} \neq 0$). Also, ∂_n is injective on the span of $u_1^n, \ldots, u_{k_n}^n$, since their images, $b_{i,n} w_i^n$, are linearly independent (as non-zero multiples of basis elements). This also means that they span a subspace complementary to $Z_n(X)$. We therefore wish to choose $\{u_i^n : 1 \leq i \leq k_n\}$ as our basis for U_n and $\{w_j^n : 1 \leq j \leq k_{n+1}\}$ as our basis for W_n, and *show* that this can be extended to a basis for $C_n(X)$ by choosing vectors lying in $Z_n(X)$; the added vectors will be a basis for V_n. Once we do this, we know that in these bases our boundary maps will have the form we desire, since the added vectors v_k map to 0 under ∂_n, as do the w_j^n we've kept, and $\partial_n u_i = b_{i,n} w_i^{n-1}$. For notational convenience in what follows, we will denote the vectors u_i, w_j that we have elected *not* to keep for our basis by u_i', w_j'.

To finish building our basis for $C_n(X)$ we essentially need to show that the span W_n of the w_j's form a direct summand of $Z_n(X) = \ker(\partial_n)$, that is, the w_j form the subset of some basis for $Z_n(X)$. To show that W_n is a summand, we show that $Q_n = Z_n(X)/W_n$ is a free abelian group; choosing a basis $\{v_k + W_n\}$ for Q_n, it then follows that $\{v_k\} \cup \{w_j\}$ is a basis for $Z_n(X)$; given $z \in Z_n(X)$, $z + W_n$ can be expressed as a combination of the $v_k + W_n$; z minus that combination therefore lies in W_n, so is a combination of the w_j. That shows that they span. Writing 0 as a linear combination and then projecting to Q_n shows that the coefficients of the v_k are 0 (since the $v_k + W_n$ are linearly independent in Q_n), and so the coeficients of the w_j are zero (since the w_j are linearly independent in W_n), proving linear independence.

And to show that $Q_n = Z_n(X)/W_n$ is a free abelian group, since it is finitely generated, it suffices to show that it contains no torsion elements. That is, if $z \in Z_n(X)$ and $cz \in W_n$ for some $c \neq 0$, then $z \in W_n$. But this is immediate, really; expressing $z = \sum c_i w_i + \sum c_i' w_i'$ (which is unique, since the w_i, w_i' form a basis for $C_n(X)$), then $cz = \sum cc_i w_i + \sum cc_i' w_i' \in W_n$ implies that the $cc_i' = 0$ (since this expression is still unique and the w_i form a basis for W_n), so the $c_i' = 0$ and $z = \sum c_i w_i \in W_n$, as desired.

The fact that the u_i and $Z_n(X)$ (hence the u_i and the basis $\{v_k, w_j\}$ for $Z_n(X)$) span $C_n(X)$ follows from the fact that the u_i' all lie in $Z_n(X)$. This implies that $C_n(X) \cong \text{span}\{u_i\} \oplus \text{span}\{v_k\} \oplus \text{span}\{w_j\}$, finishing our proof.