FPSAC 2010, San Francisco, USA DMTCS proc. **AN**, 2010, 657–664

f-vectors of subdivided simplicial complexes (extended abstract)

Emanuele Delucchi[1] and Aaron Pixton[2] and Lucas Sabalka[1]

[1] *Department of Mathematical Sciences, Binghamton University, Binghamton NY 13902-6000,*
[2] *Department of Mathematics, Princeton University, Princeton NJ 08544-1000*

Abstract. We take a geometric point of view on the recent result by Brenti and Welker, who showed that the roots of the f-polynomials of successive barycentric subdivisions of a finite simplicial complex X converge to fixed values dependinig only on the dimension of X.
We show that these numbers are roots of a certain polynomial whose coefficients can be computed explicitely. We observe and prove an interesting symmetry of these roots about the real number -2. This symmetry can be seen via a nice realization of barycentric subdivision as a simple map on formal power series. We then examine how such a symmetry extends to more general types of subdivisions. The generalization is formulated in terms of an operator on the (formal) ring on the set of simplices of the complex.

Résumé. On applie un point de vue géométrique à un récent résultat de Brenti et Welker, qui ont montré que les racines des polynômes f de subdivisions barycentriques successives d'un complexe simplicial X convergent vers des valeurs fixes, ne dépendant que de la dimension de X.
On preuve que ces nombres sont en effet eux-mêmes racines d'un polynôme dont les coefficients peuvent être calculés explicitement. De plus, on observe et on démontre une symétrie particulière de ces nombres autour du numéro -2. Cette symétrie se révèle en exprimant l'opération de subdivision barycentrique par une fonction sur des séries de puissances formelles. Une symétrie pareille existe pour des méthodes de subdivision plus générales, où elle s'exprime par des operateurs sur l'anneau des sommes formelles de simplexes du complexe.

Keywords: subdivisions of simplicial complexes, f-vectors, f-polynomials

1 Motivation and setup

This is an extended abstract of our paper Delucchi et al. (2009), to which we refer for a full exposition and the proofs of the statements. Let us begin here by stating the theorem which motivated our work.

Let X be an arbitrary finite simplicial complex of dimension $d-1$, and for convenience assume that all vectors and matrices are indexed by rows and columns starting at 0. We are interested in roots of the f-*polynomial* of X, defined as follows. Let f_i^X denote the number of i-dimensional faces of X. We declare that $f_{-1}^X = 1$, where the (-1)-dimensional face corresponds to the empty face, $\emptyset$. The *face vector*, or f-*vector* of X is the vector

$$f^X := (f_{-1}^X, f_0^X, \dots, f_{d-1}^X).$$

Let $\underline{t}$ denote the column vector of powers of t, $(t^d, t^{d-1}, \dots t^0)^T$. The *$f$-polynomial* $f^X(t)$ encodes the f-vector as a polynomial:

$$f^X(t) := \sum_{j=0}^{d} f_{j-1}^X t^{d-j} = f^X \underline{t}$$

We now focus on the recent result of Brenti and Welker Brenti and Welker (2008) that motivated our investigations. Let X' denote the barycentric subdivision of X, and more generally let $X^{(n)}$ denote the n^{th} barycentric subdivision of X.

Theorem 1.1 *Brenti and Welker (2008) Let X be a d-dimensional simplicial complex. Then, as n grows, the roots of $f^{X^{(n)}}$ converge to $d-1$ negative real numbers which depend only on d, not on X.*

This theorem may be surprising at first: there is no dependence on the initial complex X, only on the dimension d. However, geometrically this makes perfect sense. Barycentrically subdividing a simplicial complex X over and over again causes the resulting complex $X^{(n)}$ to have far more cells than the original X. Because higher-dimensional cells contribute more new cells (in every dimension) upon subdividing than lower-dimensional ones, the top-dimensional cells begin to dominate in their 'number of contributions' to subdivisions.

More precisely, each of the f_{d-1}^X top-dimensional cells of X contribute the same amount of cells to $X^{(n)}$. Since these cells eventually dominate contributions from smaller-dimensional cells, the f-polynomial for $X^{(n)}$ can be approximated by f_{d-1}^X times the f-polynomial associated to the n-fold subdivision of a single top-dimensional cell, $\sigma_d^{(n)}$. Since the roots of a polynomial are unaffected by multiplication by constants, the roots of $f^{X^{(n)}}$ converge to the roots of $f^{\sigma_d^{(n)}}$ as n increases.

We will see that both these sequences converge to the roots of a specific polynomial, and these roots satisfy an interesting symmetry.

We begin by observing the effect on f-vectors of barycentric subdivision. One key observation is that barycentric subdivision multiplies f-vectors by a fixed matrix, F_d:

Definition 1.2 *Define $\mathring{f}_i^X$ to be the number of interior i-faces of X for $i \geq 0$. We set $\mathring{f}_{-1}^X = 1$ if the dimension of X is -1, and 0 otherwise. Let σ_d denote the standard $(d-1)$-dimensional simplex. Define F_d to be the $(d+1)\times(d+1)$ matrix determined by the interior $(j-1)$-faces of the subdivided i-simplex:*

$$F_d := [\mathring{f}_{j-1}^{\sigma_i'}].$$

With this notation in place, we have the following.

Corollary 1.3 *For any $n \geq 0$,*

$$f^{X^{(n)}} = f^X F_d^n.$$

Thus, to understand barycentric subdivision, we need to understand the matrix F_d. We will compute the entries in F_d more explicitly later, but for now we simply observe a formula which follows from Inclusion-Exclusion:

Lemma 1.4 *If $j > i$ then $\mathring{f}_j^{\sigma'_i} = 0$. If $j \leq i$, then*

$$\mathring{f}_j^{\sigma'_i} = \sum_{k=0}^{i} (-1^k) \binom{i}{k} f_j^{\sigma'_{i-k}}.$$

By this lemma, F_d is lower triangular with diagonal entries $\mathring{f}_i^{\sigma'_i} = f_i^{\sigma'_i} = i!$. Thus, the eigenvalues of F_d are $0!, 1!, 2!, 3!, \ldots, d!$.

2 Main results

2.1 The limit polynomial

The goal of this section is to prove that the limit values of the roots of the $f^{\sigma_d^n}$ are themselves roots of a polynomial of which we can explititely compute the roots. The geometric intuition behing this fact is obtaining by noticing that, by definition, the coefficients of $f^{\sigma_d^{(n)}}$ record the number of cells of each dimension occurring in $\sigma_d^{(n)}$. Moreover, the number of cells in each dimension is bounded by a constant times the number of top-dimensional cells. Thus, if we normalize $f^{\sigma_d^{(n)}}$ by dividing by the number of top-dimensional cells, we have coefficients which, for each k, record the density of k-cells relative to the number of top-dimensional cells. As this density is positive but strictly decreases upon subdividing, there is a limiting value for the coefficient. Thus, there should be a limiting polynomial, with well-defined roots. Let us make this precise.

By Corollary 1.3, $f^{X^{(n)}}(t) = f^X F_d^n \underline{t}$. As the greatest eigenvalue of F_d is $d!$, we normalize $f^{X^{(n)}}(t)$ by dividing by $(d!)^n$ - let $p_n^X(t)$ denote the result:

$$p_n^X(t) := \frac{1}{(d!)^n} f^{X^{(n)}}(t).$$

Note this normalization does not alter the roots. It will also often be convenient to reverse the order of the coefficients of $p_n^X(t)$, with the effect of inverting the roots of $p_n^X(t)$ (that is, the roots of $f^{X^{(n)}}(t)$) about the unit circle in the extended complex plane:

$$q_n^X(t) := t^d p_n^X(t^{-1}).$$

To take powers of F_d, we diagonalize,

$$F_d = P_d D_d P_d^{-1},$$

where D_d is the diagonal matrix of eigenvalues $0!, 1!, \ldots, d!$ and P_d is the (lower triangular) diagonalizing matrix of eigenvectors. Thus, $F_d^n = P_d D_d^n P_d^{-1}$.

Now, let $\widetilde{D}_d := \frac{1}{d!} D_d$. Let $\overline{t}$ denote the column vector $\underline{t}$ in reverse order, $\overline{t} = (t^0, t^1, \ldots t^d)^T$. For any simplicial complex X, we thus have the following equations:

$$f^{X^{(n)}}(t) = f^X P_d D_d^n P_d^{-1} \underline{t} = (d!)^n \left(f^X P_d\right) \left(\widetilde{D}_d\right)^n \left(P_d^{-1}\right) \underline{t}$$

$$p_n^X(t) = \left(f^X P_d\right) \left(\widetilde{D}_d\right)^n \left(P_d^{-1}\right) \underline{t}, \quad q_n^X(t) = \left(f^X P_d\right) \left(\widetilde{D}_d\right)^n \left(P_d^{-1}\right) \overline{t}$$

As the eigenvalues of F_d are $0!, 1!, \ldots, d!$, for large n, D_d^n is dominated by its d^{th} diagonal entry, $(d!)^n$. In the limit, the powers of the matrix $\widetilde{D}_d = \frac{1}{d!} D_d$ converge to the matrix

$$M_{d,d} := \begin{bmatrix} 0 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1 \end{bmatrix}.$$

Thus, as n grows, the polynomials p_n^X and q_n^X respectively approach the polynomials

$$p_\infty^X(t) := \left(f^X P_d\right) M_{d,d} \left(P_d^{-1}\right) \underline{t}, \quad q_\infty^X(t) := \left(f^X P_d\right) M_{d,d} \left(P_d^{-1}\right) \overline{t}$$

in the sense that each sequence converges coefficient-wise in the vector space of polynomials of degree at most d.

By Corollary 1.3 and Lemma 1.4, we know the leading and trailing coefficients of $p_n^X(t)$ and $q_n^X(t)$: $p_n^X(t) = (d!)^{-n} t^d + \ldots + f_{d-1}^X$ and $q_n^X(t) = (d!)^{-n} + \ldots + f_{d-1}^X t^d$. Hence, in the limit, $p_\infty^X(t)$ does not have 0 as a root, but has degree less than d (one root of the p_n^X diverges to $-\infty$), while $q_\infty^X(t)$ is of degree d with 0 as a root. Because the polynomials $q_n^X(t)$ converge coefficient-wise to the polynomial $q_\infty^X(t)$ of the same degree, their roots also converge (see for instance Tyrtyshnikov (1997)):

Because the matrix P_d is lower triangular and $M_{d,d}$ has only one nonzero entry in position (d, d), we have

$$\left(f^X P_d\right) M_{d,d} = c_{X,d} e_d^T,$$

where e_d is the unit vector with a 1 in the d^{th} row, and $c_{X,d}$ is a constant depending on f^X and P_d. As both f^X and P_d do not depend on the amount of subdivision n, the roots of p_∞^X and q_∞^X do not depend on the value of $c_{X,d}$, and thus do not depend on *any* coefficient of f_d^X. This leads us to the following definition:

Definition 2.1 *Define the* limit p-polynomial *and the* limit q-polynomial *by*

$$p_d(t) := e_d^T P_d \underline{t}, \qquad q_d(t) := e_d^T P_d \overline{t}.$$

To summarize:

1. The roots of $f^{X^{(n)}}(t)$ are equal to the roots of $p_n^X(t)$.
2. The roots of $q_n^X(t)$ (resp. $p_n^X(t)$ converge to the roots of $q_d(t)$ (resp. $p_d^X(t)$), and depend only on the dimension of X.
3. The coefficient of t^i in the polynomial $p_d(t)$ (resp. $q_d(t)$) is the $(d-i)^{th}$ (resp. the i^{th}) entry in last row of P_d^{-1}.

In the full paper Delucchi et al. (2009) we derive explicit formulas for the computation of the coefficients of the matrix P_d^{-1}. We reproduce the result of some of these computations in our last section here.

Using the fact, proved by Brenti and Welker, that the limit of the roots of the f-polynomial are distinct and all real, we can summarize as follows.

Theorem A *Let X be a d-dimensional simplicial complex. Then, as n increases, the roots of $f^{X^{(n)}}$ converge to the $d-1$ (distinct) roots of a polynomial $p_d(t)$, whose coefficients can be explicitely computed and depend only on d, not on X.*

2.2 Symmetry of the limit values

Our result about the symetry of the 'limit roots' is the following.

Theorem B *For any dimension d, the $d-1$ 'limit' roots are invariant under the map $x \mapsto \frac{-x}{x+1}$.*

We will prove the corresponding symmetry for the roots of q_d instead of p_d, as it becomes a mirror symmetry instead of a Möbius invariance.

Theorem 2.2 *For every dimension d,*

$$q_\infty(t) = (-1)^d q_\infty(-1-t).$$

In particular, the roots of $q_\infty(t)$ are (linearly) symmetric with respect to $-\frac{1}{2}$.

As a first step, note that Lemma 1.4 gives the following expressions.

Lemma 2.3 *Let X be a simplicial complex. The f-polynomial of its barycentric subdivision $f^{X'}(t)$ and the corresponding $q_1^X(t)$ are given by*

$$f^{X'}(t) = \sum_{j=0}^{d} \Delta^j \{f^X(l)\}_l t^{d-j}, \qquad (d!) q_1^X(t) = \sum_{k=0}^{d} \Delta^k \{q_0^X(l)\}_l t^k.$$

This prompts us to consider barycentric subdivision as a function on polynomials in t defined by

$$b : \mathbb{Z}[t] \to \mathbb{Z}[t], \quad g(t) \mapsto \sum_{k \geq 0} \Delta^k \{g(l)\}_l t^k,$$

so that, for a simplicial complex X of dimension d we have then $b(q_j^X(t)) = d! q_{j+1}^X(t)$. The function b is linear, and thus it is given by its values on monomials, which we arrange in a formal power series in the variable x over the ring $\mathbb{Z}[t]$. We thus consider a function B on the ring $\mathbb{Z}[t][[x]]$ defined as

$$B : \quad \sum_{k \geq 0} g_k(t) x^k \quad \longmapsto \quad \sum_{k \geq 0} b\big(g_k(t)\big) x^k.$$

Theorem C *In $\mathbb{Z}[t][[x]]$, barycentric subdivision satisfies the identity*

$$B(e^{tx}) = \frac{1}{1 - (e^x - 1)t}.$$

To investigate the stated symmetry, we consider the following map

$$\iota : \mathbb{Z}[t] \to \mathbb{Z}[t], \quad g(t) \mapsto g(-1-t).$$

One readily checks by explicit calculation that $\iota B(e^{tx}) = B\iota(e^{tx})$. This suffices to prove the follwing key fact.

Lemma 2.4 *The map ι is an involution, and it satisfies*

$$\iota b \iota = b.$$

Recall that barycentric subdivision has the effect on each p- and q-polynomial of multiplying on the right by F before the $\underline{t}$ and $\overline{t}$, respectively, and rescaling by dividing by $d!$. In the limit, the limit p- and q-polynomials are invariant under barycentric subdivision up to this scaling: thus $b\big(q_\infty(t)\big) = d! q_\infty(t)$. Moreover, since the eigenvalues of F are all distinct, q_∞ is characterized by this identity and by having leading coefficient f^X_{d-1}.

A computation based on Lemma 2.4 shows $b\big(q_\infty(-1-t)\big) = d!\big(q_\infty(-1-t)\big)$ and since the lead coefficient of $q_\infty(-1-t)$ is $(-1)^d f^X_{d-1}$, the stated symmetry holds.

3 Symmetry for Other Subdivision Methods

In general, given any polynomial $g(t) \in \mathbb{Z}[t]$, we can consider the polynomial $\iota g(t) = g(-1-t)$. The coefficient of t^k in $g(t)$ contributes $(-1)^k\binom{k}{j}$ times itself to the coefficient of t^j in $\iota g(t)$: this contribution is exactly the number of $(j-1)$-dimensional faces of the $(k-1)$-dimensional simplex. Thus, we can interpret ι as a map on formal sums of simplices, as follows.

Let S be the set of simplices of a given simplicial complex X with vertex set VX. We will think of every simplex $\sigma \in S$ as a subset of VX. Now we can write

$$\iota : \mathbb{Z}[S] \to \mathbb{Z}[S], \quad \sigma \mapsto (-1)^{\dim \sigma + 1} \sum_{\tau \subseteq \sigma} \tau.$$

We will identify a subdivision of X by the triple $(X, \widetilde{X}, \phi)$, where $\widetilde{X}$ is the simplicial complex subdividing X (the 'result' of the subdivision) and $\phi : \widetilde{S} \to S$ is the function associating to each simplex $\widetilde{\sigma} \in \widetilde{S}$ its support in X. Now, a subdivision $(X, \widetilde{X}, \phi)$ induces a linear map

$$b_\phi : \mathbb{Z}[S] \to \mathbb{Z}[\widetilde{S}], \quad \sigma \mapsto \sum_{\phi(\widetilde{\sigma}) = \sigma} \widetilde{\sigma}.$$

A *subdivision method* Φ is a collection of subdivisions $\Phi := \{(\sigma_n, \widetilde{\sigma}_n, \phi_n)\}_{n \geq 0}$ such that for every k-face $i_k : \sigma_k \to \sigma_m$ of the standard m-simplex, the map ϕ_k is the restriction of ϕ_m to $i_k\sigma_k$. This ensures that, given any simplicial complex X, the complex $\Phi(X)$, called *subdivision of X according to the rule* Φ is uniquely defined by requiring that every n-simplex of X is subdivided as $(\sigma_n, \widetilde{\sigma}_n, \phi_n) \in \Phi$. A subdivision method is *nontrivial in dimension* n if ϕ_k is not the identity map for some $k \leq n$. Clearly if a subdivision is nontrivial in dimension n, then ϕ_n is not the identity map.

Given a subdivision method Φ, in view of the linearity of b_ϕ for each subdivision, it makes sense to write

$$b_\Phi(\sum_{\sigma \in X} \sigma) = \sum_{\sigma \in X} b_\Phi \sigma.$$

As with the map b given by barycentric subdivision, for any subdivision method the induced map b_Φ always commutes with the map ι:

Lemma 3.1 *For any subdivision method Φ, $\iota b_\Phi = b_\Phi \iota$.*

This commutativity was the key step in proving the symmetry for the barycentric subdivision, together with the fact that F_d had a dominating eigenvalue with geometric multiplicity 1. The latter property holds for the matrix realizing any subdivision method that is nontrivial in the top dimension. We thus have the desired result.

Theorem D *For any dimension n and any subdivision method Φ which is nontrivial in dimension n, there exists a unique 'limit polynomial' $p_{n,\Phi}(t)$, such that, for any d-dimensional simplicial complex X, the roots of $f^{\Phi^k(X)}(t)$ converge to the roots of $p_{n,\Phi}(t)$ as k increases. The roots of $p_{n,\Phi}(t)$ are invariant under the Möbius transformation $x \mapsto \frac{-x}{x+1}$.*

Remark 3.2 *Since the above interpretation is on the level of formal sums of simplices, the most natural context in which to study it seems to be the Stanley-Reisner ring $\mathbb{K}[X]$, defined by any simplicial complex X and any field $\mathbb{K}$. A good introduction to these rings can be found in Stanley (1996), where some properties of the Stanley-Reisner ring of a subdivision of a simplicial complex are explored. This brings us to ask the following question.*

Question 3.3 *Is there a (multi-)complex in each dimension whose f-polynomial is related to the limit polynomials $p_\infty^X(t)$ or $q_\infty^X(t)$? More generally, is there a geometric interpretation of the coefficients or the roots of $p_\infty^X(t)$ (equivalently, $q_\infty^X(t)$)?*

Brenti and Welker raise the question of defining a general concept of "barycentric subdivision" for a standard graded algebra. We can broaden the question to involve all subdivision methods, and ask whether the formulas developed in (Delucchi et al., 2009, Section 5) can be taken as a starting point to answer this question.

3.1 Computations

Our method allows us to explicitly compute the coefficients of $p_d(t)$, of $q_d(t)$, and thus also the limit roots. We carry out these computations in our full paper. As a sample, we give the values of the roots of $q_d(t)$ for $d \leq 10$ (computations which take less than 1 second of processor time using the formulae we derive in Delucchi et al. (2009)). The roots of $q_d(t)$ are, for $d \leq 10$, approximated by:

$d=2:$	-1	0								
$d=3:$	-1	$-.5$	0							
$d=4:$	-1	$-.76112$	$-.23888$	0						
$d=5:$	-1	$-.88044$	$-.5$	$-.11956$	0					
$d=6:$	-1	$-.93787$	$-.68002$	$-.31998$	$-.06213$	0				
$d=7:$	-1	$-.9668$	$-.79492$	$-.5$	$-.20508$	$-.0332$	0			
$d=8:$	-1	$-.98189$	$-.86737$	$-.63852$	$-.36148$	$-.13263$	$-.01811$	0		
$d=9:$	-1	$-.98996$	$-.91332$	$-.73961$	$-.5$	$-.26039$	$-.08668$	$-.01004$	0	
$d=10:$	-1	$-.99437$	$-.94277$	$-.81205$	$-.61285$	$-.38715$	$-.18795$	$-.05723$	$-.00563$	0

Acknowledgements

The authors would like to thank Laura Anderson for bringing Theorem 1.1 to their attention, Tom Zaslavsky for interesting discussions, and Dennis Pixton for introducing the authors.

References

F. Brenti and V. Welker. f-vectors of barycentric subdivisions. *Math. Z.*, 259(4):849–865, 2008. ISSN 0025-5874. doi: 10.1007/s00209-007-0251-z. URL `http://dx.doi.org/10.1007/s00209-007-0251-z`.

E. Delucchi, A. Pixton, and L. Sabalka. f-vectors of subdivided simplicial complexes. arXiv:1002.3201, 13 pp., 2009.

R. P. Stanley. *Combinatorics and commutative algebra*, volume 41 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, second edition, 1996. ISBN 0-8176-3836-9.

E. E. Tyrtyshnikov. *A brief introduction to numerical analysis*. Birkhäuser Boston Inc., Boston, MA, 1997. ISBN 0-8176-3916-0.