

Growth of Algebras, Words, and Graphs

by

Harold W. Ellingsen Jr.

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Mathematics

APPROVED:



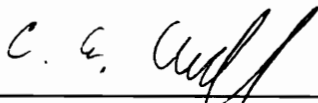
Dr. Daniel R. Farkas, Chairman



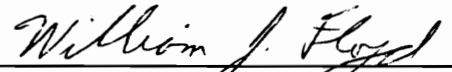
Dr. Edward L. Green



Dr. Robert F. Olin



Dr. Charles E. Aull



Dr. William J. Floyd

April 1993
Blacksburg, Virginia

Abstract

Finitely generated monomial algebras are studied. The bound in Bergman's description of algebras with GK-dimension 1 is improved and similar techniques are used to establish the equivalence of the periodicities of overlap graphs and Hilbert series coefficients.

Acknowledgements

I would like to thank Dr. Farkas, my advisor, for all of his help, guidance, and time. It was all greatly appreciated. I would like to thank Dr. Green, Dr. Olin, Dr. Aull, and Dr. Floyd for serving on my committee. Next, I want to thank my sister Vicki for all of her support. Most of all, I want to thank my mom and dad for their constant support of me in whatever I did. Finally, I would like to thank the faculty, the secretaries, and the friends I made here at Va. Tech. This was a decade I will always remember.

Contents

| | |
|-----------------------------------------------|-----------|
| Abstract | ii |
| Acknowledgements | iii |
| 1 Growth | 1 |
| 2 Words | 9 |
| 3 The Graph | 15 |
| 4 Finitely Presented Monomial Algebras | 27 |
| 4.1 Periodicity | 27 |
| 4.2 Polynomially Bounded Growth | 37 |
| 5 The Program | 40 |
| References | 44 |
| Vita | 45 |

List of Figures

| | |
|------------|----|
| Figure 3.1 | 16 |
| Figure 3.2 | 21 |
| Figure 3.3 | 22 |
| Figure 3.4 | 23 |
| Figure 3.5 | 26 |
| Figure 4.1 | 32 |
| Figure 4.2 | 35 |
| Figure 4.3 | 35 |
| Figure 4.4 | 36 |
| Figure 4.5 | 36 |
| Figure 4.6 | 36 |
| Figure 4.7 | 37 |

Chapter 1

Growth

This dissertation studies finitely generated, associative, not necessarily commutative algebras. Assume that K is a field. The algebras in question are homomorphic images of the free algebra $K \langle x_1, \dots, x_r \rangle$ on noncommuting indeterminates x_1, \dots, x_r . Suppose $\Psi : K \langle x_1, \dots, x_r \rangle \rightarrow A$ is a surjective K -algebra homomorphism. We refer to $a_1 = \Psi(x_1), \dots, a_r = \Psi(x_r)$ as generators of A ; the ideal $\text{Ker}\Psi$ is called the ideal of relations for A . A fundamental problem is to measure the “complexity” of $\text{Ker}\Psi$. One way to do this is to systematically keep track of how relations among a_1, \dots, a_r collapse to zero.

Let V be the K -vector space span of a_1, \dots, a_r in A . Define V^0 to be the field K and for $n \geq 1$, define V^n to be the K -vector space span of all the monomials in a_1, \dots, a_r of length n . Then $A = \sum_{n=0}^{\infty} V^n$. We refer to V as a generating subspace for A . Note that if $1 \in V$, then $V^n \subset V^{n+1}$ for all $n \geq 0$ and we may write $A = \bigcup_{n=0}^{\infty} V^n$. We should also note that if $V^n = V^{n+1}$ for some n , then $V^m = V^n$ for all $m \geq n$ and A is finite dimensional.

Now define a growth function $d_V : \mathbb{N} \rightarrow \mathbb{N}$ by $d_V(n) = \dim \sum_{i=0}^n V^i$. Observe that d_V is a nondecreasing function which gives us some handle on the size and nature of $\text{Ker}\Psi$. Unfortunately, d_V is an invariant of the particular presentation of

A by generators and relations and not a genuine invariant of A itself. We can fix this by using the following equivalence relation.

Definition 1.1 Let Φ be the set of all functions $f : \mathbb{N} \rightarrow \mathbb{R}$ that are eventually positive and nondecreasing. For $f, g \in \Phi$, define $f \prec g$ if and only if there exist positive integers l, m, n_0 such that $f(n) \leq lg(mn)$ for all $n \geq n_0$. Define $f \sim g$ iff $f \prec g$ and $g \prec f$.

It is easy to check that \sim is an equivalence relation on Φ . The next result tells us that d_V is an invariant of A after all, at least up to equivalence.

Proposition 1.1 [3] Let A be a finitely generated K -algebra with finite-dimensional generating subspaces V and W . Then $d_V \sim d_W$.

Proof: Since W is finite-dimensional, there is a $t \in \mathbb{N}$ such that $W \subseteq \sum_{i=0}^t V^i$. Then $W^n \subseteq \sum_{i=0}^{tn} V^i$. Thus $d_W(n) \leq d_V(tn)$ and so $d_W \prec d_V$. Similarly we have that $d_V \prec d_W$. Therefore, $d_V \sim d_W$. \square

Example 1.1 [3] If f and g are polynomials, then $f \sim g$ if and only if they have the same degree. So if the degree of f (as a function of n) is d , then $f \sim n^d$. If a and b are any real numbers greater than one, then $a^n \sim b^n$. In this case we have an equivalence of functions $a^n \sim 2^n$.

Definition 1.2 Let d_V be a growth function for the finitely generated algebra A . If $d_V \sim n^d$ for some positive integer d , then A is said to have polynomial growth. If $d_V \sim 2^n$, then A is said to have exponential growth. (Intermediate growths are possible.) In general, the Gelfand-Kirillov dimension of A is given by $\text{GK-dim} A = \limsup_{n \rightarrow \infty} \log_n d_V(n)$.

One can show that if $f, g \in \Phi$ and $f \sim g$, then $\limsup_{n \rightarrow \infty} \log_n f(n) = \limsup_{n \rightarrow \infty} \log_n g(n)$ [3]. So by Proposition 1.1, GK-dimension is an invariant of a finitely generated algebra. If A is a finitely generated algebra with growth function d_V such that $d_V \sim n^d$ then it is easy to see that $\text{GK-dim} A = d$. One of the goals of the dissertation is to re-examine a remarkable early observation about GK-dimension. According to G. Bergman, there is no finitely generated algebra whose GK-dimension is strictly between 1 and 2 [3].

In general, growth functions are very difficult to calculate. One way around this is to construct a new finitely generated algebra from the original one, whose ideal of relations is easier to understand but whose growth function is the same. The simpler algebra we shall study is called a monomial algebra. A monomial algebra is a finitely generated algebra which has a presentation by monomial relations. In other words, the algebra takes the form $K \langle x_1, \dots, x_r \rangle / (S)$ where S is a set of monomials in the letters x_1, \dots, x_r . This algebra is isomorphic to the vector space over K with basis B consisting of all monomials which are not relations under the multiplication rule

$$b_1 \bullet b_2 = \begin{cases} b_1 b_2 & \text{if } b_1 b_2 \in B \\ 0 & \text{if } b_1 b_2 \notin B \end{cases}$$

for $b_1, b_2 \in B$. So calculating the growth function for this monomial algebra will be the same as counting the number of words in B of length at most n .

We must first describe how to construct a special set S of monomials from an arbitrary finitely generated algebra A . When we use this set S for the set of relations which define a monomial algebra, it will turn out that B is the complement of S . We borrow the idea of a “leading term” for a commutative polynomial so that S will be the set of words that are “leading terms.” Recall that the free monoid M

is the set of all monomials in the letters x_1, \dots, x_r . It is traditional to refer to a monomial as a “word” in the letters x_1, \dots, x_r . This point of view gives rise to an alternate way of viewing the free algebra, in which we regard $K \langle x_1, \dots, x_r \rangle$ as the semigroup algebra of M .

The length of a word w in M is the number of letters in it, counting repetitions, and will be denoted $|w|$. The empty word, or word of length zero, will be denoted 1 . We order all the words in M as follows. First we say that $1 < x_1 < x_2 < \dots < x_r$. Let $u = x_{f(1)} \cdots x_{f(n)}$ and $v = x_{g(1)} \cdots x_{g(m)}$ be words in M . We say $u < v$ if $|u| < |v|$ or in the case $|u| = |v|$, if there is a t with $1 \leq t \leq n$, such that $x_{f(i)} = x_{g(i)}$ for $i = 1, \dots, t-1$ and $x_{f(t)} < x_{g(t)}$. We say that $u \leq v$ provided $u < v$ or $u = v$. This total ordering is called the length-lexicographic ordering of M ; one can prove that M is well-ordered by \leq . One can also check that \leq is compatible with the multiplication of the words in M . That is, if $a < b$ and u and v are words, then $uav < ubv$.

An arbitrary element x in the free algebra looks like $\sum_{i=1}^m \alpha_i m_i$, where each $\alpha_i \in K$ is nonzero and each $m_i \in M$. The tip of x , denoted $\text{TIP}(x)$, is the largest of the m_i 's under the length-lexicographic ordering. If Y is a subset of the free algebra, then $\text{TIP}(Y)$ will denote the set of all monomials that are tips of some element in Y ([1]).

Lemma 1.1 If I is an ideal of $K \langle x_1, \dots, x_r \rangle$ then $\text{TIP}(I)$ is a semigroup ideal of the free monoid M on the letters x_1, \dots, x_r .

Proof: Let $w \in \text{TIP}(I)$ and let u and v be words in M . Then there is a element $x = \alpha w + \sum_{i=1}^l \alpha_i m_i \in I$ where $\alpha, \alpha_i \in K$ with $\alpha \neq 0$ and the m_i 's are words such that each $m_i < w$. Since I is an ideal, $uxv \in I$. By the compatibility of

the order with multiplication, $\text{TIP}(uxv) = uvv$. Thus $uvv \in \text{TIP}(I)$. \square

If w is not the tip of any element in Y , then we say that w is a nontip. The set of words in M that are not tips of any element in Y will be denoted $\text{NONTIP}(Y)$, which is, of course, the complement of $\text{TIP}(Y)$ in M . This set of words has a very nice property.

Lemma 1.2 The complement of a semigroup ideal in M is closed under taking subwords. In particular, $\text{NONTIP}(I)$ has this property.

Proof: Suppose that J is a semigroup ideal of M , that $w \notin J$, and that v is a subword of w . Since w is in the semigroup ideal generated by v , we see that $v \notin J$. \square

Let $A = K \langle x_1, \dots, x_r \rangle / I$ be the given finitely generated algebra and I_T be the ideal in $K \langle x_1, \dots, x_r \rangle$ generated by $\text{TIP}(I)$. Then the monomial algebra we want to study is $G = K \langle x_1, \dots, x_r \rangle / I_T$. The connection between these two algebras will become clear with the following results.

Lemma 1.3 $\text{TIP}(I_T) = \text{TIP}(I)$.

Proof: Notice that I_T is spanned by $\text{TIP}(I)$ since $\text{TIP}(I)$ is a semigroup ideal of M . Hence $\text{TIP}(I_T) \subseteq \text{TIP}(I)$. The reverse inclusion follows since a single word is its own tip. \square

The following proposition makes an important connection between $\text{NONTIP}(I)$ and A . If $w \in M$, we will denote the image of w in A by \bar{w} .

Proposition 1.2 The image of $\text{NONTIP}(I)$ in A is a basis for A .

Proof: Let T be the image of $\text{NONTIP}(I)$ in A . Assume that $I \neq K < x_1, \dots, x_r >$. Since the image of M spans A , we first show that the image of each word in M can be written as a linear combination of elements in T . Since M is a well-ordered set, we can proceed by induction on the words in M . Clearly, $\bar{1} \in T$. Now assume that the image of every word in M less than w is in $\text{span}(T)$. If $w \in \text{NONTIP}(I)$, then $\bar{w} \in \text{span}(T)$. If $w \in \text{TIP}(I)$, then $w - \sum_{i=1}^l \alpha_i m_i \in I$ where the α_i 's are scalars and the m_i 's are words in M with each $m_i < w$. Then \bar{w} is a linear combination of the images of smaller words, each of which is in the span of T by induction. Thus $\bar{w} \in \text{span}(T)$.

For linear independence, assume that $m_1 < \dots < m_l \in \text{NONTIP}(I)$ and $\sum_{i=1}^l \alpha_i \bar{m}_i = 0$ where the α_i are nonzero scalars. Then $\alpha_l \neq 0$ tells us that $m_l \in \text{TIP}(I)$. Hence the image of $\text{NONTIP}(I)$ is a basis for A . \square

Corollary 1.1 The image of $\text{NONTIP}(I)$ in G is a basis for G .

Proof: Proposition 1.2 implies that the image of $\text{NONTIP}(I_T)$ in G is a basis for G . But Lemma 1.3 tells us that $\text{NONTIP}(I_T) = \text{NONTIP}(I)$. \square

Proposition 1.3 A and G have the same growth function.

Proof: Let V be the span of the images of x_1, \dots, x_r in A and let W be the span of the images of x_1, \dots, x_r in G . Then V is a finite-dimensional generating subspace for A and W is a finite-dimensional generating subspace for G . In the proof of Proposition 1.2, the inductive step shows that the images of the words in

NONTIP(I) of length at most n is a basis for $\sum_{i=0}^n V^i$. Corollary 1.1 and that same inductive step show that the images of the words in NONTIP(I) of length at most n is a basis for $\sum_{i=0}^n W^i$. Hence, $d_V(n) = d_W(n)$ for all n . Therefore, $d_V = d_W$. \square

From this point on, when we discuss the growth properties of a finitely generated algebra, we will replace the original algebra with a monomial algebra without any comments.

If $G = K \langle x_1, \dots, x_r \rangle / (S)$ is a given monomial algebra with the distinguished basis B of nontips, let h_n denote the number of words in B with length n . The corresponding Hilbert series is the formal power series $H(t) = \sum_{n=0}^{\infty} h_n t^n$. Let V be the span of the images of x_1, \dots, x_r in G and d_V be the growth function for G . Then it is easy to check that d_V and H are related by $\sum_{n=0}^{\infty} d_V(n) t^n = \frac{H(t)}{1-t}$. It is a far less trivial fact that $H(t)$ is a rational function [4].

The major problem examined in this thesis is the behavior of the coefficients of a Hilbert series. More precisely, what sequences h_0, h_1, h_2, \dots can appear for a finitely generated monomial algebra? In his proof that no finitely generated algebra has GK-dimension between 1 and 2, Bergman partially answered this question by showing that if $h_d \leq d$ for some $d \geq 2$, then $h_m \leq d^3$ for all $m \geq d$. This result links the following elementary observations:

1. If the GK-dimension of a monomial algebra is less than 2, then $h_d \leq d$ for some $d \geq 1$.
2. If h_m is bounded, then the GK-dimension is at most 1.

The first is a consequence of the fact that if $h_n > n$ for all n , then $d_V(n) = \sum_{i=0}^n h_i \geq \sum_{i=0}^n i = \frac{n(n+1)}{2}$. The second comes from the observation that if there is a constant

C such that $h_n \leq C$ for all n , then $d_V(n) \leq (C + 1)n$.

The main theorem we prove is an improvement of Bergman's bound: if $h_d \leq d$ for some $d \geq 2$, then $h_m \leq \left(\frac{d+1}{2}\right)^2$ for all $m \geq 2d$. Moreover, we show that this bound is sharp: given any $d \geq 2$, we construct a monomial algebra with exactly $\left\lfloor \left(\frac{d+1}{2}\right)^2 \right\rfloor$ words of length n for any $n \geq 2d$. This settles a conjecture of D. Anick.

Chapter 2

Words

In this chapter we motivate the construction of a certain graph which simplifies book-keeping on words. We first introduce some useful definitions and notation about words in a free monoid on a finite set of letters.

Definition 2.1 Let u and v be words in the letters x_1, \dots, x_t . We say that u is a prefix of v , written $u]v$, provided there is a word w such that $v = uw$ and we say that u is a suffix of v , written $u[v$, provided there is a word z such that $v = zu$.

To illustrate the above definition, let $v = xy y x x y$ be a word in the letters x and y . The prefixes of v are $1, x, xy, xyy, xy y x, xy y x x, \text{ and } v$. The suffixes of v are $1, y, xy, xxy, y x x y, y y x x y, \text{ and } v$.

Definition 2.2 Let $w = a_1 \cdots a_n$ be a word of length n . If there is a positive integer p such that $a_i = a_{i+p}$ for $i = 1, \dots, n - p$, then we say that w is a periodic word with period p . In this case we write $w = (a_1 \cdots a_p)^j a_1 \cdots a_t$ where $n = pj + t$. We call the word $a = a_1 \cdots a_p$ a base for w . We will suggestively write $w]a^\infty$.

For example, let x and y be letters. Then the word $xy y x x y$ has period 3 with base xy and the word $xy y x x y y x x y$ has period 4 with base $xy y x$. We should note

that a period and base of a word are not necessarily unique: the word $xyxxyx$ can be written two different ways, $(xyx)(xyx) = (xyxxy)(x)$. We see that w has period 3 for the base xyx as well as period 5 for the base $xyxxy$. However, we shall prove later that if w has two bases and is sufficiently long with respect to the bases, then there is an essential uniqueness.

From now on, A will denote a finitely generated monomial algebra. We will be sloppy and identify A with the span of its nontips B in the free algebra. Let $H(t) = \sum_{n=0}^{\infty} h_n t^n$ be its Hilbert series. (Recall that h_n is the number of words in B of length n .) Lemma 1.2 says that B is closed under taking subwords and, thus, a word in B is a product of smaller words in B . So suppose that we know all of the words of some length d . Because a word of length $d + 1$ has a prefix and a suffix of length d , we can relate the words of length $d + 1$ to these smaller words of length d . We would like to be able to systematically construct all of the words in B of length greater than d in some inductive way, but, as we shall see, there are too many possibilities to handle naively.

In light of Bergman's Theorem, our running hypothesis is that B contains at most d words of length d for some $d \geq 1$. In other words, $h_d \leq d$. We will refer to this hypothesis as " B is fixed at d ." If there are no words of length d , we have $h_d = 0$ and so $h_n = 0$ for all $n \geq d$. In this case, A is finite dimensional and its Hilbert series is just a polynomial whose coefficients are positive integers. To handle the case that $d = 1$, we prove a useful lemma about periodic words.

Lemma 2.1 Let u, v be words such that $|u| < |v|$, $u|v$, and $u|v$. Then $v|w^\infty$, where $v = wu$.

Proof: We must consider two cases. The first occurs when $v = uv'u$. That is, the prefix u does not overlap the suffix u . Here we take $w = uv'$.

The second case occurs when the prefix u does overlap the suffix u . We induct on $|v|$. If $|v| = 1$ then let $w = v$. Assume that the lemma is true for words of length less than $|v|$. Since there is an overlap in v , write $v = wcy$ where $u = wc = cy$. Then the word c is both a prefix and a suffix of u and $|c| < |u|$. Also, $|u| < |v|$. By the inductive hypothesis, $u]w^\infty$ where $u = wc$. Since $v = wcy$ and $u = cy$, we have $v]w^\infty$. \square

Proposition 2.1 Let A be as above. If, for some $d \geq 1$, we have $h_d = 1$ then $h_n \leq 1$ for all $n \geq d$.

Proof: Let u be the unique word of length d . If $d = 1$, then u is a single letter and so the only potential word of length n is u^n . Thus $h_n \leq 1$. Next assume that $d \geq 2$. A word v of length $d + 1$ must have u as both a prefix and a suffix which overlap. By Lemma 2.1, $v = a^{d+1}$ where a is the first letter of v . Again, there is at most one word of length $d + 1$, forcing $h_{d+1} \leq 1$. Repeating this argument shows that $h_n \leq 1$ for all $n \geq d$. \square

Now we move on to the case in which we have two words of a fixed length at least 2. That is, $h_d = 2$ for some $d \geq 2$. To find out how many words of length $d + 1$ there are in B , we check all the possible prefixes and suffixes in a word of length $d + 1$. The following theorem shows that there are at most two words of any length greater than d . Furthermore, the proof shows that in order for there to be words of lengths greater than d , patterns in the words must occur. This idea has already appeared in Proposition 2.1: there, a word of length $n \geq d + 1$ had the form a^n .

Theorem 2.1 If $h_d = 2$ for some $d \geq 2$ then $h_n \leq 2$ for all $n \geq d$.

Proof: If $h_{d+1} = 0$, then $h_n = 0$ for all $n \geq d + 1$, and hence the theorem is true. If $h_{d+1} = 1$ then, by Proposition 2.1, $h_n \leq 1$ for all $n \geq d + 1$ and hence the theorem is true. So we assume that $h_{d+1} \geq 2$ and let w_1 and w_2 be the two words of length d . A word of length $d + 1$ has two possible prefixes of length d and two possible suffixes of length d . So we can have at most four words of length $d + 1$. Here is a list of the possible cases:

| prefix | suffix |
|--------|--------|
| w_1 | w_1 |
| w_1 | w_2 |
| w_2 | w_1 |
| w_2 | w_2 |

We will consider two possibilities at a time and show that no more than two can happen.

Let u_1 and u_2 be two of the words of length $d + 1$. Suppose they have the same prefix; call it w_1 . Then they must have different suffixes, say $w_1]u_1$ and $w_2]u_2$. By Proposition 2.1, $w_1 = a^d$ and $u_1 = a^{d+1}$ for some letter a . Since $w_1]u_2$, $u_2 = a^d b$ where $b \neq a$. So $w_2 = a^{d-1}b$. Let u_3 be a third word of length $d + 1$. If it is the case that $w_1]u_3$, then $u_3 = a^d c$ for some letter c . Thus $a^{d-1}c$ is either w_1 or w_2 . So c is either a or b . Hence u_3 is either u_1 or u_2 . On the other hand, if $w_2]u_3$, then $u_3 = a^{d-1}be$ for some letter e . So $a^{d-2}be$ is one of w_1 or w_2 . Each case forces b to be a , which is impossible. Thus the only possibilities for the words of length $d + 1$ are a^{d+1} and $a^d b$. Hence $h_{d+1} \leq 2$.

Now suppose that u_1 and u_2 have different prefixes, say $w_1]u_1$ and $w_2]u_2$. They may or may not have the same suffix. Suppose they do have the same suffix, say w_1 . Then, by Proposition 2.1, $u_1 = a^{d+1}$ and $w_1 = a^d$. Since $w_1]u_2$, $u_2 = ba^d$ and $w_2 = ba^{d-1}$ where $b \neq a$. Now let u_3 be another word of length $d + 1$. If $w_1]u_3$ then $u_3 = a^d c$ for some letter c . Hence $a^{d-1}c$ has to be either a^d or ba^{d-1} . Thus $c = a$

and $u_3 = u_1$. If $w_2]u_3$, then $u_3 = ba^{d-1}e$ for some letter e . Thus $a^{d-1}e$ is either a^d or ba^{d-1} . Hence $e = a$ and $u_3 = u_2$. Thus the only possibilities for words of length $d + 1$ are a^{d+1} and ba^d . In this case, $h_{d+1} \leq 2$.


Now assume that the suffixes are different. One possibility is that $w_1[u_1$ and $w_2]u_2$. Then by Proposition 2.1, $u_1 = a^{d+1}$ and $u_2 = b^{d+1}$. Hence the only possible words of length $d + 1$ are a^{d+1} and b^{d+1} . So $h_{d+1} \leq 2$.

Finally, suppose $w_2]u_1$ and $w_1[u_2$. Then $u_1 = w_1a = xw_2$ and $u_2 = w_2b = yw_1$ for some letters a, b, x, y . Consequently, $w_1ab = xw_2b = xyw_1$. Proposition 2.1 now tells us that $w_1](ab)^\infty$. So $u_1](ab)^\infty$ as well. A similar argument shows that $w_2](ba)^\infty$ and $u_2](ba)^\infty$. Let u_3 be another word of length $d + 1$. Suppose that $w_1]u_3$. Then $u_3 = w_1c$ for some letter c . Write $w_1 = (ab)^ja^\mu$ and $w_2 = (ba)^jb^\mu$ where $\mu \in \{0, 1\}$ and $\mu + 2j = d$. The suffix of u_3 with length d is $(ba)^iba^\mu c$ where $i + \mu + 2 = d$. Since $a \neq b$, this suffix must be w_2 . Hence $u_3 = u_1$. That is, there is at most one word of length $d + 1$ with prefix w_1 . Similarly, there is at most one word of length $d + 1$ with prefix w_2 , namely u_2 . Thus $h_{d+1} \leq 2$.

We have shown that in every case, there are at most two words of length $d + 1$. Iterating the argument shows that $h_n \leq 2$ for all $n \geq d$. \square

Bergman proved that if a set of words W is closed under taking subwords and is fixed at some d , then W contains at most d^3 words of any length greater than d . Theorem 2.1 suggests that this bound can be improved. However, as we saw in the proof of Theorem 2.1, the book-keeping on the prefixes and suffixes of the words of length $d + 1$ can be quite cumbersome. It gets worse when we assume that there are three words of length three in B . We obviously need a better system to handle all of the cases.

Our first, crude attempt at organizing cases appears in the table of prefixes and suffixes in Theorem 2.1. Two prefix-suffix links seemed to be enough to determine future patterns of words. The new idea is to visualize this link by associating each word of length d with a vertex and drawing a directed edge from w_i to w_j when there is a word of length $d + 1$ whose prefix is w_i and whose suffix is w_j . For example, when we analyzed the case $w_1]u_1, w_1[u_1, w_2]u_2, w_2[u_2$ we were essentially looking at the picture



The heart of the argument was to show that these loops were “independent” in that they were powers of different letters. Visually, there are no further “connections” between these two loops.

The graph we have hinted at already appears in work of Ufnarovski ([4]). Even though his graph uses words of a specified length, we will borrow his idea of an “overlap” graph.

Chapter 3

The Graph

In this chapter we construct the “overlap” graph Γ for the words of length d in B . Then we show that there are several restrictions on the structure of Γ . It is these restrictions that allow us to calculate our bound.

Let the words in B of length d be the vertices of Γ . The assumption that B is fixed at d remains in effect. We draw an arrow from u to v provided there is a word in B of length $d + 1$ with u as its prefix and v as its suffix. Equivalently, there are letters a and b such that $ua = bv \in B$. Now we can prove a proposition that relates the number of paths of length j in Γ to the number of words of a fixed length.

Proposition 3.1 If $u \in B$ and $|u| = n \geq d$, then there is a unique path in Γ with $n - d$ arrows corresponding to u .

Proof: Write $u = y_1 \cdots y_n$ where the y_i 's are letters. Let $u_i = y_i \cdots y_{i+d-1}$ for $i = 1, \dots, n - d + 1$. Notice that the u_i 's are the d -letter subwords of u . Since $u \in B$, each $u_i \in B$. Thus the u_i 's are vertices in Γ . By construction, $u_i y_{i+d} = y_i u_{i+1}$ for $i = 1, \dots, n - d$. Hence, $u_1 \rightarrow u_2 \rightarrow \cdots \rightarrow u_{n-d+1}$ is a path in Γ . Since u is uniquely determined by its d -letter subwords, this path is unique. \square

The following is an example of an overlap graph.

Example 3.1 Let $A = K \langle a, b, c, d, e \rangle / I$ where

$$I = \langle a^2, ac, ad, ae, b^2, bd, ca, cb, c^2, cd, ce, db, dc, d^2, de, \\ ea, eb, ec, ed, e^2, bababe, dababc \rangle .$$

$\text{NONTIP}(I)$ has five words of length 5: $ababa, babab, ababe, dabab, ababc$. When $d = 5$ the graph Γ is




Figure 3.1: Example of an overlap graph

In general, let w and v be words of length d in Γ which are connected by an edge. That is, there is either an arrow from w to v or from v to w . In moving from w to v , regardless of the direction of the arrow, we delete a letter from one side of w and attach another letter to the other side. (This is well illustrated in the graph above with the three left-most vertices.) Hence v contains a subword of w of length $d - 1$.

Now consider Γ as an undirected graph. Let $u_0 - u_1 - \dots - u_r$ be a chain in Γ where the edges denote arrows in unspecified directions. If we start at u_0 , then at each vertex as we move from u_0 toward u_r at most one letter is deleted from some side of u_0 . Thus u_r contains a subword of length at least $d - r$, which is also a subword of u_0 . Note that since $d \geq r + 1$, we have $d - r \geq 1$. Thus the shared subword is nonempty.

Now we show what words on circuits look like.

Lemma 3.1 Let Γ be the overlap graph for words of length d and assume that $w_1 \rightarrow \cdots \rightarrow w_m \rightarrow w_1$ is a circuit in Γ with $m \leq d$. Then each w_i is periodic with a base of length m . Moreover, bases of length m for any two words in the circuit are cyclic permutations of each other.

Proof: By the definition of overlap, there are letters a_i and b_i such that $w_i a_i = b_i w_{i+1}$ for $i = 1, \dots, m-1$ and $w_m a_m = b_m w_1$. Since $m \leq d$ we may set $u_i = a_i \cdots a_m a_1 \cdots a_{i-1}$ and $v_i = b_i \cdots b_m b_1 \cdots b_{i-1}$ for each i . Direct calculation shows that $w_i u_i = v_i w_i$ and, so, $w_i (u_i)^r = (v_i)^r w_i$ by induction for all $r \geq 1$. By taking a very large r we see that $w_i \in v_i^\infty$. The last statement of the lemma follows because any two of the v_i 's are cyclic permutations of each other. \square

Notice that, as a consequence of the lemma, every word on a circuit C with m vertices can be uniquely reconstructed from its “circuit base” (its prefix with m letters). Since circuit bases for distinct words on C are related by cyclic shift, it immediately follows that the m cyclic shifts of a base must be distinct. Therefore, the circuit base cannot be a power of a shorter word.

The following lemma is a purely graph-theoretic result. It says that if a directed graph has two intersecting circuits, then the the first one may be perturbed so that there is a particularly simple intersection: the common vertices are consecutive vertices when regarded on subpaths of either circuit.

Lemma 3.2 Let G be a directed graph. If G has two intersecting circuits, then there are distinct circuits $p_1 \rightarrow p_2 \rightarrow \cdots \rightarrow p_k \rightarrow p_1$ and $q_1 \rightarrow q_2 \rightarrow \cdots \rightarrow q_l \rightarrow q_1$ such that for some $r \leq \min\{k, l\}$

(i) $p_i = q_i$ for $i = 1, \dots, r$ and

(ii) $\{p_{r+1}, \dots, p_k\} \cap \{q_{r+1}, \dots, q_l\} = \emptyset$.

Proof: Suppose the two original circuits are $u_1 \rightarrow \dots \rightarrow u_s \rightarrow u_1$ and $v_1 \rightarrow \dots \rightarrow v_t \rightarrow v_1$ with $s \leq t$ and $u_1 = v_1$. If $u_i = v_i$ for $1 \leq i \leq s$, we already have the desired circuits. Otherwise, look at the first time the circuits differ: suppose $u_i = v_i$ for $1 \leq i \leq m$ with $m < s$ and $u_{m+1} \neq v_{m+1}$. The u -circuit that left the v -circuit must eventually return. Let $n \geq m$ be the smallest integer such that $u_{n+1} \in \{v_1, \dots, v_t\}$. (We allow the possibility that $n = s$ by setting $u_{n+1} = u_1$. It is also possible that $n = m$.) Set $u_{n+1} = v_h$.

The first circuit we want is

$$v_h \rightarrow \dots \rightarrow v_t \rightarrow v_1 \rightarrow \dots \rightarrow v_h,$$

a shift of the original v -circuit. The second circuit is the closed path

$$v_h \rightarrow \dots \rightarrow v_t \rightarrow v_1 \rightarrow \dots \rightarrow v_m \rightarrow u_{m+1} \rightarrow \dots \rightarrow u_n \rightarrow v_h.$$

It is clear that this, indeed, is a circuit (*i.e.*, has no self-intersections) by the choice of v_h . It is also clear that these two circuits are distinct. \square

We now prove the lemmas about the uniqueness of the period and base for a periodic word.

Lemma 3.3 Let m and n be positive integers and set $g = \gcd(m, n)$. Define $\gamma : Z_m \rightarrow Z_m$ by $\gamma(j) = j + n$. Then

(i) for each j the γ -orbit

$$j, \gamma(j), \dots, \gamma^{m/g-1}(j)$$

coincides with the coset $j + gZ_m$.

(ii) every element of Z_m appears as the final member of some γ -orbit.

Proof:

(i) Note that $\gamma^t(j) = j + tn$ for all t and the number of elements in $j + gZ_m$ is $\frac{m}{g}$.

Since $j + tn = j + g\left(\frac{tn}{g}\right)$, $j + tn$ is an element of the coset $j + gZ_m$. Since the coset and orbit have the same number of elements, they must be the same set.

(ii) Let $i \in Z_m$. Then direct calculation shows that $i = \gamma^{m/g-1}(i - (m/g - 1)n)$. \square

Lemma 3.4 Let u and v be nonempty words with $|u| = m$ and $|v| = n$. Assume that $w]u^\infty$ and $w]v^\infty$. If $|w| = m + n - 1$ then u and v are powers of the same word of length $g = \gcd(m, n)$.

Proof: Since $|w| \geq |u|$ and $|w| \geq |v|$, both u and v are prefixes of w . Let $1 \leq j < m$. Then the j -th letter of u is the j -th letter of v , which is the j -th letter of w . Since w has period n , the j -th letter of u is the $(j + n)$ -th letter of w . Since $w]u^\infty$, the j -th letter of u is the $(j + n)$ -th letter of u , where $(j + n)$ is interpreted modulo m if necessary. Notice that this conclusion cannot be extended to the case $j = m$ since there is no $(m + n)$ -th letter of w . In other words, if we borrow the notation of the previous lemma and proceed along a γ -orbit then positions in the spelling of u which are consecutive numbers in the orbit are occupied by the same letter so long as we do not encounter m in the middle of the orbit. This seeming obstacle disappears if we apply (ii) of Lemma 3.3 to ensure that m is the last member of its orbit.

We conclude that there is a single letter which repeatedly appears in the j -th position of u as j ranges over a γ -orbit. Part (i) of Lemma 3.3 now makes it clear

that $u = z^{m/g}$ for some word z of length g . Recall that $v]w]u^\infty$. Hence $v]z^\infty$. Since $|z|$ is a factor of $|v|$, we see that v is a power of z as well. \square

Here is an alternative proof of Lemma 3.4. It is due to the anonymous referee for the *Journal of Algebra* who reviewed the paper based on this dissertation.

Proof: Since $w]u^\infty$ and $|w| \geq |u|$, we can write $w = uw'$ where $w']u^\infty$ and $|w'| = |v| - 1$. Thus $w']w$. Since $w]v^\infty$, w' is the first $|v| - 1$ letters of v . Hence w is the first $|u| + |v| - 1$ letters of uv . A similar argument shows that w is also the first $|u| + |v| - 1$ letters of vu . So uv and vu differ in at most their last letter. But since the number of occurrences of each letter in uv and vu is the same and $w]uv$ and $w]vu$, the last letters of uv and vu must be that same. Hence $uv = vu$.

If $|u| = |v|$, then $u = v$ and we are done. Now we assume that $|u| < |v|$ and prove the lemma by inducting on $|v|$. If $|v| = 2$, then $v = u^2$ and we are done. Now suppose that the lemma is true for $|v| \leq n$ and assume that $|v| = n + 1$. Since $|u| \leq |v|$ and $uv = vu$, we have that $u]v$. So $v = ut$ for some word t of length $|v| - |u|$. Then both u and t have length less than v and so, by the inductive hypothesis, both u and t are powers of a common subword of length $g = \gcd(|u|, |t|)$. Clearly v is also a power of this subword and it is easy to check that $g = \gcd(|u|, |v|)$. \square

Theorem 3.1 Let A be a finitely generated monomial algebra with basis of words B . Assume that for some d , B is fixed at d . If Γ is the overlap graph for words of length d , then no topologically connected component of Γ contains more than one circuit.

Proof: Let $C_1 : u_1 \rightarrow \cdots \rightarrow u_m \rightarrow u_1$ and $C_2 : v_1 \rightarrow \cdots \rightarrow v_n \rightarrow v_1$ be two circuits in some component of Γ . Assume that $m \leq n$. Suppose that the circuits C_1

and C_2 do not intersect. We know that the circuits are connected by a sequence of $r + 1$ edges for some $r \geq 0$. Without loss of generality, we will assume that the only circuit vertices on this path are u_1 and v_1 . Here is a picture to illustrate what we may have:



Figure 3.2: Two connected circuits

Note that $d \geq m + n + r$. By the observation preceding Lemma 3.1, there is a word of length $d - (r + 1)$ which is a subword of both u_1 and v_1 . Since $d - (r + 1) \geq m + n - 1$, we may assume that their common subword w has length $m + n - 1$. By Lemma 3.1, each u_i has a base p_i of length m and each v_j has a base q_j of length n . Since $|w| \geq m$ and $|w| \geq n$, the common word w contains some permutation of the bases of u_1 and v_1 as prefixes. It follows that $w]p_i^\infty$ and $w]q_j^\infty$ for some i and j . Then by Lemma 3.4 and the comment after Lemma 3.1, p_i and q_j must be the same word, so $u_i = v_j$. This means that the two circuits intersect, contrary to our initial assumption.

Now assume that C_1 and C_2 intersect at u_1 which is the same as v_1 . If $m = n$, then the circuit bases of u_1 and v_1 are the same and thus $u_i = v_i$ for all i . This forces the two circuits to be the same. Now let $m < n$. By Lemma 3.2 we can assume that the intersection of C_1 and C_2 is just a common path, say $u_1 = v_1, \dots, u_s = v_s$ for some s with $1 \leq s \leq m$.

Here is a generic picture:




Figure 3.3: Intersecting circuits

In particular, we have $d \geq m + n - s$. If $s = 1$, we could use Lemma 3.4 to directly reach a contradiction. If $2 \leq s \leq m$, we fool s into thinking it is 1 by collapsing the intersection of the circuits to one “fat” point. The path from u_1 to u_s corresponds to a word $w = u_1 a$ where $|a| = s - 1$. So $|w| \geq m + n - 1$. Since this path lies on C_1 , we conclude from Lemma 3.1 that $w]p_1^\infty$. Similarly, $w]q_1^\infty$. As we argued above, Lemma 3.4 forces $p_1 = q_1$. It follows from Lemma 3.1 that C_1 and C_2 are the same circuits, so $m = n$. Thus the case of intersecting circuits cannot occur either. \square

According to the previous theorem, each connected component of Γ has at most one circuit with paths entering and leaving it. The following theorem tells us something more.

Theorem 3.2 Fix a circuit C in Γ . If w is a word not on C , then there is at most one path from w to C or from C to w , but not both.

Proof: Assume that C has m vertices and that there is a directed path from w to C with r vertices between w and C . Let p be the first word on C that this path hits.




Figure 3.4: A path from w to C

In this case $d \geq r + m + 1$. The path from w to p represents the word $wa = bp$ where a and b are words of length $r + 1$. Let q be the suffix of w with length m . Then $|qa| = m + r + 1 \leq d$. Thus $qa[p$. Since p is periodic with a base of length m , it follows that q is a cyclic permutation of this base and that $qa]q^\infty$. Thus a is uniquely determined by w as the shortest prefix of q^∞ such that the suffix of wa with length d lies on C . Similarly, there is at most one path from C to w .

If there were paths from w to the circuit and from the circuit to w , then Γ would have two intersecting circuits. Theorem 3.1 says that this is impossible. \square

Theorem 3.3 Suppose that Γ is topologically connected with e vertices. Then for any $n \geq e$, there are at most $\left(\frac{e+1}{2}\right)^2$ directed paths with n arrows.

Proof: When Γ has no circuits, there are no paths with e or more arrows. So let C be the unique circuit in Γ with c vertices, let L be the set of l vertices on paths that enter C , and let R be the set of r vertices on paths that leave C . According to Theorem 3.2, these three sets are mutually disjoint. Then the number of vertices in Γ is at least $l + r + c$. Thus $e \geq l + r + c$. Let $n \geq e$. We count the number of paths with n arrows.

The first paths we count are those that are on the circuit C . Pick a vertex x on C . Then there is exactly one path of fixed length on C that begins at x . Consequently, there are c paths on C which have n arrows. Now we count the number of paths

that start at a vertex in L and end in C . Say y lies in L . By Theorem 3.2, there is only one path from y to C . There can be only one way to extend this path further into C . Hence we have l such paths that start in L and end in C . Similarly, there are r paths with the appropriate number of arrows which start in C and end in R . Finally we make a gross estimate for the number of paths with n arrows which begin in L and end in R , namely lr . The total number of paths P is at most $lr + l + r + c$. We bound this expression.

First we see that $P \leq lr + e$. For a fixed c , elementary calculus shows that lr achieves its maximum when $l = r = \frac{e-c}{2}$. Thus $lr + e \leq \left(\frac{e-c}{2}\right)^2 + e$. Since $1 \leq c \leq e$, the maximum occurs when $c = 1$. This tells us that there are at most $\left(\frac{e+1}{2}\right)^2$ paths with n arrows. \square

Corollary 3.1 Let Γ be an overlap graph with at most d vertices consisting of words with length d . Then there are at most $\left(\frac{d+1}{2}\right)^2$ paths with n arrows for any $n \geq d$.

Proof: Suppose Γ has k topologically connected components with d_i vertices for $i = 1, \dots, k$. So $d \geq \sum_{i=1}^k d_i$. We are done once we prove that $\sum_{i=1}^k \left(\frac{d_i+1}{2}\right)^2 \leq \left(\frac{1 + \sum_{i=1}^k d_i}{2}\right)^2$. This is equivalent to showing that $(\sum_{i=1}^k d_i^2) + k \leq (\sum_{i=1}^k d_i)^2 + 1$. This is clear if $k = 1$. Now let $k \geq 2$. Then

$$\begin{aligned} \left(\sum_{i=1}^k d_i\right)^2 + 1 &> \sum_{i=1}^k d_i^2 + \sum_{1 \leq i \neq j \leq k} d_i d_j \\ &\geq \sum_{i=1}^k d_i^2 + \sum_{1 \leq i \neq j \leq k} 1 \\ &= \sum_{i=1}^k d_i^2 + k(k-1) \end{aligned}$$

$$\geq \sum_{i=1}^k d_i^2 + k$$

□

Recall that A is a monomial algebra with a basis $\text{NONTIP}(I)$ that contains at most d words of length d . We have shown that the overlap graph Γ corresponding to A has at most $\left(\frac{d+1}{2}\right)^2$ paths with n arrows for any $n \geq d$. Since every word in $\text{NONTIP}(I)$ of length $d+j$ corresponds to a path in Γ with j arrows, we see that $\text{NONTIP}(I)$ has at most $\left(\frac{d+1}{2}\right)^2$ words of length m for any $m \geq 2d$.

Example 3.2 For an arbitrary $d \geq 2$ we construct a monomial algebra with d words of length d and $\lfloor \left(\frac{d+1}{2}\right)^2 \rfloor$ words of length n for all large n . This will show that our bound is sharp. Set

$$l = \begin{cases} \frac{d-1}{2} & \text{when } d \text{ is odd} \\ \frac{d-2}{2} & \text{when } d \text{ is even} \end{cases} \quad \text{and} \quad r = \begin{cases} l & \text{when } d \text{ is odd} \\ l+1 & \text{when } d \text{ is even} \end{cases}.$$

Note that we have $d = l + r + 1$. Consider the free algebra on the variables $a_1, \dots, a_l, b, c_1, \dots, c_r$ and let I be the ideal in the free algebra generated by

$$a_i b^k c_s, c_s a_i, a_i a_j, c_s c_t, b a_i, c_s b$$

where $i, j = 1, \dots, l$; $s, t = 1, \dots, r$; and $k = 0, \dots, d-2$. (If $d = 2$, the a_i variables do not appear.) The algebra we want is $K \langle a_1, \dots, a_l, b, c_1, \dots, c_r \rangle / I$. Its obvious basis has d words of length d : $a_i b^{d-1}$, b^d , and $b^{d-1} c_j$, where $i = 1, \dots, l$ and $s = 1, \dots, r$. The overlap graph Γ looks like




Figure 3.5: A sharp example

There are $l + r + 1 = d$ words of length d . More importantly, every directed path in Γ does actually represent a word in $\text{NONTIP}(I)$. If $n \geq 3$, then the number of directed paths with n arrows is $l + r + 1 + lr$. (This can be read off of the proof of Theorem 3.3.) When d is odd, this number is $\left(\frac{d+1}{2}\right)^2$; when d is even, it is $\left(\frac{d+1}{2}\right)^2 - \frac{1}{4}$.

Chapter 4

Finitely Presented Monomial Algebras

4.1 Periodicity

In this section we assume that A is a finitely presented monomial algebra with basis of words B that is fixed at some d . Since A is finitely presented, $A = K \langle x_1, \dots, x_r \rangle / I$ where I is an ideal of $K \langle x_1, \dots, x_r \rangle$ generated by a finite set of monomials. Choose a minimal generating set of words for I and let f be the length of its longest generator. We will assume that $d \geq f - 1$. Then, if Γ_d is the overlap graph for the words in B of length d , it follows that there is a one-to-one correspondence between paths with j arrows in Γ_d and words in B of length $d+j$ [4]. In this case we will say that Γ_d is faithful. This property holds in each of the higher overlap graphs.

Lemma 4.1 Assume that B is fixed at d . Let $j \geq 0$ and u, v be words of length d . Then there is at most one path in Γ_d with j arrows from u to v .

Proof: First, we assume that $j \geq l - p$, where p is the number of vertices on the circuit in Γ_d . Then if a path from u to v has j arrows, then this path must

contain a circuit vertex. Since there is at most one path from a vertex to the circuit and from the circuit to a vertex by Theorem 3.2, this path is unique.

Rather than assume, next, that $j < l - p$ we will suppose $j < d$. Let $u \rightarrow u_2 \rightarrow \cdots \rightarrow u_j \rightarrow v$ and $u \rightarrow v_2 \rightarrow \cdots \rightarrow v_j \rightarrow v$ be two paths in Γ_d with j arrows from u to v . Let $u_1 = v_1 = u$ and $u_{j+1} = v_{j+1} = v$. Then there exist letters a_i, b_i, y_i, z_i such that $u_i a_i = b_i u_{i+1}$ and $v_i y_i = z_i v_{i+1}$. Then $u a_1 \cdots a_j = b_1 \cdots b_j v$ and $u y_1 \cdots y_j = z_1 \cdots z_j v$. Since $j < d$, $a_1 \cdots a_j$ and $y_1 \cdots y_j$ are suffixes of v and $b_1 \cdots b_j$ and $z_1 \cdots z_j$ are prefixes of u . Thus $a_i = y_i$ and $b_i = z_i$ for $i = 1, \dots, j$. Recall that $u a_1 = b_1 u_2$ and $u y_1 = z_1 v_2$. Since the prefixes of u_2 and v_2 of length $d - 1$ are the suffix of u of the same length and $a_1 = y_1$, we have $u_2 = v_2$. Continuing in this manner, we can conclude that $u_i = v_i$ for $i = 2, \dots, j$. Hence the two paths are actually the same. \square

Lemma 4.2 Assume that B is fixed at d . Let Γ_{d+j} be the overlap graph for the words in B of length $d + j$. If Γ_d is faithful, then for each $j \geq 1$, Γ_{d+j} is also faithful.

Proof: We need only consider the case that $j = 1$. If w is a word in B of length $n \geq d + 1$, write $w = a_1 a_2 \cdots a_n$. For $i = 1, \dots, n - d$ let $q_i = a_i \cdots a_{i+d}$. Then $q_1 \rightarrow q_2 \rightarrow \cdots \rightarrow q_{n-d}$ is a path in Γ_{d+1} that corresponds to w . By Lemma 4.1, this path is uniquely determined by w .

Let $u_1 \rightarrow u_2 \rightarrow \cdots \rightarrow u_t$ be a path in Γ_{d+1} . Let p_i be the prefix of u_i of length d and let s_i be the suffix of u_i of length d . Since $u_i \in B$, then $p_i, s_i \in B$ and hence there is an arrow $p_i \rightarrow s_i$ in Γ_d . Also, by the definition of overlap, $s_i = p_{i+1}$ for $i = 1, \dots, t - 1$. Thus, $p_1 \rightarrow s_1 \rightarrow s_2 \rightarrow \cdots \rightarrow s_t$ is a path in Γ_d and the word it

corresponds to is the same as the word corresponding to the path in Γ_{d+1} . \square

We are interested in the asymptotic behavior of the sequence of overlap graphs $\{\Gamma_{d+j}\}_{j=1}^{\infty}$. The simplest question one can ask is whether the sequence is periodic: when are two of these graphs isomorphic? Let $\sum_{n=0}^{\infty} h_n t^n$ be the Hilbert series for A . We can ask a question which appears weaker: is the sequence of coefficients h_0, h_1, h_2, \dots periodic? If B is fixed at d , then B contains l words of length d for some $l \leq d$. Now we construct the adjacency matrix G for Γ_d . Its size is $l \times l$ and its (i, j) entry is 1 if there is an arrow from w_i to w_j and it is 0 otherwise. If we assume that Γ_d is connected, then Γ_d has at most one circuit (by Theorem 3.3). Thus we can arrange the vertices in such a way that if w_i and w_j are not on the circuit and there is an arrow from w_i to w_j , then $i < j$. It is not difficult to see that the sum of the entries in G^j is the number of paths with j arrows in Γ_d [4]. The following results will show that the sequence $\{G^j\}_{j=1}^{\infty}$ is eventually periodic. The first observation is immediate.

Lemma 4.3 Let $u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_p \rightarrow u_1$ be a circuit in Γ_d . Then the adjacency matrix for the circuit is of the form

$$\begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ & & & \ddots & & \\ 0 & 0 & \cdots & 0 & 1 & \\ 1 & 0 & \cdots & 0 & 0 & \end{bmatrix}$$

\square

Now G is a $(0,1)$ -matrix of the form:

$$\begin{bmatrix} S & * & * \\ 0 & C & * \\ 0 & 0 & T \end{bmatrix}$$

where S and T are strictly upper triangular matrices and C is the circuit block. A routine calculation shows that the characteristic polynomial for the circuit block is $x^p - 1$, where p is the number of vertices on the circuit. It follows that the characteristic polynomial for G is $x^l - x^{l-p}$.

Proposition 4.1 Assume that B is fixed at d and Γ_d is connected and faithful. Let $\sum_{n=0}^{\infty} h_n t^n$ be the Hilbert series for A . Then the sequence of coefficients h_0, h_1, h_2, \dots is eventually periodic and has period p where p is the number of vertices on the circuit in Γ_d .

Proof: Let G be the adjacency matrix for Γ_d as described before. By the Cayley-Hamilton Theorem, we have $G^l = G^{l-p}$. Thus $G^{l+j} = G^{l-p+j}$ for all $j \geq 0$. Consequently, the number of paths with $l+j$ arrows is the same as the number of paths with $l-p+j$ arrows. Since Γ_d is faithful, we have that $h_{d+l+j} = h_{d+l-p+j}$ for all $j \geq 0$. This shows that the sequence h_0, h_1, \dots is eventually periodic. \square

Proposition 4.1 shows that a period of the coefficients is the number of vertices on the circuit in Γ_d . However, this may not be the minimal period (although the minimal period will be a divisor of p) and the periodicity may start before the $(d+l-p)$ -th coefficient. Here is an example that illustrates these possibilities. Suppose B contains the following 7 words of length 7: $yzababa, xzababa, zababab,$

abababa, bababab, abababc, bababad. Then the adjacency matrix for Γ_7 is

$$G = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

According to Proposition 4.1, $G^7 = G^5$, the period of the coefficients is 2, and the periodicity appears by $n = 12$. But further examination of this case shows that the period is actually 1 and the periodicity begins at $n = 9$. The Hilbert series past degree 6 is $7t^7 + 8t^8 + 10t^9 + 10t^{11} + \dots$. The calculations in this example were done by a Matlab program. A detailed description can be found in Chapter 5.

Now we consider the periodicity of the sequence of overlap graphs. We start off with the same assumptions as before. Let v be a word of length $d + j$ where $j \geq 1$. Since the path corresponding to v is unique, then the first and last vertices on the path uniquely determine v . These words are the prefix and suffix of v of length d , respectively. Thus each word in B of length greater than d corresponds to an ordered pair of vertices in Γ_d . We should note that not every pair of vertices will correspond to a word; this occurs when there is no path between two vertices. If $j \geq l - p$, then the path corresponding to v contains a circuit vertex. Let v' be the word corresponding to the same path as v with an extra trip around the circuit. Then $|v'| = d + j + p$. Since v' and v begin and end at the same vertices, the correspondence $v \mapsto v'$ is one-one. We will call this function the “extra trip” function and we will denote the image of v by v' . The following example illustrates this function. Assume that we have the following words of length 4: *xaba, abab, baba, abax*. Their overlap graph may look like




Figure 4.1: An illustration of the “extra trip” function

In this case, $d = l = 4$, $p = 2$, and $l - p = 2$. Using the Matlab program, we have that $v = xababax$ is a word of length 7. Its path in Γ_4 is $xaba \rightarrow abab \rightarrow baba \rightarrow abax$. Here is the path with an extra trip around the circuit: $xaba \rightarrow abab \rightarrow baba \rightarrow [abab \rightarrow baba] \rightarrow abax$. The brackets denotes the extra trip around the circuit. Then $v' = xabababax$.

Suppose that $v \rightarrow w$ is an edge in Γ_{d+j} . Let $v_1 \rightarrow \dots \rightarrow v_{j+1}$ and $w_1 \rightarrow \dots \rightarrow w_{j+1}$ be the paths in Γ_d which correspond to v and w respectively. We will show that these paths “overlap,” that is, $v_{i+1} = w_i$ for $i = 1, \dots, j$. By the definition of overlap there exist letters a, b, a_i, b_i, y_i, z_i , $i = 1, \dots, j$ such $va = bw$, $v_i a_i = b_i v_{i+1}$, and $w_i y_i = z_i w_{i+1}$. Then $va = v_1 a_1 \dots a_j = b w_1 y_1 \dots y_j = bw$. So $b_1 v_2 a_2 \dots a_j = b w_1 y_1 \dots y_j$. Since b and b_1 are letters and $|v_2| = |w_1|$, we see that $v_2 = w_1$. Continuing this way gives us the overlap of paths. The path that corresponds to the word va is $v_1 \rightarrow \dots \rightarrow v_{j+1} \rightarrow w_{j+1}$. It is easy to see that the converse of this is also true: if the paths for v and w in Γ_d overlap, then there is an arrow $v \rightarrow w$ in Γ_{d+j} .

By Proposition 4.1, the “extra trip” function is bijective and the number of vertices in Γ_{d+j} and Γ_{d+j+p} are the same. Now in order to show that Γ_{d+j} is isomorphic to Γ_{d+j+p} , we need to show that $v \rightarrow w$ is an edge in Γ_{d+j} if and only if $v' \rightarrow w'$ is an edge in Γ_{d+j+p} . Let v correspond to the path $v_1 \rightarrow \dots \rightarrow v_{j+1}$.

If $v \rightarrow w$, then the previous paragraph shows that the path corresponding to w is $v_2 \rightarrow \cdots \rightarrow v_{j+1} \rightarrow w_{j+1}$ where w_{j+1} is the prefix of w of length d . So the paths corresponding to v' and w' are $v_1 \rightarrow \cdots \rightarrow [\text{extra trip}] \rightarrow v_{j+1}$ and $v_2 \rightarrow \cdots \rightarrow [\text{extra trip}] \rightarrow v_{j+1} \rightarrow w_{j+1}$. Clearly, these two paths overlap and thus there is an arrow $v' \rightarrow w'$ in Γ_{d+j+p} . The converse of this is also clear.

We have proved the following proposition:

Proposition 4.2 Assume that B is fixed at d and that its overlap graph Γ_d is connected, faithful, and has a circuit containing p vertices. Then the sequence of overlap graphs $\{\Gamma_{d+j}\}_{j=0}^{\infty}$ is eventually periodic with period p starting at $j = l - p$.

Now we can prove a general theorem.

Theorem 4.1 Let A be a finitely presented monomial algebra with basis of words B and Hilbert series $\sum_{n=0}^{\infty} h_n t^n$. Then the following are equivalent:

- (i) A has linearly bounded growth.
- (ii) The coefficients of the Hilbert series for A are eventually periodic.
- (iii) The sequence of overlap graphs is eventually periodic.

Proof:

i \Rightarrow **ii** Suppose A has linearly bounded growth. Then B is fixed at some d , where d is large enough to allow us to assume that Γ_d is faithful. Let $\Gamma_d^{(i)}$, $i = 1, \dots, k$ be the connected components of Γ_d where the i -th component has d_i vertices and circuit length p_i . So $d = \sum_{i=1}^k d_i$. If a $\Gamma_d^{(i)}$ has no circuit, then there can be no paths with more than $d_i - 1$ arrows in this component. So its contribution to the Hilbert series will be nothing for $n \geq d + d_i - 1$. We will assume, then,

that each component has a circuit. Let $h_n^{(i)}$ be the number of words of length n that come from paths in $\Gamma_d^{(i)}$. Then by Proposition 4.1, the $h_n^{(i)}$ have period p_i starting at $n = d + d_i - p_i$. Since $h_n = \sum_{i=1}^k h_n^{(i)}$, the periodicity of the h_n 's starts at $N = \text{maximum of } \{d + d_i - p_i\}_{i=1}^k$ and a period p will be the least common multiple of the p_i 's.

i \Rightarrow **iii** We will use the same notation and assumptions as in the proof of **i** \Rightarrow **ii**.

Also, let $\Gamma_{d+j}^{(i)}$ be the contribution of $\Gamma_d^{(i)}$ to the graph Γ_{d+j} . By Proposition 4.2 the sequence $\{\Gamma_{d+j}^{(i)}\}_{j=0}^{\infty}$ has period p_i starting at $j = d_i - p_i$. Since each Γ_{d+j} is the union of the $\Gamma_{d+j}^{(i)}$, the periodicity of the Γ_{d+j} 's begins at $J = \text{maximum of } \{d_i - p_i\}$ and a period p is the least common multiple of the p_i 's.

iii \Rightarrow **ii** Suppose that the sequence of overlap graphs is eventually periodic. There exist N and p such that Γ_n is isomorphic to Γ_{n+p} for all $n \geq N$. Then the number of vertices in Γ_n is the same as the number in Γ_{n+p} . Since the number of vertices in each Γ_m is h_m , we have that $h_n = h_{n+p}$ for all $n \geq N$.

ii \Rightarrow **i** If the h_n 's are eventually periodic, they must be bounded. Thus B is fixed at some d . This means that A has linearly bounded growth.

□

The assumption that A is finitely presented is very important to the theorem. It allows us to assume that Γ_d is faithful. Without this assumption, the powers of the adjacency matrices may have entries that do not correspond to words in B . (We should note here that the hypothesis that A be finitely presented is not used in proving **iii** \rightarrow **ii** and **ii** \rightarrow **i**.) We show how to construct examples of algebras that

have periodicity greater than the number of vertices on the circuit in Γ_d and that have no periodicity at all. Each of these algebras will be infinitely presented.

Choose $p \geq 1$. Let $A = K \langle x, y \rangle / I$, where I is the ideal generated by $xyx, yxy, \{xy^{pn}x\}_{n=0}^\infty$. The basis of words B contains three words of length three: y^3, xy^2, y^2x . We always have




Figure 4.2: A subgraph of any Γ_{d+j}

as part of any Γ_{d+j} . We have an arrow from xy^{2+j} to $y^{2+j}x$ when p does not divide $2 + j$ and there is an isolated vertex $xy^{1+j}x$ when p does not divide $1 + j$. So for, $p = 1$ we have Γ_{3+j} isomorphic to the above subgraph for all $j \geq 0$. Hence the Hilbert series is $1 + 2t + 3t^2 + 3t^3 + \dots$ and the periods of both the coefficients and the graphs is one.

If $p = 2$, and j is even, then Γ_{3+j} is




Figure 4.3: Γ_{d+j} when $p = 2$ and j is even

When j is odd, Γ_{3+j} is




Figure 4.4: Γ_{d+j} when $p = 2$ and j is odd

In this case the Hilbert series is $1 + 2t + 3t^2 + 3t^3 + 4t^4 + 3t^5 + 4t^6 + \dots$ and the periods of both the coefficients and the graphs is two.

For $p \geq 3$ and $j \equiv 0, 1, \dots, p - 3(\text{mod } p)$, Γ_{3+j} is




Figure 4.5: Γ_{d+j} when $p \geq 3$ and $j \equiv 0, 1, \dots, p - 3(\text{mod } p)$

(Note that for the j 's described above, p divides neither $j + 1$ nor $j + 2$.) If $j \equiv p - 2(\text{mod } p)$, then Γ_{3+j} is




Figure 4.6: Γ_{d+j} when $p \geq 3$ and $j \equiv p - 2(\text{mod } p)$

If $j \equiv p - 1(\text{mod } p)$, then Γ_{3+j} is




Figure 4.7: Γ_{d+j} when $p \geq 3$ and $j \equiv p - 1 \pmod{p}$

The periodicity starts when $j = 1$ and the Hilbert series is $1 + 2t + 3t^2 + 3t^3 + 4t^4 + \dots + 4t^{p+1} + 3t^{p+2} + 4t^{p+3} + \dots$. The period of the coefficients and graphs is p .

Now we construct an infinitely presented algebra whose growth is linearly bounded and whose coefficients and graphs have no period at all. Let A be the algebra $K \langle x, y \rangle / I$ where I is the ideal generated by

$$yxy, x^2, \{xy^{2^m}x\}_{m=0}^{\infty}.$$

The words of length $n \geq 4$ look like y^n , xy^{n-1} , $y^{n-1}x$ if $n - 2$ is a power of 2; if $n - 2$ is not a power of 2, we have these three words and $xy^{n-2}x$. One can see that the Hilbert coefficients have no period. Since the coefficient h_n is the number of vertices in Γ_n , the sequence of graphs cannot be periodic.

4.2 Polynomially Bounded Growth

In this section we assume that A is a finitely presented monomial algebra with basis of words B , but instead of assuming that B is fixed at d , we will assume that B contains at most $d + i$ words of length d for some $d \geq 2$ and $i \geq 0$. We now try to generalize Bergman's theorem. We can still construct the overlap graph Γ for the words of length d . If we assume that d is large enough (as in [4]), then Γ is faithful.

The theorem of this section depends on what happens to the graph obtained from Γ by removing a vertex and all arrows to and from this vertex. Let Γ^* be such

a graph, where the removed vertex is v . If $A = K \langle x_1, \dots, x_r \rangle / (m_1, \dots, m_s)$, then Γ^* is an overlap graph for the algebra $A^* = K \langle x_1, \dots, x_r \rangle / (m_1, \dots, m_s, v)$. The basis of words B^* consists of those words in B which do not contain v as a subword. A path in Γ^* is also a path in Γ and so it will correspond to a unique word w in B . The following lemma is now clear.

Lemma 4.4 Let Γ be an overlap graph for a finitely presented monomial algebra A with basis of words B . If Γ is faithful, then so is any subgraph of Γ . \square

Theorem 4.2 Let A be a finitely presented monomial algebra with basis of words B . Assume that B contains at most $d + i$ words of length d for some $d \geq 2$ and $i \geq 0$ and Γ_d is faithful. Then the growth of A is either exponential or bounded by a polynomial of degree $i + 1$.

Proof: We induct on i . When $i = 0$ the theorem above is Bergman's Theorem ([3]). Now assume that the desired result is true for i ; we assume Γ has $d + i + 1$ vertices. Suppose that the growth of A is not exponential. Since A is finitely presented, its growth is polynomial and Γ has no intersecting circuits ([4]). By [4], the degree of growth of A is the maximum of the number of circuits that lie in any simple directed path in Γ . (A path is simple if it contains no repeated vertices) How many circuits can be on a path in Γ ? Lemma 4.4 says that any subgraph of Γ is faithful and so, by the inductive hypothesis, any subgraph of Γ with $d + i$ vertices can have at most $i + 1$ circuits on a path. We reach a contradiction by supposing that Γ has $i + 3$ circuits on a simple path. Remove a vertex and its incident edges from the last circuit on this path to produce a subgraph with $i + 2$ circuits and at most $d + i$ vertices. The contradiction shows that Γ can have at most $i + 2$ circuits

on a path: the growth of A is bounded by a polynomial of degree $i + 2$. \square

To construct a counter-example for the infinitely presented case, we use Warfield's construction of an algebra with growth n^r for any $r \geq 2$ ([3]). The algebra is of the form $A = K \langle x, y \rangle / I$. Let Y_n be the set of words in x and y of length n with at least two appearances of y , e.g., $y^2 \in Y_2$ and $yxxy \in Y_3$. The ideal I will be generated by $Z = \bigcup_{n=2}^{\infty} Z_n$ where $Z_n \subseteq Y_n$. Let $d(m, Z)$ be the number of words not in Z of length at most m . Notice that $d(m, Z)$ is the growth function for A . The Z_n 's are constructed inductively in such a way that for all m , $d(m, Z) \leq m^r + 1$ and for infinitely many m , $m^r \leq d(m, Z)$. For any real number $r \geq 2$, we let $Z_2 = Y_2$, $Z_3 = Y_3$ and $Z_4 = Y_4$. Suppose that only finitely many Z_n 's are required to achieve the desired growth, say $Z = \bigcup_{n=2}^p Z_n$. We will show that the growth of A is exponential. Consider the two words x^p and yx^{p-1} . Neither lies in any Z_n . Let w be a word in the free monoid M on x^p and yx^{p-1} . Then at most one y can appear in any subword of w of length at most p . Thus $w \notin Z$. For any $n \geq 1$, there are 2^n words of length pn in M . Thus $d(pn, Z) \geq 2^n$. This means that A has exponential growth. Thus, to achieve the desired growth using this process, an infinite number of Z_n 's must be used. We get our counter-example by using Warfield's process with $r = 3$: we have an algebra with 3 words of length 2 (x^2, xy, yx) and cubic growth.

Chapter 5

The Program

While investigating the periodicity of the sequence of graphs $\{\Gamma_{d+j}\}_{j=0}^{\infty}$, we realized that analyzing the graphs by hand was becoming very complicated. We needed a better way to do computations. Matlab has a nice feature called “string comparison.” This is exactly how we determine if there is an arrow between two vertices. So creating a function to construct the adjacency matrices was easy. However, a problem arises because there are too many ways to label the vertices of a graph. We standardize the labelling of vertices. We want to order them so that the adjacency matrix is strictly upper triangular with the exception of at most one entry. First we compute the adjacency matrix for the words in no particular order. Then we classify the words as follows: “initial” (no arrows leading to them), “terminal” (no arrows leading from them), “isolated” (no arrows leading to or from them), and “other”. The words are then put in the following order: initial, other, terminal, isolated. We do a modified depth-first search [2] on the list without the isolated words (since they do not affect the words of next length). The input for the search is a vertex and the adjacency list for the graph with the new list of words. An adjacency list L for a graph is a 2-column matrix such if there is an arrow from the i -th word to the j -th word, then i, j is a row of L . The function finds all vertices that have an arrow

coming from v and then it does the search on these vertices. The search stops when all vertices have been found. The output is a vector whose i -th entry is the vertex that was found i -th in the search. We use this list to reorder the words again.

Then we pass this new list to a function which computes the words and adjacency matrices for words of lengths d up to $2d$. Between the j -th and $(j + 1)$ -st steps we reorder the words as in the modified depth-first search. Given the final output of the program, the sequence of adjacency matrices, we compare the matrices to see if they are the same or not. We then noticed that if two matrices were the same, then they were either consecutive in the sequence or a distance apart equal to the number of vertices on the circuit. This is what lead us to the conjecture for components (and then the theorem) that a period for the graphs was the number of vertices on the circuit.

The program does have its limitations. One is that it needs to have initial vertices. It does not handle unconnected graphs well. It may also be possible that there is some ordering of words that it could not straighten out as we would like it to. Now we give the algorithm for the program:

Algorithm for computing the sequence of overlap graphs

MAIN

- 1 read the words in W , the list containing at most d words of length d
- 2 $(C,I,R) = \text{LIST}(W)$
- 3 $T = \text{LOOP}(C,R)$
- 4 adjoin the isolated words to the end of T
- 5 $\text{GRAPHS}(T)$

LIST

1 form the adjacency matrix of the overlap graph for W

2 for $j = 1, \dots, \text{number of words in } W$

 if j -th column = 0 and j -th row $\neq 0$

 add j to the initial list

 else if j -th row = 0 and j -th column $\neq 0$

 add j to the terminal list

 else if j -th row = 0 and j -th column = 0

 add j to the isolated list

 else add j to the other list

3 R is the list of words in the order initial, other, terminal

I is the isolated words

C is the initial words

LOOP

1 $m = \text{number of words in } R$

$\text{ind} = 1$

$\text{dfi} = \text{zero vector of length } m$

$A = \text{adjacency list for } R$

2 while there is a v such that $\text{dfi}(v) = 0$, do

$F(1) = v$

$\text{MDFS}(v)$

 if the length of $F > 1$

 adjoin all but the first entry of F to C

3 $T =$ words in R reordered according to C and with the isolated words adjoined to the end

MDFS

1 $dfi(v) = ind$

2 $ind = ind + 1$

3 using A , let B be all the vertices that have arrows coming from v

4 for each w in B , MDFS(w)

GRAPHS

1 $temp = T$

$t = 1$

$d0 =$ number of words in T

2 for $t = 1, \dots, d0$

form the adjacency matrix of the overlap graph for the words in $temp$

compute the words of length $d0 + t$ and call them new

print out the adjacency matrix

$(C,I,R) = LIST(temp)$

$C = LOOP(C,R)$

$T =$ the words in R reordered according to C and the isolated words adjoined

at the end

$temp =$ words of length $d0 + t$

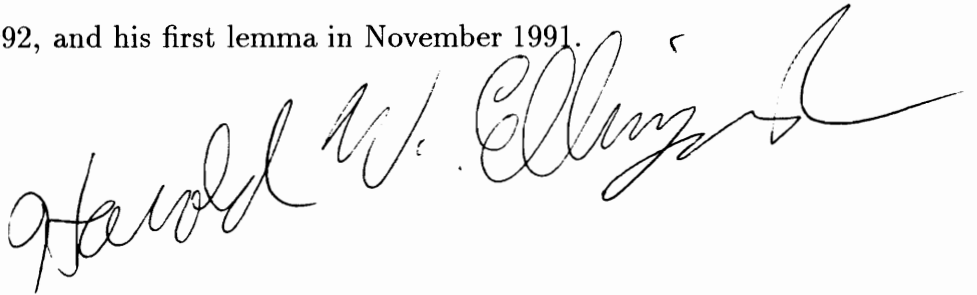
$t = t + 1$

Bibliography

- [1] D. Farkas and E.L. Green. Do subspaces have distinguished bases? *Rocky Mountain Journal*, to appear.
- [2] Alan Gibbons. *Algorithmic Graph Theory*. Cambridge University Press, Cambridge, England, 1985.
- [3] G.R. Krause and T.H. Lenagan. *Growth of algebras and Gelfand-Kirillov dimension*, volume 116 of *Research Notes in Mathematics*. Pitman Publishing Inc., London, 1985.
- [4] V.A. Ufnarovski. A growth criterion for graphs and algebras defined by words. *Math. Notes*, 31(3):238–241, March 1982.

VITA

Harold W. Ellingsen Jr. was born on May 28, 1965 in Cape May Court House, N.J. He graduated from Wildwood High School in 1983. He then went to Virginia Polytechnic Institute and State University where he received his B.S. in mathematics in 1987, his M.S. in mathematics in 1989, and his Ph.D. in mathematics in 1993. He caught his first drumfish in May 1980, his first tuna in July 1985, his first striper in November 1992, and his first lemma in November 1991.

A handwritten signature in cursive script, reading "Harold W. Ellingsen Jr.", written in black ink.