

INDEPENDENT SETS FROM AN ALGEBRAIC PERSPECTIVE

ALICIA DICKENSTEIN* and ENRIQUE A. TOBIS^{†,‡}

*Departamento de Matemática
Facultad de Ciencias Exactas y Naturales
Universidad de Buenos Aires, Buenos Aires, Argentina*
*alidick@dm.uba.ar
[†]etobis@dc.uba.ar

Received 10 June 2010
Accepted 15 April 2011
Published 23 February 2012

Communicated by S. Margolis

In this paper, we study the basic problem of counting independent sets in a graph and, in particular, the problem of counting antichains in a finite poset, from an algebraic perspective. We show that neither independence polynomials of bipartite Cohen–Macaulay graphs nor Hilbert series of initial ideals of radical zero-dimensional complete intersections ideals, can be evaluated in polynomial time, unless $\#P = P$. Moreover, we present a family of radical zero-dimensional complete intersection ideals J_P associated to a finite poset P , for which we describe a universal Gröbner basis. This implies that the bottleneck in computing the dimension of the quotient by J_P (that is, the number of zeros of J_P) using Gröbner methods lies in the description of the standard monomials.

Keywords: Independent sets; antichains; posets; Gröbner basis; Hilbert series; complexity.

Mathematics Subject Classification 2010: 05C69, 06A07, 06A11, 13P99, 68Q17

1. Introduction

We approach the basic problem of counting independent sets in a graph and, in particular, the problem of counting antichains in a finite partially ordered set, from an algebraic perspective. We derive structural considerations and complexity results.

The use of algebraic methods in the study of discrete problems, in particular problems in graph theory, was pioneered by Stanley [24] and Lovász [20], from the combinatorics side, and Herzog and Hibi [17], Simis and Vasconcelos [23] and

[‡]Present address: Children’s Hospital Boston, 300 Longwood Ave., Boston, MA, USA.

Villarreal [26] from the commutative algebra side. The enumeration of independent sets has been approached using Reverse Search (see [10]), the Belief Propagation heuristic (see [5]) and Binary Decision Diagrams (see [27]), to name a few techniques.

The main algebraic object we will use is the Hilbert Series of the initial monomial ideal associated with a graph. The problem of computing a Hilbert Series is NP-Complete (see [1]). There is a standard algorithm (first proposed in [21]) for computing the Hilbert Series of a quotient $\mathbb{C}[\mathbf{x}]/I$, where I is a homogeneous ideal in $\mathbb{C}[\mathbf{x}]$. There are some classes of ideals for which this algorithm finishes in time polynomial in the input, e.g. Borel (see [1]) and Borel-type ideals (see [14]). Open computer algebra systems (CoCoA [6], SINGULAR [13], Macaulay2 [12]) implement the standard algorithm in subtly different ways. We suggest [18, Chap. 5] as a general reference on Hilbert Series.

The connection between independent sets and commutative algebra is spearheaded by the following construction.

Definition 1.1. Let $G = (V, E)$ be a graph, with $V = \{v_1, \dots, v_n\}$. The *edge ideal* ([23, 26]) $I'_G \subseteq \mathbb{C}[x_1, \dots, x_n]$ of G is defined as

$$I'_G = \langle x_i x_j, \text{ for all } (v_i, v_j) \in E \rangle. \tag{1}$$

This ideal links independent sets in G and certain monomials. If x^α is a monomial *not* in I'_G (termed a *standard monomial*), then it encodes an independent set S of G in this way:

$$v_i \in S \Leftrightarrow x_i \mid \mathbf{x}^\alpha. \tag{2}$$

This encoding is not one-to-one. For example, the monomials x_1 and x_1^2 represent the same independent set: $\{v_1\}$. We introduce a slightly modified version of I'_G , with which we obtain a bijective encoding.

Definition 1.2. Let $G = (V, E)$ be a graph. We define the *modified edge ideal* I_G of G as

$$I_G = I'_G + \langle x_i^2, \text{ for all } v_i \rangle. \tag{3}$$

Notice that I_G is zero-dimensional (the origin is the only root), and that its standard monomials are square-free, with the degree of a monomial equal to the size of the corresponding independent set. The number of independent (or stable) sets in G thus coincides with the k -vector space dimension of the quotient of the polynomial ring in n variables over any field k by the ideal I_G . This dimension is computed in [6, 12, 13] using the additivity of the Hilbert function in short exact sequences.

In Sec. 2, we recall the definition of the Hilbert function (see (4)) and we analyze the instantiation of the standard algorithm for computing the Hilbert Series for the ideals I_G . Our main result in this section shows that the recursive calls simply correspond to counting independent sets of G that contain a pivot vertex,

and those that do not contain it. In Sec. 3, we turn our attention to the subproblem of counting the antichains of a finite poset. We present the universal reduced Gröbner Basis for a family of zero-dimensional radical ideals derived from posets. In Sec. 4, we specialize our study in the case of Cohen–Macaulay bipartite graphs, corresponding to Cohen–Macaulay ideals I'_G . Using the characterization in [17], we show that counting independent sets in such graphs is equivalent to evaluating at 2 the independence polynomial of the comparability graph of a general finite poset. Section 5 contains our complexity study. We prove that antichain polynomials cannot be evaluated in polynomial time at any non-zero rational number t unless $P = \#P$. When combined with the algebraic results from the previous sections we deduce Corollaries 5.4 and 5.5 on the intractability of the evaluation of Hilbert Series of initial ideals of zero-dimensional complete intersections and independence polynomials of Cohen–Macaulay bipartite graphs. We close with a few experimental observations in Sec. 6.

2. Counting Independent Sets via the Computation of Hilbert Series

We start by recalling a few definitions concerning Hilbert Series. Let M be a positively graded finitely generated $\mathbb{C}[\mathbf{x}]$ -module (e.g. the quotient $\mathbb{C}[\mathbf{x}]/I_G$ for some graph G). We can write

$$M = \bigoplus_{0 \leq i} M_i,$$

where M_i is the subspace of M of degree i . The Hilbert Function (HF_M) of M maps i onto $\dim_{\mathbb{C}}(M_i)$. The Hilbert Series (HS_M) of M is the generating function

$$\text{HS}_M(z) = \sum_{0 \leq i} \text{HF}_M(i)z^i. \tag{4}$$

If $M = \mathbb{C}[\mathbf{x}]/I$ for a monomial ideal I , then $\text{HF}_M(i)$ is the number of standard monomials of degree i (that is, monomials which *are not* in I).

If we take $I = I_G$ for some graph G , as we mentioned in Sec. 1, $\text{HF}_M(i)$ is then the number of independent sets of size i in G . In this case, the Hilbert Series of $\mathbb{C}[\mathbf{x}]/I_G$ is a polynomial, called the *independence polynomial* of G . As usual, we denote this polynomial by $I(G, x)$. We refer the reader to [19] for a comprehensive survey of independence polynomials.

The standard algorithm for computing HS_M hinges on the following property. If we have a homogeneous exact sequence of finitely generated graded $\mathbb{C}[\mathbf{x}]$ -modules

$$0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0, \tag{5}$$

then

$$\text{HS}_M(z) = \text{HS}_{M'}(z) + \text{HS}_{M''}(z). \tag{6}$$

Given a finitely generated graded $\mathbb{C}[\mathbf{x}]$ -module M and $f \neq 0$ a homogeneous polynomial of degree d , we have the following *multiplication sequence*

$$0 \rightarrow [M/(0 :_M (f))(-d)] \xrightarrow{\varphi} M \rightarrow M/fM \rightarrow 0, \quad (7)$$

where φ is induced by multiplication by f . Here, $(0 :_M (f)) = \{g \in M, \text{ such that } gf = 0\}$, and $(-d)$ induces a degree shift, so that φ is a homogeneous map of degree 0. Rewriting equation (6) we obtain

$$\text{HS}_M(z) = \text{HS}_{M/fM}(z) + z^d \text{HS}_{(0 :_M (f))}. \quad (8)$$

The polynomial f above is called a *pivot*.

Actually, the standard algorithm does not directly compute the Hilbert Series. We can see in [18, Theorem 5.2.20] that in the case of the modified edge ideal I_G , the Hilbert Series of $M = \mathbb{C}[x_1, \dots, x_n]/I_G$ has the form

$$\text{HS}_M = \frac{\text{HN}_M(z)}{(1-z)^n}, \quad (9)$$

where $\text{HN}_M(z)$ is called the *Hilbert Numerator*. The algorithm computes $\text{HN}_M(z)$, and the series is then obtained by dividing it by $(1-z)^n$.

We reproduce the algorithm for computing the Hilbert Numerator of a monomial ideal (see [18, Theorem 5.3.7]).

Algorithm 2.1. Algorithm to compute the Hilbert Numerator of a monomial ideal I (called HN).

Require: A set of minimal monomial generators for the ideal I .

Ensure: The Hilbert Numerator of $\mathbb{C}[\mathbf{x}]/I$.

- 1: **if** the minimal generators of I are pairwise coprime **then**
 - 2: **return** $\prod_{i=1}^s (1 - z^{d_i})$, where d_i is the degree of the i th generator of I .
 - 3: **else**
 - 4: Choose a monomial p as pivot.
 - 5: $f_1 \leftarrow \text{HN}(I : p)$.
 - 6: $f_2 \leftarrow \text{HN}(I + p)$.
 - 7: **return** $z^{\deg(p)} f_1(z) + f_2(z)$.
 - 8: **end if**
-

Notice that the sets of generators of I'_G and I_G described in (1) and (3) are minimal. The process of obtaining minimal sets of generators for the ensuing recursive calls can be optimized by performing careful interreductions.

The choice of pivot must satisfy one condition. Namely,

$$\sum \deg(I : p) < \sum \deg(I) \quad \text{and} \quad \sum \deg(I + p) < \sum \deg(I). \quad (10)$$

Here, $\sum \deg(I)$ denotes the sum of the degrees of all the minimal monomial generators of I . Intuitively, this condition says that the recursive calls are made on “smaller” ideals, and shows that the algorithm terminates.

The program CoCoA implements this algorithm, and uses a certain strategy for the choice of pivot in step 4. First, it chooses any variable x_i appearing in the most number of generators of I_G . Then it picks two random generators containing that variable. The pivot is the highest power of x_i that divides *both* random generators.

We present a specialized version of Algorithm 2.1, suited for the computation of the Hilbert Series of $\mathbb{C}[\mathbf{x}]/I_G$ for any graph G .

Theorem 2.2. *Let I_G be the modified edge ideal of a graph G . The general algorithm for computing the Hilbert Series of $\mathbb{C}[\mathbf{x}]/I_G$ has the specialized version presented in Algorithm 2.3.*

This algorithm has an obvious graphical interpretation. The choice of step 4 corresponds to choosing a node v of the graph. The recursive calls of step 7 correspond to counting the independent sets of G that contain v (HS_{Colon}) and those that do not contain v (HS_{Plus}).

Algorithm 2.3. Specialized algorithm to compute the HS of $\mathbb{C}[\mathbf{x}]/I_G$.

Require: The list L of minimal monomial generators of I_G described in (3).

Ensure: The Hilbert Series of $\mathbb{C}[\mathbf{x}]/I_G$.

- 1: **if** L consists only of variables and squares of variables **then**
 - 2: **return** $(1 + z)^k$, where k is the number of variables which appear squared in L .
 - 3: **else**
 - 4: Choose a variable x_i that appears squared in L .
 - 5: Colon \leftarrow a minimal set of monomial generators of $\langle L \rangle : x_i$.
 - 6: Plus \leftarrow a minimal set of monomial generators of $\langle L, x_i \rangle$.
 - 7: **return** $z\text{HS}_{\text{Colon}}(z) + \text{HS}_{\text{Plus}}(z)$
 - 8: **end if**
-

Proof. Algorithm 2.3 differs from Algorithm 2.1 in two key steps. In step 1, the special version does not check coprimality, as is done in Algorithm 2.1. The other difference is in step 4: The specialized version chooses a variable, instead of an arbitrary monomial.

We make a claim that helps us understand why this specialized version is correct. In every call to the algorithm, each of the n variables appears in L raised to the first or to the second power. Furthermore, in each call, L contains only the powers just mentioned and the “edge monomials” $x_i x_j$ of G such that both x_i and x_j appear squared in L . This leads to an obvious graphical interpretation: The list L represents the subgraph of G induced by those variables that appear squared in L .

We prove the correctness of the algorithm by showing that the choice of a pivot in Algorithm 2.1 must always yield a variable when applied to a modified edge ideal, and that the claim of the previous paragraph is true.

When the algorithm is originally invoked, every variable appears squared in L . Besides the squares of variables, L contains the “edge monomials” $x_i x_j$ for every edge (i, j) of G . This proves that the claim above holds in the first call.

Assuming that the elements of L have the structure we claim, let us show that any choice of pivot yields a variable. Suppose that we employ any conceivable strategy for the choice of pivot, always subject to condition (10). The pivot p cannot be a multiple of any monomial in L . If it is, then $\text{Plus} = L$, and the decreasing total degree condition (10) is not satisfied. The pivot p must then be a product of variables that appear squared in L , but it must not be divisible by any “edge monomial.” Suppose that the pivot is the product of at least two variables. That is, $x_i x_j | p$, where x_i^2 and x_j^2 are in L , and $x_i x_j$ is not in L . Then Plus violates the decreasing total degree condition (10), because it has the same generators as L , plus p . If $p = 1$, then $\text{Colon} = L$, and this violates the decreasing total degree condition. The only valid choice is then $p = x_i$, for some x_i that appears squared in L .

Once we know that the pivot is always a variable, we can show that the claim above holds for Plus and for Colon . In doing so, we also explain the second part of the theorem.

The list of minimal monomial generators for Plus contains all the variables that were raised to the first power in L . Furthermore, it must also contain the pivot x_i . The square of x_i is not in Plus , because Plus is minimal, and the “edge monomials” that contained x_i , are not present in Plus . The rest of the generators in L are unaffected. Therefore, we have that every variable appears in Plus either squared or raised to the first power, as we wanted to show. Plus corresponds to the graph obtained by removing the node that corresponds to x_i and all the edges incident with it.

The analysis of Colon is somewhat similar. To obtain a minimal set of monomial generators, we just cross out the pivot x_i from every generator in L that contains it, and then eliminate multiples. If we had an “edge monomial” $x_i x_j$, then x_j is in Colon . Therefore, the square of x_j is no longer a generator, and all the “edge monomials” containing x_j are also missing from Colon . Again, every variable appears either squared or raised to the first power. In this case, we remove the node corresponding to x_i , all its adjacent nodes and all the edges incident with x_i or with any node adjacent to x_i .

Let v be the node of G associated with the pivot x_i . The combination step of the algorithm reflects the meaning of Colon and Plus : The independent sets of G are those of Plus (i.e. those *do not* that contain v) and those of Colon (i.e. those that contain v).

The algorithm terminates when there are no more “edge monomials”. Since all the generators are variables, or squares of variables, then they are pairwise coprime and satisfy the stopping criterion of Algorithm 2.1.

A note is in order about the value returned in the base case. Algorithm 2.1 returns

$$\prod_{i=1}^n (1 - z^{d_i}), \tag{11}$$

where d_i is the degree of the i th generator. Since in the specialized case the generators are of the form x_i or x_i^2 , expression (11) has the form

$$(1 - z)^n (1 + z)^k, \tag{12}$$

where k is the number of variables that appear squared in L . According to formula (9), the value returned by Algorithm 2.3 is the Hilbert Series of $\mathbb{C}[x_1, \dots, x_n]/I_G$.

All these observations show that the graphical interpretation is accurate and that the specialized version is indeed correct. \square

3. Partially Ordered Sets and Gröbner Bases

In this section, we study a family of zero-dimensional radical complete intersection polynomial ideals associated with posets, first proposed in [4].

Recall that a *poset* (or *partially ordered set*) is a set P , together with a (partial order) relation \leq satisfying

- $a \leq a$, for all $a \in P$;
- $a \leq b$ and $b \leq a$ implies $a = b$, for all a and b in P ;
- $a \leq b$ and $b \leq c$ implies $a \leq c$ for all a, b and c in P .

Two elements a and b of P are *comparable* if $a \leq b$ or if $b \leq a$. Otherwise, they are *incomparable*. We will usually just write P and drop the partial order relation from the notation.

We can associate to a poset P its comparability graph.

Definition 3.1. Let P be a poset. The comparability graph $G(P)$ has the set P as nodes and there is an edge between two different nodes a and b if and only if a, b are comparable in P .

A subset S of a poset P is an *antichain* if all the elements of S are pairwise incomparable in P . We write $\mathcal{A}(P)$ for the set of antichains of P . Note that $S \in \mathcal{A}(P)$ if and only if S is an independent set of $G(P)$.

Definition 3.2. For any poset P we define the *antichain polynomial* $A(P, x)$ by

$$A(P, x) = I(G(P), x).$$

Thus, the k th coefficient of $A(P, x)$ equals the number of antichains of P with k elements and the cardinal $|\mathcal{A}(P)|$ is given by the evaluation $A(P, 1)$.

Given a finite poset (P, \leq) , we define a polynomial ideal $J_P \subset \mathbb{C}[x_1, \dots, x_n]$ by:

$$J_P = \left\langle x_i - x_i \prod_{v_j \leq v_i} x_j, \text{ for all } v_i \in P \right\rangle. \quad (13)$$

Lemma 3.3. *Let P be a finite poset. Then the elements of $V(J_P)$ are strings of 0's and 1's.*

Proof. Let $a \in V(J_P)$. Suppose that an element $v_i \in P$ is minimal. Then $x_i - x_i^2 \in J_P$, hence a_i is 0 or 1. Now, take any v_i , and assume that for every $v_j < v_i$ we know that a_j is 0 or 1. Note that $x_i - x_i \prod_{v_j \leq v_i} x_j = x_i(1 - x_i \prod_{v_j < v_i} x_j)$. If any a_j is 0, then a_i must be 0 too. If all a_j are 1 then $a_i(1 - a_i) = 0$. \square

Moreover, we have the following theorem.

Theorem 3.4 ([4]). *For any finite poset P , J_P is a radical zero-dimensional ideal. Then, it has a finite number of simple zeros. Furthermore,*

$$|V(J_P)| = |\mathcal{A}(P)|. \quad (14)$$

We now show that we can present J_P as a zero-dimensional complete intersection by means of generators of lower degree. A standard alternate way of dealing with a poset P is to look at the *cover* relation. Given a and b in P , we say that $a \prec b$ (read “ b covers a ”) if and only if $a < b$ and there is no $c \in P$ such that $a < c < b$. Using this relation we define the ideal

$$J'_P = \left\langle x_i - x_i \prod_{v_j \preceq v_i} x_j, \text{ for all } v_i \in P \right\rangle. \quad (15)$$

Lemma 3.5. *Let P be a finite poset. Then*

$$J_P = J'_P. \quad (16)$$

Proof. It is straightforward to see that the varieties of J_P and J'_P coincide. We show that J'_P is radical. Since we already know that J_P is radical, this proves the equality.

It is enough to prove that the square-free polynomial $x_i - x_i^2$ is in J'_P for all $v_i \in P$. We know this to be true for the minimal elements of P , by the very definition of J'_P . Suppose we have a non-minimal element v_i in P . Let v_{j_1}, \dots, v_{j_r} be the elements such that $v_{j_k} \prec v_i$, and assume that $x_{j_l} - x_{j_l}^2$ is in J'_P for all l . First, we observe that $x_i x_{j_l} - x_i$ is in J'_P for all l . Indeed,

$$(x_{j_l} - 1) \left(x_i - x_i^2 \prod_{k=1}^r x_{j_k} \right) - \left(x_i^2 \prod_{\substack{k=1 \\ k \neq l}}^r x_{j_k} \right) (x_{j_l} - x_{j_l}^2) = x_i x_{j_l} - x_i.$$

Now, consider the following step:

$$\left(x_i - x_i^2 \prod_{k=1}^r x_{j_k} \right) - (x_i - x_i x_{j_r}) x_i \prod_{k=1}^{r-1} x_{j_k} = x_i - x_i^2 \prod_{k=1}^{r-1} x_{j_k}.$$

Since $(x_i - x_i^2 \prod_{k=1}^r x_{j_k})$ and $(x_i - x_i x_{j_r})$ are in J'_P , we have that $x_i - x_i^2 \prod_{k=1}^{r-1} x_{j_k}$ is also in J'_P . If we apply this procedure repeatedly, we eliminate variables from the product, and eventually find that $x_i - x_i^2$ is in J'_P . \square

We now take Theorem 3.4 one step further, and give an explicit bijection between $\mathcal{A}(P)$ and $V(J_P)$.

Proposition 3.6. *Let P be a finite poset. Define the function $f : V(J_P) \rightarrow \mathcal{A}(P)$ by*

$$f(a) = \{v_i \in P, \text{ such that } a_i = 1 \text{ and } a_j = 0 \text{ for all } v_j > v_i\}.$$

The map f is bijective, and its inverse $g : \mathcal{A}(P) \rightarrow V(J_P)$ is defined by

$$g(S) = a', \text{ where } a'_i = 1 \text{ if } \exists v_j \in S \text{ such that } v_i \leq v_j \text{ and } a'_i = 0 \text{ otherwise.}$$

Proof. It is clear from the definition of f that no pair of elements of the subset $f(a)$ can be comparable for any $a \in V(J_P)$, that is, that $f(a)$ is indeed an antichain. Reciprocally, let S be an antichain of P and let $a' = g(S)$. We need to see that $a'_i(1 - a'_i \prod_{v_k \leq v_i} a'_k) = 0$ for all i . This is clear if $a'_i = 0$. When $a'_i = 1$, there exists $v_j \in S$ with $v_j \geq v_i$. By the transitivity of the order relation we deduce that $a'_k = 1$ for all $v_k \leq v_i$ and so the equation is satisfied.

Let a be an element of $V(J_P)$. Let $S = f(a)$ and $a' = g(S)$. We want to show that $a = a'$. Suppose that $a'_i = 1$. Then $\exists v_j \in S$ such that $v_i \leq v_j$, and therefore $a_i = 1$. By a similar argument, if $a'_i = 0$, then $a_i = 0$. \square

We now describe the universal reduced Gröbner basis of $J_P = J'_P$.

Proposition 3.7. *The universal, reduced Gröbner Basis of J_P is the set Gb_P of polynomials*

$$\begin{aligned} gb_i &= x_i^2 - x_i \quad \forall v_i \in P, \\ gb_{(j,i)} &= x_i x_j - x_i \quad \forall v_j \leq v_i. \end{aligned}$$

Proof. Lemma 3.3 shows that the elements of $V(J_P)$ are strings of 0's and 1's. The polynomials $x_i^2 - x_i$ are in Gb_P , and therefore the elements of $V(Gb_P)$ are also strings of 0's and 1's. Let $\mathbf{x} = (x_i)_{v_i \in P}$ be a string of 0's and 1's. $\mathbf{x} \in V(J_P)$ if and only if $\forall v_i \in P, (x_i = 0 \Leftrightarrow (\exists v_j \leq v_i \text{ such that } x_j = 0))$. But this is equivalent to $\mathbf{x} \in V(Gb_P)$. Then, Gb_P is zero-dimensional, and contains a square-free univariate polynomial in each variable (gb_i) . Therefore it is also radical. This shows that the ideal generated by Gb_P coincides with J_P .

We now prove that Gb_P is a Gröbner Basis for any monomial order $<$. Recall that, given $<$ and a non-zero polynomial p , $LT_{<}(p)$ denotes the largest term of p ,

with respect to $<$. Clearly, $\text{LT}_{<}(gb_i) = x_i^2$ and $\text{LT}_{<}(gb_{(j,i)}) = x_i x_j$. Given any two polynomials in the set Gb_P , we show that their S -polynomial is divisible by the polynomials in (Gb_P) . If we let $p = x_i x_j - x_i$ and $q = x_k x_\ell - x_k$ be two polynomials in Gb_P , all possible combinations of the indices i, j, k and ℓ boil down to the following non-trivial possibilities for (p, q) (with i, j, k, ℓ all different):

- (1) $(x_i^2 - x_i, x_k x_i - x_k)$ or $(x_i x_j - x_i, x_k x_j - x_k) \Rightarrow S(p, q) = 0$.
- (2) $(x_i^2 - x_i, x_i x_\ell - x_i) \Rightarrow S(p, q) = gb_i - gb_{(\ell,i)}$.
- (3) $(x_i x_j - x_i, x_i x_\ell - x_i) \Rightarrow S(p, q) = gb_{(j,i)} - gb_{(\ell,i)}$.
- (4) $(x_i x_j - x_i, x_k x_i - x_k) \Rightarrow S(p, q) = gb_{(j,k)} - gb_{(i,k)}$.
- (5) $(x_i x_j - x_i, x_j x_\ell - x_j) \Rightarrow S(p, q) = gb_{(\ell,i)} - gb_{(j,i)}$.
- (6) $(x_i^2 - x_i, x_k^2 - x_k)$ or $(x_i^2 - x_i, x_k x_\ell - x_k)$ or $(x_i x_j - x_i, x_k x_\ell - x_k)$. In all three cases, since the leading monomials of p and q are coprime, $S(p, q)$ is divisible by (p, q) .
- (7) $(x_i x_j - x_i, x_j x_i - x_j)$. This can only hold if $v_i \leq v_j$ and $v_j \leq v_i$, that is, $v_i = v_j$.

In cases (4) and (5) above, we know that $gb_{(j,k)}$ and $gb_{(\ell,i)}$, respectively, are in Gb_P , because a partial order relation is *transitive*. Therefore, Gb_P is a Gröbner Basis.

Finally, none of the polynomials are redundant, all of the leading coefficients are one, and the “other” monomial in each polynomial of Gb_P has degree 1, so it cannot be divisible by any leading monomial of Gb_P . Therefore, Gb_P is a reduced universal Gröbner Basis of J_P . \square

We can count the antichains of P by studying J_P . We have seen that $|\mathcal{A}(P)| = |V(J_P)|$. It is well-known (see [8, Theorem 2.2.10]) that as J_P is radical, it holds that

$$|V(J_P)| = \dim_{\mathbb{C}}(\mathbb{C}[\mathbf{x}]/J_P).$$

The Hilbert Series algorithm could help us to compute $\dim_{\mathbb{C}}(\mathbb{C}[\mathbf{x}]/J_P)$, but it requires that the ideal J_P be homogeneous, which is not the case. This is circumvented by considering an initial ideal of J_P . If $<$ is a monomial order and I is an ideal, the initial ideal of $\text{LT}_{<}(I)$ is defined by

$$\text{LT}_{<}(I) = \langle \text{LT}_{<}(p), p \in I \rangle.$$

By [7, Chap. 5, Sec. 3]

$$\dim_{\mathbb{C}}(\mathbb{C}[\mathbf{x}]/I) = \dim_{\mathbb{C}}(\mathbb{C}[\mathbf{x}]/\text{LT}_{<}(I)).$$

In particular, we have the following equality

$$|\mathcal{A}(P)| = \dim_{\mathbb{C}}(\mathbb{C}[\mathbf{x}]/\text{LT}_{<}(J_P)). \tag{17}$$

Let Gb_P be the (universal) Gröbner basis of J_P in the statement of Proposition 3.7. By the definition of a Gröbner Basis, it holds that for any monomial order, $\text{LT}_{<}(J_P) = \langle \text{LT}_{<}(g), g \in Gb_P \rangle$. Note that this initial ideal has the same structure of the ideals I_G in Sec. 2. In fact, it equals $I_{G(P)}$.




Fig. 1. The path of length three P_3 .

4. Independent Sets in Bipartite Cohen–Macaulay Graphs

Let G be a graph, and I'_G its edge ideal. We say that G is a Cohen–Macaulay graph if $\mathbb{C}[\mathbf{x}]/I'_G$ is a Cohen–Macaulay $\mathbb{C}[x]$ -module. The quotient $\mathbb{C}[\mathbf{x}]/I_G$ is always Cohen–Macaulay, because I_G is zero-dimensional. Cohen–Macaulay rings and modules are extensively studied in [16], and the article [25] covers Cohen–Macaulay graphs.

Not every graph is Cohen–Macaulay, of course. For example, the path of length three (see Fig. 1) has the edge ideal $J_{P_3} = \langle x_1x_2, x_2x_3 \rangle$, defined in $\mathbb{C}[x_1, x_2, x_3]$. The quotient $\mathbb{C}[x_1, x_2, x_3]/J_{P_3}$ is not Cohen–Macaulay. It is not even equidimensional, since the zero set of J_{P_3} consists of the plane $x_2 = 0$, together with the line $x_1 = x_3 = 0$.

One particularly interesting subfamily of Cohen–Macaulay graphs are *bipartite Cohen–Macaulay graphs*.

Definition 4.1. Let $G = (V_1 \sqcup V_2, E)$ be a bipartite graph. Then G is a Cohen–Macaulay graph if and only if $\mathbb{C}[\mathbf{x}]/I'_G$ is a Cohen–Macaulay $\mathbb{C}[x]$ -module.

There is an equivalent characterization, given by the following result.

Theorem 4.2 ([17]). *Let $G = (V_1 \sqcup V_2, E)$ be a bipartite graph. We say that G is a Cohen–Macaulay graph if $|V_1| = |V_2|$, and the vertices $V_1 = \{x_1, \dots, x_n\}$ and $V_2 = \{y_1, \dots, y_n\}$ can be labeled in such a way that*

- (1) $(x_i, y_i) \in E$ for all $i = 1, \dots, n$;
- (2) if $(x_i, y_j) \in E$, then $i \leq j$;
- (3) if (x_i, y_j) and (x_j, y_k) are edges, then (x_i, y_k) is also an edge.

There are two ways of seeing a bipartite Cohen–Macaulay graph $G = (V_1 \sqcup V_2, E)$ as a poset. The obvious way is to set the following partial order on the nodes of G : $x \leq y$ if and only if $x = y$ or $x \in V_1, y \in V_2$ and (x, y) is an edge of G . That is, one chooses one of the parts as the “upper” one.

The other way, which we will consider here, involves a different construction. Let $G = (V_1 \sqcup V_2, E)$ be a bipartite Cohen–Macaulay graph. We define a poset P_G as follows. The elements of P_G are those of V_1 . Given x_i and x_j , we set $x_i \leq x_j$ if and only if the edge (x_i, y_j) is in E . From the transitivity of bipartite Cohen–Macaulay graphs, we see that P_G is a poset.

Conversely, let P be a poset, with elements x_1, \dots, x_r . We build a bipartite graph $G_P = (V, E)$ as follows. We set $V = V_1 \sqcup V_2$, with $V_1 = \{x_1, \dots, x_r\}$ and $V_2 = \{y_1, \dots, y_r\}$. We put the edges (x_i, y_i) in E for all i , and we have the edge (x_i, y_j) if and only if $x_i \leq x_j$ in P . In this case, the transitivity of \leq ensures that G_P is a bipartite Cohen–Macaulay graph.

The following lemma is straightforward.

Lemma 4.3. *The two transformations*

$$P \mapsto G_P \quad \text{and} \quad G \mapsto P_G$$

are inverses.

We now compare the independence polynomial of a bipartite Cohen–Macaulay graph G with the antichain polynomial of the poset P_G .

Lemma 4.4. *Let $I(G, x)$ be the independence polynomial of a bipartite Cohen–Macaulay graph G and let $A(P_G, x)$ be the antichain polynomial of its associated poset P_G . Then*

$$I(G, x) = A(P_G, 2x).$$

Proof. The construction outlined above expands every element of the poset P_G into a segment in the bipartite Cohen–Macaulay graph G . An antichain S of size k in P_G gives rise to 2^k independent sets of size k in the bipartite graph G , since we can replace any $x_i \in S$ by either the node x_i or the node y_i of G . It is clear that any independent set of G can be seen in this way for a unique antichain S of P_G . \square

5. Complexity Results

Is it classically known that it is not possible to count the number of antichains of a general poset (that is, to evaluate its antichain polynomial at 1) in polynomial time unless $P = \#P$ [22]. We extend this result in Theorem 5.3 to the evaluation at any non-zero rational number t , by a translation and specialization of [3, Theorem 2.2] to the context of finite posets. We then use our previous results to deduce in Corollaries 5.4 and 5.5 the hardness of evaluating the Hilbert function of initial ideals of zero-dimensional radical ideals and the independence polynomial of Cohen–Macaulay bipartite graphs.

We start with some definitions.

Definition 5.1. We define the lexicographic product poset $P_1[P_2]$ of two finite posets P_1 and P_2 as the set $P_1 \times P_2$, ordered by the relation $(x, i) \leq (y, j)$ if $x \leq y \wedge x = y \Rightarrow i \leq j$. Similarly, we define the lexicographic product graph $G_1[G_2]$ of two graphs as the set $G_1 \times G_2$ with (i, j) adjacent to (k, l) if and only if i is adjacent to k or if $i = k$ and j is adjacent to l .

It is easy to check that $P_1[P_2]$ is indeed a poset.

Given a natural number m , denote by \tilde{K}_m the poset given by the set $\{1, \dots, m\}$, ordered with the usual \leq relation. The associated comparability graph is the complete graph K_m in m nodes, whose independence polynomial equals $I(K_m, x) = 1 + mx$.

It is straightforward to check that the comparability graph of the lexicographic product $P_1[P_2]$ of two posets equals the lexicographic product $G(P_1)[G(P_2)]$ of the

respective comparability graphs. We therefore have:

Lemma 5.2. *For any poset P and $m \in \mathbb{N}$, the comparability graph of the lexicographic product $P[\tilde{K}_m]$ equals the lexicographic product of the graphs $G(P)[K_m]$.*

We are now ready to prove the following theorem.

Theorem 5.3. *Evaluating the antichain polynomial of any finite poset P at any non-zero rational number t is $\#P$ -hard.*

Proof. We mimic the arguments in [3, Theorem 2.2]. Suppose, on the contrary, that given any poset P on n vertices, there exists an $O(n^k)$ -algorithm to compute $A(P, t)$ for some constant k . Then, given a poset P with n vertices, consider the posets $P[\tilde{K}_m]$ for $m = 1, \dots, n + 1$. It follows from Lemma 5.2 that we can use the reasoning in [3, Theorem 2.2] to deduce that $A(P, mt) = A(P[\tilde{K}_m], t)$. In fact, by [2, Theorem 1]), $A(P[\tilde{K}_m], t) = A(P, A(\tilde{K}_m, t) - 1) = A(P, mt)$. As the posets $P[\tilde{K}_m]$, $m = 1, \dots, n + 1$ can be constructed in polynomial time from the data of P , it would be possible to compute $A(P, mt)$ in polynomial time for $m = 1, \dots, n + 1$. But then, the coefficients i_j of $A(P, x) = \sum_{j=0}^n i_j x^j$ would be computed in polynomial time by solving the $(n + 1) \times (n + 1)$ linear system with invertible matrix $M = (M_{ij})$ given by $M_{ij} = (jt)^{i-1}$, $i, j = 1, \dots, n + 1$.

It follows that the number of antichains $|\mathcal{A}(P)|$ of P would be computable in polynomial time by adding $\sum_j i_j$. But this counting problem is $\#P$ -complete [22]. □

Combining this complexity results with the algebraic results of the previous sections, we have the following two corollaries.

Corollary 5.4. *No algorithm can evaluate the Hilbert Series at a fixed non-zero rational number t in polynomial time when applied to initial ideals of radical zero-dimensional complete intersections, unless $\#P = P$.*

Proof. By the results of Secs. 2 and 3, the Hilbert Series of the initial ideal $I_{G(P)}$ of the radical zero-dimensional ideal J_P associated to any poset, equals the antichain polynomial $A(P, x)$. The result follows from Theorem 5.3. □

Corollary 5.5. *There can be no polynomial algorithm to evaluate at any non-zero rational number t the independence polynomial of bipartite Cohen–Macaulay graphs unless $\#P = P$.*

Proof. By Lemma 4.4, the independence polynomial $I(G, p)$ of a bipartite Cohen–Macaulay graph G and the antichain polynomial $A(P_G, x)$ of its associated poset P_G are related by the equality $I(G, x) = A(P_G, 2x)$. So, any polynomial algorithm to evaluate $I(G, t/2)$ in polynomial time for any Cohen–Macaulay graph G , would allow us to evaluate $A(P, t)$ in polynomial time for any poset by Lemma 4.3. The result now follows from Theorem 5.3. □

6. Some Experimental Observations

We tested the three Computer Algebra Systems CoCoA, SINGULAR and Macaulay 2. The examples we used were the posets consisting of the power set of $\{1, \dots, n\}$, ordered by inclusion (Boolean lattice). Of the three systems, only CoCoA managed to count the antichains for $n = 7$. These numbers (called Dedekind numbers) are known for n up to 13. However, those computations required many hours of super-computer time [15].

The strategy employed by CoCoA for the Hilbert Numerator algorithm seems to be generally good. We made some observations about it in [9]. We have also tested a recent software package, EdgeIdeals [11]. EdgeIdeals allows us to compute the Hilbert Series of a modified edge ideal of a graph G by computing the f -vector of the simplicial complex associated with the edge ideal of G . The simplicial complex also contains a description of the standard monomials of the modified edge ideal of G . The computation of both objects (the f -vector and the standard monomials) was faster using EdgeIdeals for the Boolean lattice, compared to the native Macaulay 2 implementation of Hilbert Series, for n up to 6.

Acknowledgments

Both authors were partially supported by UBACYT X064, CONICET PIP 112-200801-00483 and ANPCyT, Argentina.

References

- [1] D. Bayer and M. Stillman, Computation of Hilbert functions, *J. Symb. Comput.* **14**(1) (1992) 31–50.
- [2] J. I. Brown, C. A. Hickman and R. J. Nowakowski, On the location of roots of independence polynomials, *J. Algebr. Combin.* **19** (2004) 273–282.
- [3] J. Brown and R. Hoshino, Independence polynomials of circulants with an application to music, *Discrete Math.* **309**(8) (2009) 2292–2304.
- [4] E. Cattani and A. Dickenstein, Counting solutions to binomial complete intersections, *J. Complexity* **23**(1) (2007) 82–107.
- [5] V. Chandreskaran, M. Chertkov, D. Gamarnik, D. Shash and J. Shin, Counting independent sets using the Bethe approximation (2009), http://www-math.mit.edu/~jinwoos/submit_bp.pdf.
- [6] CoCoA Team, CoCoA: A system for doing computations in commutative algebra, <http://cocoa.dima.unige.it>.
- [7] D. Cox, J. Little and D. O’Shea, *Ideals, Varieties, and Algorithms*, Undergraduate Texts in Mathematics (Springer, 1997).
- [8] D. Cox, J. Little and D. O’Shea, *Using Algebraic Geometry*, 2nd edn., Graduate Texts in Mathematics, Vol. 185 (Springer, New York, 1998).
- [9] A. Dickenstein and E. A. Tobis, Algebraic methods for counting antichains, in *Advances in Graph Theory and Applications*, eds. J. Scharcasnki and V. Trevisan (UFRGS, Porto Alegre, 2007).
- [10] D. Eppstein, All maximal independent sets and dynamic dominance for sparse graphs, in *Proc. Sixteenth Annual ACM-SIAM Symp. Discrete Algorithms* (SIAM, 2005), pp. 451–459.

- [11] C. A. Francisco, A. Hoefel and A. V. Tuyl, EdgeIdeals: A package for (hyper)graphs, *J. Softw. Algebra Geom.* **1** (2009).
- [12] D. R. Grayson and M. E. Stillman, Macaulay 2, a software system for research in algebraic geometry, <http://www.math.uiuc.edu/Macaulay2/>.
- [13] G.-M. Greuel, G. Pfister and H. Schönemann, Singular 3.1.0 — A computer algebra system for polynomial computations (2009), <http://www.singular.uni-kl.de>.
- [14] A. Hashemi, Polynomial complexity for Hilbert series of Borel type ideals, *Albanian J. Math.* **1**(3) (2007) 145–155.
- [15] J. Heitzig and J. Reinhold, The number of unlabeled orders on fourteen elements, *Order* **17** (1999) 333–341.
- [16] J. Herzog and W. Bruns, *Cohen–Macaulay Rings*, Cambridge Studies in Advanced Mathematics, Vol. 39 (Cambridge University Press, 1993).
- [17] J. Herzog and T. Hibi, Distributive lattices, bipartite graphs and Alexander duality, *J. Algebr. Combin.* **22**(3) (2005) 289–302.
- [18] M. Kreuzer and L. Robbiano, *Computational Commutative Algebra 2* (Springer-Verlag, Heidelberg, 2005).
- [19] V. E. Levit and E. Mandrescu, The independence polynomial of a graph — A survey, in *Proc. 1st Int. Conf. Algebraic Informatics* (Aristotle University, Thessaloniki, 2005), pp. 233–254.
- [20] L. Lovász, Stable sets and polynomials, *Discrete Math.* **124** (1994) 137–153.
- [21] F. Mora and H. M. Möller, The computation of the Hilbert function, in *EUROCAL 83*, Lecture Notes in Computer Science, Vol. 162 (Springer-Verlag, 1983), pp. 157–167.
- [22] J. S. Provan and M. O. Ball, The complexity of counting cuts and of computing the probability that a graph is connected, *SIAM J. Comput.* **12**(4) (1983) 777–788.
- [23] A. Simis, W. Vasconcelos and R. H. Villarreal Rodríguez, On the ideal theory of graphs, *J. Algebra* **167**(2) (1994) 389–416.
- [24] R. P. Stanley, *Combinatorics and Commutative Algebra*, 2nd edn., Progress in Mathematics, Vol. 41 (Birkhäuser Boston, Boston, MA, 1996).
- [25] R. H. Villarreal Rodríguez, Cohen–Macaulay graphs, *Manuscripta Math.* **66**(1) (1990) 1432–1785.
- [26] R. H. Villarreal Rodríguez, *Monomial Algebras*, Pure and Applied Mathematics (CRC, 2001).
- [27] A. B. Yedidia, Counting independent sets and kernels of regular graphs, preprint (2009), <http://arxiv.org/pdf/0910.4664v1>.