COHEN-MACAULAY MONOMIAL IDEALS ARISING FROM DIRECTED HYPERGRAPHS

CHASE MEADORS

ABSTRACT. In [HH05], Herzog and Hibi study the Hibi ideal of a poset and describe its resolution in detail. Furthermore, they use this notion to classify Cohen-Macaulay bipartite graphs. In [CF07], Carra'Ferro and Farrarello demonstrate a construction associating a bipartite undirected graph G to a directed graph D, and re-cast the aforementioned classification in terms of the directed graph D. In particular, they show that G is Cohen-Macaulay if and only if D is acyclic and transitive. We study how and to what extent this result generalizes to to directed *hypergraphs*. Though we do not achieve a characterization as in the case of directed graphs, we demonstrate a class of directed hypergraphs and a sufficient condition for the associated undirected hypergraph to be Cohen-Macaulay.

1. INTRODUCTION

In [CF07], Carra'Ferro and Farrarello introduce the following construction for associating a bipartite undirected graph to a directed graph. Since we will be dealing with both kinds of graphs, we establish now the notational convention of using the notation $(x \to y)$ for directed edges in a directed graph, and $\{x, y\}$ for undirected edges in an undirected graph.

Definition 1.1. Let D be a simple directed graph on the vertex set X with edge set E_D. Let $Y = \{\overline{x} : x \in X\}$ be a set consisting of formal copies of the variables in X. Then, the undirected graph $\mathcal{G}(D)$ is a graph on the vertex set $X \sqcup Y$ with edge set E_G, where:

$$E_G = \{\{x, \overline{x}\} : x \in X\} \cup \{\{x, \overline{y}\} : (x \to y) \in E_D\}$$

Note that $\mathcal{G}(D)$ is always a bipartite graph (with bipartition $X \sqcup Y$) with a perfect matching.

In [HH05], Herzog and Hibi classify all bipartite Cohen-Macaulay graphs. That is, a necessary and sufficient set of conditions is given for a bipartite graph to be Cohen-Macaulay (that is, having a Cohen-Macaulay edge ideal). In [CF07], these results are re-interpreted – in particular, we can describe when $\mathcal{G}(D)$ is Cohen-Macaulay in terms of properties of the graph D:

Theorem 1.2. *Let D be a simple directed graph. Then D is transitive if and only if $\mathcal{G}(D)$ is Cohen-Macaulay.*

Remark 1.3. Some state this theorem with the requirement that D also be acyclic. However, we require that D be a simple graph, and it is a fact from graph theory that any simple transitive directed graph is necessarily acyclic.

Example 1.4. Let $X = \{x_1, x_2, x_3\}$ and D be a directed graph on X with edge-set $\{x_1 \to x_2, x_2 \to x_3, x_1 \to x_3\}$. For this (and other) examples, we will set $\overline{x_i} = y_i$ to improve readability. Below are D and $\mathcal{G}(D)$, respectively:

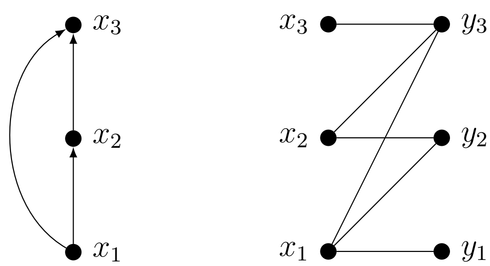

We have, in $k[x_1, x_2, x_3, y_1, y_2, y_3]$,

$$I(\mathcal{G}(D)) = (x_1y_1, x_2y_2, x_3y_3, x_1y_2, x_2y_3, x_1y_3)$$

Here D is transitive, so $I(\mathcal{G}(D))$ and thus $\mathcal{G}(D)$ is Cohen-Macaulay.

2. Recovery of Original Result

If we assume Herzog and Hibi's results from [HH05], the proof of Theorem 1.2 is straightforward. Of course, most of the technical content of the proof is contained there. As such, we give an alternative presentation of the (partial) proof from first principles. By doing this, we extract the parts of the proof that will generalize to the directed hypergraph case.

Before starting, we will need some notation. Throughout, D will be a directed graph on the vertex set $X = \{x_1, \ldots, x_n\}$, and $\mathcal{G}(D)$ will be the associated undirected graph using the same notation as 1.1. Let $S = k[x_1, \ldots, x_n, \overline{x_1}, \ldots, \overline{x_n}]$ be a polynomial ring in $2n$ variables. We will often be considering the edge ideal of $\mathcal{G}(D)$, i.e., $I(\mathcal{G}(D))$. For ease of notation we set $I_D = I(\mathcal{G}(D))$.

We also use $I_D^\vee$ to refer to the *Alexander dual* of I_D. Recall that the Alexander dual of I is the Stanley-Reisner ideal of $\Delta^\vee$, where Δ is the Stanley-Reisner simplicial complex of I.

We will need a bit more notation. Let $A \subseteq X$ be an arbitrary subset of X. We will speak about several monomials in S constructed from such a subset. In particular, we let:

$$p(A) = \prod_{x \in A} x \qquad \overline{p}(A) = \prod_{x \in A} \overline{x} \qquad u(A) = \prod_{x \in A} x \prod_{x \in X \setminus A} \overline{x}$$

Our goal in this section is to prove the forward direction of Theorem 1.2. That is,

Theorem 2.1. *For a simple digraph D, if D is transitive, then $\mathcal{G}(D)$ is Cohen-Macaulay.*

The main strategy of the proof will be to determine $I_D^\vee$ explicitly (determine its minimal generators), and show that I_D has a linear resolution. We start with a lemma:

Lemma 2.2. *Consider $J = (x\overline{x} \ : \ x \in X)$, an ideal in S. Then the Alexander dual of J is $J^\vee = (u(A) \ : \ A \in \mathscr{P}(X))$.*

Proof. Note that we can express the Alexander dual of J in a straightforward way – that is $J^\vee = \bigcap_{x \in X} (x, \overline{x})$. We now show this expression is equal to the one given in the lemma.

Suppose $m \in \bigcap_{x \in X} (x, \overline{x})$. Then, either x or $\overline{x}$ must divide m for each $x \in X$. Let $A = \{x \in X : x \mid m\}$. Then, for $x \notin A$, we must necessarily have $\overline{x} \mid m$. Thus, $u(A) \mid m$, and we have $m \in (u(A) \ : \ A \in \mathscr{P}(X))$.

Suppose $m \in (u(A) \ : \ A \in \mathscr{P}(X))$. Then it follows immediately that, for each $x \in X$, at least one of x or $\overline{x}$ must divide m, and so $m \in \bigcap_{x \in X} (x, \overline{x})$. □

This lemma is important because J is precisely I_D for the directed graph D with no edges. In fact, I_D for arbitrary D with edge set E_D will have the form:

$$I_D = J + (x\overline{y} \ : \ (x \to y) \in E_D)$$

This helps us understand the Alexander dual of I_D. In fact, the Alexander dual will have the form:

$$I_D^\vee = J^\vee \cap \bigcap_{(x \to y) \in E_D} (x, \overline{y})$$

Now, we show another lemma, elucidating the structure of $I_D^\vee$:

Lemma 2.3. *Consider a simple transitive directed graph D with edge-set E_D. If m is a minimal generator of $I_D^\vee$, then $m = u(A)$ for some subset $A \subseteq X$.*

Proof. From the previous lemma and following discussion, recall that

$$I_D^\vee = J^\vee \cap \bigcap_{(x \to y) \in E_D} (x, \overline{y})$$

Suppose m is a minimal generator of $I_D^\vee$. To show that $m = u(A)$, it suffices to show that, for any $x \in X$, at least one of x or $\overline{x}$ divide m, but not both.

Suppose towards contradiction that *neither* x or $\overline{x}$ divide m. But, then m would not be in $J^\vee$, because $u(A)$ for any subset A contains either x or $\overline{x}$, and no $u(A)$ would divide m.

Suppose that *both* x and $\overline{x}$ divide m for some x. We show that we must be able to safely remove one of the two variables from m and obtain a monomial still in the ideal, contradicting the minimality of m.

Suppose that there is an edge $(a \to x) \in E_D$ such that $a \nmid m$ (this means that we cannot possibly remove $\overline{x}$ from m, as the resulting monomial would no longer be in the ideal associated to this edge). We show that we must be able to safely remove x: Suppose towards contradiction that there was some edge $(x \to b) \in E_D$ with $\overline{b} \nmid m$. But, this cannot be the case since $(a \to b)$ is necessarily an edge, and m would fail to be in the ideal $(a, \overline{b})$, a contradiction. Thus, there are no such edges, and we find that $\overline{b} \mid m$ for every ideal $(x, \overline{b})$ in the intersection. Thus, setting $m' = m/x$, we have that m' is still in $I_D^\vee$ – certainly it is still in $J^\vee$, and the only possible additional ideals that could be affected are ideals of the form $(x, \overline{b})$, but we just established that $\overline{b} \mid m$ and thus m' in these cases, so m' is still in each of these ideals.

We use an entirely symmetric argument to show that in the opposite case (that is, there is an edge $(x \to b)$ where $\overline{b} \nmid m$), the monomial $m' = m/\overline{x}$ is still in the ideal.

Of course, if neither of the suppositions in the preceding two cases are true, it would not matter which variable we remove. So, we conclude that in any case, we must be able to remove either x or $\overline{x}$ from the ideal and obtain a monomial still in the ideal, contradicting the minimality of m. ☐

Now, we know the form of the minimal generators of $I_D^\vee$. The next step is how to characterize precisely *which* monomials $u(A)$ make up the generators. For the current case of directed *graphs*, there is a very nice interpretation, which is the key to the proof. Before the next theorem, recall that, for a directed acyclic graph D, the reachability relation in D forms a partial order – recall that a vertex y is *reachable* from a vertex x if there exists a (directed) path from x to y. In particular, if D is transitive, the reachability relation *is* the edge set. So, in this case, we can speak of the partial order on X defined by D, where $x < y$ if and only if $(x \to y) \in E_D$. Finally, recall that an *order ideal* with respect to the partial order $<$ (also called a downset) is a subset $A \subseteq X$ such that if $b \in A$ and $a < b$, then $a \in A$.

Theorem 2.4. *Consider a simple transitive directed graph D with edge set E_D. Then the generators of $I_D^\vee$ are precisely those monomials $u(A)$ where A is an order ideal with respect to the partial order on X defined by D.*

Proof. Let $m = u(A)$ be a generator of $I_D^\vee$ (from the lemma, all generators are of this form). Suppose that A is *not* an order ideal. Then, there is some $a \in A$ and $r \in X$ such that $r < A$ but $r \notin A$. But, since $r < a$, $(r \to a)$ is an edge, and m must be in the ideal $(r, \overline{a})$. But, since $r \notin A$, we have $\overline{r} \mid m$ and necessarily $r \nmid m$. Since $a \in A$, we have $a \mid m$ and necessarily $\overline{a} \nmid m$. Then, m cannot be in the ideal $(r, \overline{a})$, a contradiction.

Let A be an order ideal and set $m = u(A)$. We show it's in the ideal (and thus a minimal generator, due to its form). Certainly, $m \in J^\vee$. Take any of the additional ideals in the intersection $(x, \overline{y})$. This means $(x \to y)$ is an edge and $x < y$. In the case that $y \in A$, so is x, and so necessarily $x \mid m$. In the case that $y \notin A$, we have $\overline{y} \mid m$. Either way, $m \in (x, \overline{y})$. Then, $m \in I_D^\vee$. □

It is at this point that, using the language of [HH05], we realize that $I_D^\vee$ is precisely the *Hibi ideal* of the poset defined by D – i.e., the ideal generated by monomials $u(A)$ where A ranges over the order ideals of the poset defined by D. It is shown in [HH05] that Hibi ideals have linear resolutions (in fact, the resolution is computed explicitly). Thus, we finish the proof of Theorem 2.1 by noting that since $I_D^\vee$ has a linear resolution, I_D is Cohen-Macaulay (by Eagon-Reiner).

The converse, through true, uses more graph theory than algebra, and since we do not aim to generalize it, we will not present it here. Interested readers should refer to [CF07] for such a proof.

3. Generalization

We now study how and to what extent Theorem 2.1 generalizes to directed *hypergraphs*. First, we present the relevant definitions to remove ambiguity:

Definition 3.1. A *hypergraph* is a pair $G = (V, E)$ of a vertex set V and an edge-set E, where elements of E are subsets of V. A *simple* hypergraph satisfies the condition that $x \not\subseteq y$ for any $x, y \in E$.

There is some question as to what a *directed* hypergraph should be. One generalization, and indeed the one that works well for our purposes, is the following:

Definition 3.2. A *directed hypergraph* is a pair $D = (V, E)$ of a vertex set V and an edge-set E, where elements of E are tuples (S, T), where S and T are disjoint subsets of V. The elements of S are the *sources* the edge, and the elements of T are the *targets* of the edge. A *simple* hypergraph satisfies the condition that for any $(S, T), (S', T') \in E$, $S \cup T \not\subseteq S' \cup T'$.

Remark 3.3. This generalization of a directed graph makes sense in the following way: A edge in a directed graph adds structure to the plain undirected edge by specifying a source and a target for the edge (giving it an orientation). Here, an edge in a directed hypergraph adds the same structure by *partitioning* the plain undirected edge into sources and targets.

We will also need the notion of a cycle in a directed hypergraph:

Definition 3.4. A *cycle* in a directed hypergraph D is a sequence of vertices $\{v_1, \ldots, v_n\}$ such that $v_1 = v_n$ and for each $1 \leq i < n$, there exists an edge $(S \to T)$ such that $v_i \in S$ and $v_{i+1} \in T$.

This notion of directed hypergraph is well-known an actively studied, particularly in the Computer Science literature (see [GLPN93], [AFF01] for broad surveys of the topic).

The construction to produce an undirected hypergraph from a directed one carries over nicely:

Definition 3.5. Let D be a simple directed hypergraph on the vertex set X with edge-set E_D. Let $Y = \{\overline{x} : x \in X\}$ be a set consisting of formal copies of the variables in X. Then, the undirected graph $\mathcal{G}(D)$ is a graph on the vertex set $X \sqcup Y$ with edge set E_G, where:

$$E_G = \{\{x, \overline{x}\} : x \in X\} \cup \{S \cup \{\overline{x} : x \in T\} : (S \to T) \in E_D\}$$

The next question that arises is what it should mean for a directed hypergraph to be *transitive*. There are ostensibly many different ways to generalize transitivity to directed hypergraphs. The following example provides motivation for the belief that *some* notion of transitivity will serve to generalize Theorem 2.1:

Example 3.6. Consider the directed graph D on vertex set $\{x_1, \ldots, x_7\}$ with edge set $\{\{x_1, x_2\} \to \{x_3, x_4\}, \{x_4, x_5\} \to \{x_6, x_7\}\}$:

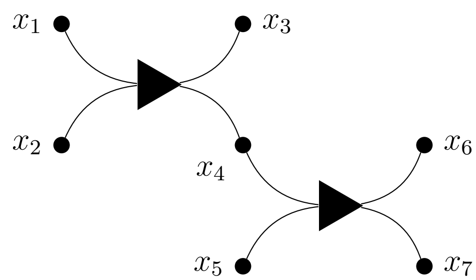

Again, set $\overline{x_i} = y_i$. The edge ideal of the associated hypergraph is then

$$I_D = I(\mathcal{G}(D)) = (x_1y_1, x_2y_2, x_3y_3, x_4y_4, x_5y_5, x_6y_6, x_7y_7, x_1x_2y_3y_4, x_4x_5y_6y_7)$$

which is *not* Cohen-Macaulay. However, consider the following two closely related graphs. The left, D_1, is obtained by adding the edge $\{x_1\} \to \{x_6\}$ to D. The right, D_2, is obtained by adding two new vertices x_8, x_9 and the edge $\{x_1, x_8\} \to \{x_6, x_9\}$ to D.

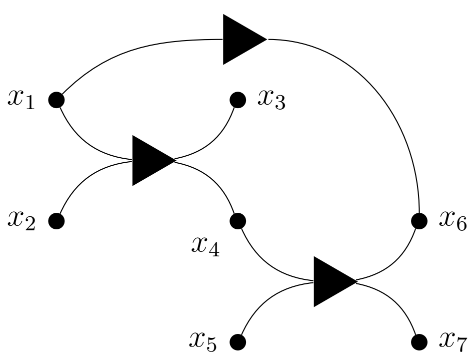

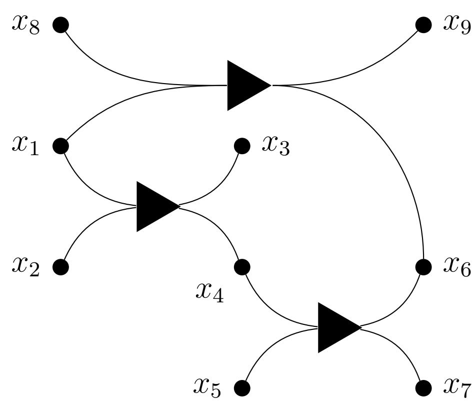

One can verify that I_{D_1} *is* Cohen-Macaulay, while I_{D_2} is again *not* Cohen-Macaulay.

Motivated by this example, we make the following definition:

Definition 3.7. A directed hypergraph D is *intersection containment transitive* or *IC-transitive* if, whenever (S_1, T_1) and (S_2, T_2) are edges with $T_1 \cap S_2 \neq \varnothing$, then there exists an edge (S_1', T_2'), where $S_1' \subseteq S_1$ and $T_2' \subseteq T_2$

Note that the graph D_1 in Example 3.6 is the only one satisfying IC-transitivity, and also the only one to produce a Cohen-Macaulay undirected hypergraph.

For digraphs, we made use of the fact that a simple transitive digraph was necessarily acyclic. In the present setting however, an analogous statement is not true. The following is an example of a simple IC-transitive directed hypergraph with a cycle:

Example 3.8. Let D be the directed hypergraph on $\{x_1, \dots, x_9\}$ with edge set:

$$\{x_1, x_4\} \to \{x_2, x_5\}$$
$$\{x_2, x_6\} \to \{x_3, x_7\}$$
$$\{x_3, x_8\} \to \{x_1, x_9\}$$
$$\{x_4\} \to \{x_7\}$$
$$\{x_6\} \to \{x_9\}$$
$$\{x_8\} \to \{x_5\}$$

One can verify that this graph is indeed simple and satisfies IC-transitivity, but has a cycle $\{x_1, x_2, x_3, x_1\}$.

We will indeed require acyclicity in the proofs of the next section, but due to the preceding discussion we must explicitly require it instead of simply specifying IC-transitivity. From now on, we let $\mathfrak{D}$ refer to the class of simple, acyclic, IC-transitive directed hypergraphs. In the next section, we mimic the proofs of section 2 in this generalized setting, and see that they mostly carry over.

4. Proof of Generalized Theorem

Our goal in this section is to prove the generalization of Theorem 2.1:

Conjecture 4.1. *If $D \in \mathfrak{D}$, then I_D is Cohen-Macaulay.*

With Theorem 4.10 at the end of this section, we manage to prove the conjecture for two broad subclasses of graph. First, we see that the Alexander dual of the associated edge ideal has the same kind of generators:

Lemma 4.2. *Let $D \in \mathfrak{D}$ with edge-set E. If m is a minimal generator of $I_D^\vee$, then $m = u(A)$ for some $A \subseteq X$.*

Proof. The proof proceeds in the same way as Lemma 2.3. We have, using notation introduced in the last section:

$$I_D^\vee = J^\vee \cap \bigcap_{(S \to T) \in E} \mathfrak{p}(p(S)\overline{p}(T))$$

where $\mathfrak{p}(m)$ denotes the prime monomial ideal generated by the support of m.

So, suppose m is a minimal generator of $I_D^\vee$. Again, we show that for every $x \in X$, either $x \mid m$ or $\overline{x} \mid m$ but not both. Similarly to before, if neither x or $\overline{x}$ divides m, then m would fail to be in $J^\vee$, so this cannot be the case.

Suppose *both* x and $\overline{x}$ divide m for some x. Again, we show that we can remove one of the two and obtain a monomial still in the ideal, contradicting the minimality of m.

Suppose there is an edge $(A \to U)$ with $x \in U$ such that for all $y \in A$, $y \nmid m$ and for all $z \in U \setminus \{x\}$, $\overline{z} \nmid m$. This means that we cannot possibly remove $\overline{x}$ from m, as the resulting monomial would no longer would no longer be in the ideal associated to this edge. Then, we show that we must be able to safely remove x: Suppose towards contradiction that there were some edge $(U' \to B)$ with $x \in U'$ such that for all $y \in U' \setminus \{x\}$, $y \nmid m$ and for all $y \in B$, $\overline{y} \nmid m$. But, since $U \cap U'$ is non-null (it contains at least x), we must have an edge $(A' \to B')$ with $A' \subseteq A$ and $B' \subseteq B$, but m would fail to be in the ideal $\mathfrak{p}(p(A')\overline{p}(B'))$, a contradiction. Then, setting $m' = m/x$, we find that m' is still in the ideal. Certainly, it is still in $J^\vee$, and and the only possible additional ideals in the intersection that could be affected are ones of the form $\mathfrak{p}(p(U')\overline{p}(B))$, but we established that in all of these cases, some variable in the prime ideal besides x must divide m and thus m', so m' is still in these ideals.

Again, an entirely symmetric argument shows the opposite case (that is, if we cannot remove x then we must be able to remove $\overline{x}$).

So, in any case, we must be able to remove either x or $\overline{x}$ from m and obtain a monomial still in the ideal, contradicting the minimality of m. □

So the generators of $I_D^\vee$ look the same as in the directed graph case. The concern again becomes determining precisely *which* monomials $u(A)$ make up the generators. The characterization here is not quite as nice, but is still relatively easy to describe:

Definition 4.3. For a subset $A \subseteq X$, and an edge $(S \to T)$ in D such that $T \subseteq A$, A is said to satisfy the *target containment property with repsect to* T if, for all edges $(S' \to T)$, $A \cap S' \neq \varnothing$. If A satisfies the target containment property with respect to every target contained in it, A is simply said to satisfy the *target containment property.*

Theorem 4.4. *Let $D \in \mathfrak{D}$. For $A \subseteq X$, $u(A)$ is a minimal generator of $I_D^\vee$ if and only if A satisfies the target containment property.*

Remark 4.5. Notice that subsets A not containing any whole target of an edge vacuously satisfy the target containment property – as such, all such subsets are represented in the generators.

Proof. Let $m = u(A)$ be a generator of $I_D^\vee$ (from the lemma, all generators are of this form). Suppose that A does *not* satisfy the property. Then, there exists an edge $(S \to T)$ and an edge $(S' \to T)$ such that $T \subseteq A$ but $A \cap S' = \varnothing$. But, in particular, this means $S' \nsubseteq A$. Thus, for all $y \in S'$, $\overline{y} \mid m$ and thus $y \nmid m$. Since $T \subseteq A$, for all $z \in T$, we have $z \mid m$ and so $\overline{z} \nmid m$. But then, m fails to be in the ideal $\mathfrak{p}(p(S')\overline{p}(T))$ – a contradiction.

Let A be a subset satisfying the property and set $m = u(A)$. We show it's in the ideal (and thus a minimal generator, due to its form). Certainly, $m \in J^\vee$. Consider an arbitrary additional ideal in the intersection, say $\mathfrak{p}(p(S)\overline{p}(T))$, corresponding to an edge $(S \to T)$. In the case that $T \subseteq A$, then there necessarily exists some $x \in S \cap A$, and $x \mid m$, so m is in the prime ideal. If $T \nsubseteq A$, then there exists some $x \in T \backslash A$, and $\overline{x} \mid m$, so m is again in the prime ideal. □

At this point in section 2, we relied on $I_D^\vee$ being the Hibi ideal of a poset – an object with a known linear resolution. Here, we do not have the same luxury, and we must show that $I_D^\vee$ has a linear resolution via other means.

Many examples and computations give reason to believe that $I_D^\vee$ has linear quotients. Since $I_D^\vee$ is always generated in a single degree, this implies that $I_D^\vee$ has a linear resolution. At this time, we cannot determine a linear quotients order in full generality (that is, for any $D \in \mathfrak{D}$). However, if we again add one more restriction to the graph D, we can show linear quotients for this special case.

Definition 4.6. For a directed hypergraph D, D is a *single-sourced* graph if, for every edge $(S \to T)$, we have $|S| = 1$. Likewise, we call D a *single-targeted* graph if, for every edge $(S \to T)$, we have $|T| = 1$.

In preparation for the next theorem, recall the following combinatorial condition equivalent to linear quotients for a squarefree monomial ideal ([HH11]):

Lemma 4.7. *Suppose $I = (m_1, \dots, m_n)$ is a squarefree monomial ideal, and let $M_i = \operatorname{supp}(m_i)$ for all i. I has linear quotients with respect to this generator ordering if and only if, for all i and all $j < i$, there exists an $x \in M_j \backslash M_i$ and an integer $k < i$ so that $M_k \backslash M_i = \{x\}$.*

Now, we demonstrate a linear quotients order for single-sourced graphs in $\mathfrak{D}$:

Theorem 4.8. *Let $D \in \mathfrak{D}$ be a single-sourced graph and suppose $I_D^\vee = (m_1, \dots, m_k)$. Let $\{A_1, \dots, A_k\}$ be the subsets such that $m_i = u(A_i)$. Order the generators of $I_D^\vee$ such that $|A_i| \geq |A_{i-1}|$ for all $i \geq 2$. Then $\{m_1, \dots, m_k\}$ is a linear quotients order for $I_D^\vee$.*

Proof. Since D is single-sourced, we abuse notation throughout and write $(s \to T)$ for edges, where $s \in X$ and $T \subseteq X$. For $A_1 = \varnothing$, there is nothing to check. So, let $i \geq 2$ and $j < i$ be arbitrary.

Consider $A_i \backslash A_j$. We claim that there exists an $x \in A_i \backslash A_j$ such that, for all edges $(s \to T)$ with $T \subseteq A_i$, $x \neq s$. Notice that, intuitively, this means removing x will not affect the target containment property with respect to T for any T contained in A_i.

For the moment, assume this claim. Note that $x \in A_i \backslash A_j$ implies that $\overline{x} \mid m_j$ but $\overline{x} \nmid m_i$ (i.e., $\overline{x} \in M_j \backslash M_i$). Write $B = A_i \backslash \{x\}$. B still satisfies the target-containment property by the choice of x. In particular then, $u(B)$ is among the minimal generators, and $m_k = u(B)$ for some k strictly less than i (because the size of B is smaller). Then, since B and A_i agree on all of their members except for x, we have that $\overline{x}$ is the *only* variable dividing m_k that does not divide m_i (i.e., $M_k \backslash M_i = \{\overline{x}\}$), satisfying the linear quotients property.

Now, to prove the claim, suppose it is not true. Then, for every $x \in A_i \backslash A_j$, there exists *some* edge $(x \to T_x)$ with $T_x \subseteq A_i$ (note that $A_i \backslash A_j$ is non-empty, since A_j is a distinct set of equal or smaller cardinality). Observe that we cannot have $T_x \subseteq A_j$ (if it were, x is also necessarily in A_j by the target containment property, but x was stated *not* to be in A_j). Then, for every $x \in A_i \backslash A_j$, there must exist a $t_x \in T_x$ such that $t_x \notin A_j$ – and in particular $t_x \in A_i \backslash A_j$. Now we have a well-defined mapping $A_i \backslash A_j \to A_i \backslash A_j$ by $x \mapsto t_x$. Clearly, it must be the case that x and t_x are distinct since the sources and targets of an edge are required to be disjoint. But, such a map corresponds to a cycle in D – a walk that starts at some vertex and iteratively follows edges from x to t_x will reach a vertex that has already been encountered. Since graphs in $\mathfrak{D}$ are acyclic, this is a contradiction. $\square$

Next, we demonstrate that the opposite ordering is always a linear quotients order for single-targeted graphs:

Theorem 4.9. *Let $D \in \mathfrak{D}$ be a single-targeted graph and suppose $I_D^\vee = (m_1, \dots, m_k)$ with $m_i = u(A_i)$ as before. Order the generators such that $|A_i| \leq |A_{i-1}|$ for all $i \geq 2$. Then $\{m_1, \dots, m_k\}$ is a linear quotients order for $I_D^\vee$.*

Proof. The proof proceeds in a complementary way to that of Theorem 4.8. Again, we abuse notation and write $(S \to t)$ for edges where $S \subseteq X$ and $t \in X$. For $A_1 = X$, there is nothing to check. So, let $i \geq 2$ and $j < i$ be arbitrary.

Consider $A_j \backslash A_i$. We claim there is an $x \in A_j \backslash A_i$ such that, for all edges $(S \to x)$ (if any), $A_i \cap S \neq \varnothing$. This means that *adding* x to A_i will produce a set still satisfying the target containment property.

For the moment, assume this claim. $x \in A_j \backslash A_i$ means that $x \mid m_j$ but $x \nmid m_i$ (i.e., $x \in M_j \backslash M_i$). Write $B = A_i \cup \{x\}$. B satisfies the target containment property by the choice of x as discussed earlier. As such, $u(B)$ is among the minimal generators and $m_k = u(B)$ for some k strictly less than i (because the size of B is *greater*). Then, since B and A_i agree on all their members except for x, we have that x is the only variable dividing m_k that does not divide m_i i.e., $M_k \backslash M_i = \{x\}$), satisfying the linear quotients property.

To prove the claim, suppose that it is not true. Then, for every $x \in A_j \backslash A_i$, there exists some edge $(S_x \to x)$ with $A_i \cap S_x = \varnothing$. Then, because $x \in A_j$, we must have some $s_x \in S_x$ with $s_x \in A_j$. In particular, $s_x \notin A_i$ because of the null intersection, so $s_x \in A_j \backslash A_i$. Then again we have a map $A_j \backslash A_i \to A_j \backslash A_i$ by $x \mapsto s_x$. Clearly, x and s_x must be distinct since the sources and targets of an edge are required to be disjoint. But such a map corresponds to a cycle in D – a sequence of vertices formed by iteratively applying the map $x \mapsto s_x$ will reach a vertex that has already been encountered – following said sequence backwards would correspond to following the edges of a cycle in D – again a contradiction. □

So, in these cases $I_D^\vee$ has a linear resolution. Then, we have shown the following special case of Conjecture 4.1:

Theorem 4.10. *If $D \in \mathfrak{D}$ is single-sourced or single-targeted, then I_D is Cohen-Macaulay.*

5. Conclusion

For graphs in $\mathfrak{D}$ that are neither single-sourced nor single-targeted, there is strong evidence to believe that there always exists a linear quotients order, but such orders are always non-monotone with respect to the cardinality of the subsets. Whereas the previous proofs started with a relatively simple order and showed it was sufficient, a proof for the general case would likely have to construct a unique ordering for each possible graph, depending on intricate combinatorial properties of the graph.

Readers familiar with Stanley-Reisner theory will recognize that investigating linear quotients orders for $I_D^\vee$ is equivalent to investigating shelling orders for the Stanley-Reisner complex of I_D – call it Δ_D. In fact, in this case, the facets of Δ_D are obtained from the generators of $I_D^\vee$ by simply 'toggling the bar' on the variables. In other words, if $I_D^\vee = (u(A_1), \dots, u(A_k))$, then the facets of Δ_D are $\{u(A_1^c), \dots, u(A_k^c)\}$. One conjecture that is supported by a considerable amount of evidence is the following:

Conjecture 5.1. *If $D \in \mathfrak{D}$, Δ_D is vertex-decomposable. Furthermore, any $x \in X$ such that $x \notin T$ for any edge $(S \to T)$ is a shedding vertex.*

Since Δ_D is pure (equivalently, $I_D^\vee$ is generated in a single degree), this would imply the shellability of Δ_D and complete the proof in the general case.

References

[AFF01] Giorgio Ausiello, Paolo G. Franciosa, and Daniele Frigioni. Directed hypergraphs: Problems, algorithmic results, and a novel decremental approach. In *Lecture Notes in Computer Science*, pages 312–328. Springer Berlin Heidelberg, 2001. https://doi.org/10.1007/3-540-45446-2_20.

[CF07] Giuseppa Carra'Ferro and Daniela Farrarello. Cohen-macaulay graphs arising from digraphs. https://arxiv.org/abs/math/0703417, 2007.

[GLPN93] Giorgio Gallo, Giustino Longo, Stefano Pallottino, and Sang Nguyen. Directed hypergraphs and applications. *Discrete Applied Mathematics*, 42(2-3):177–201, Apr 1993. https://doi.org/10.1016/0166-218X(93)90045-P.

[HH05] Jürgen Herzog and Takayuki Hibi. Distributive lattices, bipartite graphs and alexander duality. *Journal of Algebraic Combinatorics*, 22(3):289–302, Nov 2005. https://doi.org/10.1007/s10801-005-4528-1.

[HH11] Jürgen Herzog and Takayuki Hibi. *Monomial Ideals*. Springer London, 2011. https://doi.org/10.1007/978-0-85729-106-6.