Journal of Pure and Applied Algebra 219 (2015) 3889–3912

Contents lists available at ScienceDirect

Journal of Pure and Applied Algebra

www.elsevier.com/locate/jpaa

Path ideals of weighted graphs

Bethany Kubik [a], Sean Sather-Wagstaff [b,*,1]

[a] USMA Department of Mathematical Sciences, 601 Thayer Rd., West Point, NY 10996, USA
[b] Department of Mathematics, NDSU Dept. #2750, PO Box 6050, Fargo, ND 58108-6050, USA

ARTICLE INFO

Article history:
Received 7 August 2014
Received in revised form 17 November 2014
Available online 12 January 2015
Communicated by S. Iyengar

MSC:
Primary: 05C22; 05E40; 13F55; 13H10; secondary: 05C05; 05C69

ABSTRACT

We introduce and study the weighted r-path ideal of a weighted graph G_ω, which is a common generalization of Conca and De Negri's r-path ideal for unweighted graphs and Paulsen and Sather-Wagstaff's edge ideal of the weighted graph. Over a field, we explicitly describe primary decompositions of these ideals, and we characterize Cohen–Macaulayness of these ideals for trees (with arbitrary r) and complete graphs (for $r = 2$).

0. Introduction

Assumption. Throughout this paper, let G be a (finite, simple) graph with vertex set $V = V(G) = \{v_1, \dots, v_n\}$ of cardinality $n \geqslant 1$ and edge set $E(G) = E$. Let A be a non-zero commutative ring, and set $S = A[X_1, \dots, X_n]$ unless otherwise specified. Fix an integer $r \in \mathbb{N} = \{1, 2, \dots\}$.

Commutative algebra and combinatorics have a rich history of fruitful interactions. In this paper, we focus on the connections between commutative algebra and graph theory. For our purposes, this begins with Villarreal's notion [16,17] of the edge ideal associated with the graph G, which is the ideal $I(G)$ in S "generated by the edges of G". Much research has been done on the relations between the combinatorial properties of G and the algebraic properties of $I(G)$; see, e.g., [3–6,8–10,13–15]. For instance, it is straightforward to show that, when A is a field, an irredundant primary decomposition of the ideal $I(G)$ is determined by "vertex covers" of the graph G. Thus, given decomposition information about $I(G)$, one can deduce combinatorial information about G, and vice versa.

Recently, this construction has been generalized in two different directions relevant to our work. First, Conca and De Negri [2] introduce the r-path ideal of G, when G is a tree. This is the ideal $I_r(G)$ of

* Corresponding author.
E-mail addresses: bethany.kubik@usma.edu (B. Kubik), sean.sather-wagstaff@ndsu.edu.
URL: http://www.ndsu.edu/pubweb/~ssatherw/ (S. Sather-Wagstaff).

1 Sean Sather-Wagstaff was supported in part by a grant from the NSA.

http://dx.doi.org/10.1016/j.jpaa.2014.12.027
0022-4049/

S "generated by the paths in G of length r". This recovers Villarreal's edge ideal as the special case $I_1(G) = I(G)$. See also [1] for useful properties of this construction, including a characterization of the Cohen–Macaulay property of $I_r(G)$.

Next, Paulsen and Sather-Wagstaff [11] introduce the edge ideal of a weighted graph G_ω, i.e., a graph G equipped with a function $\omega: E \to \mathbb{N}$ that assigns to each edge e of G a weight $\omega(e)$. The edge ideal $I(G_\omega)$ in this case is generated by the weighted edges of G_ω. In particular, if $1: E \to \mathbb{N}$ is the constant function $1(e) = 1$, then $I(G_1) = I(G)$. See Section 1 for foundational material about weighted graphs.

In the current paper, we introduce and study a common generalization of these two constructions, the *weighted r-path ideal* associated to G_ω. This is the ideal $I_r(G_\omega)$ of S that is "generated by the weighted paths of length r of G":

$$I_r(G_\omega) = \left(X_{i_1}^{e_{i_1}} \cdots X_{i_{r+1}}^{e_{i_{r+1}}} \;\middle|\; \begin{array}{l} v_{i_1} \dots v_{i_{r+1}} \text{ is a path in } G \text{ with } e_{i_1} = \omega(v_{i_1} v_{i_2}), \\ e_{i_j} = \max\{\omega(v_{i_{j-1}} v_{i_j}), \omega(v_{i_j} v_{i_{j+1}})\} \text{ for } 1 < j \leqslant r \\ \text{and } e_{i_{r+1}} = \omega(v_{i_r} v_{i_r+1}) \end{array} \right) S.$$

As before, this recovers the previous constructions as special cases with $I_r(G_1) = I_r(G)$ and $I_1(G_\omega) = I(G_\omega)$.

We investigate foundational properties of $I_r(G_\omega)$ in Section 2. In particular, the following decomposition result is proved in Theorem 2.7.

Theorem A. *Given a weighted graph G_ω one has*

$$I_r(G_\omega) = \bigcap_{(W,\sigma)} P_{(W,\sigma)} = \bigcap_{(W,\sigma)\min} P_{(W,\sigma)}$$

where the first intersection is taken over all weighted r-path vertex covers of G_ω, and the second intersection is taken over all minimal weighted r-path vertex covers of G_ω. Moreover, the second intersection is irredundant.

(See Section 1 for definitions of terms like "weighted r-path vertex cover".) When A is a field, this result yields a primary decomposition of $I_r(G_\omega)$.

In Section 3 we turn our attention to Cohen–Macaulayness of $I_r(G_\omega)$ when the underlying graph G is a tree. The main result of this section is the following, which is proved in Theorem 3.11.

Theorem B. *Assume that G_ω is a weighted tree and that A is a field. Then the following conditions are equivalent:*

(i) *$I_r(G_\omega)$ is Cohen–Macaulay;*
(ii) *$I_r(G_\omega)$ is m-unmixed; and*
(iii) *there is a weighted tree Γ_μ and an r-path suspension H_λ of Γ_μ such that H_λ is obtained by pruning a sequence of r-pathless leaves from G_ω and for all $v_i v_j \in E(\Gamma_\mu)$ we have $\omega(v_i v_j) \leqslant \min\{\omega(v_i y_{i,1}), \omega(v_j y_{j,1})\}$.*

Note that this shows that Cohen–Macaulayness of path ideals of weighted trees is independent of the characteristic of A.

Section 4 is devoted to Cohen–Macaulayness of $I_r(K^n_\omega)$, where $G = K^n$ is complete, i.e., an n-clique. Note that it is straightforward to show that the edge ideal $I_1(K^n_\omega)$ is always Cohen–Macaulay, since it is unmixed of dimension 1. On the other hand, the case of $I_r(G_\omega)$ with $r \geqslant 2$ is more complicated. We deal with the case $r = 2$, the proof of which takes up most of Section 4; see Theorems 4.7 and 4.12.

Theorem C. *Assume that $n \geqslant 3$, and let K_ω^n be a weighted n-clique. Assume that A is a field. Then the ideal $I_2(K_\omega^n)$ is Cohen–Macaulay if and only if every induced weighted sub-3-clique $K_{\omega'}^3$ of K_ω^n has $I_2(K_{\omega'}^3)$ Cohen–Macaulay.*

As in Theorem B, this shows that the Cohen–Macaulay property is characteristic-independent for cliques. Unlike Theorem B, though, it does not say that Cohen–Macaulayness is equivalent to unmixedness. See Example 4.10 for a weighted 4-clique that is unmixed but not Cohen–Macaulay.

Finally, we note that in Sections 1 and 2 we deal with a more general situation than the one described in this introduction. It uses the following.

Notation. Throughout this paper, G_ω is a weighted graph. Let $\mathcal{P}_2(\mathbb{N})$ denote the set of subsets $U \subset \mathbb{N}$ such that $|U| \leqslant 2$. Fix a function $f: \mathcal{P}_2(\mathbb{N}) \to \mathbb{N}$, and write $f\{a,b\}$ in place of $f(\{a,b\})$. For instance, f may be max, min, gcd, or lcm.

1. Weighted graphs and weighted r-path vertex covers

In this section, we develop the graph theory used in the rest of the paper, beginning with the unweighted situation.

Definition 1.1. An *r-path* in G is a sequence $v_{i_1} \dots v_{i_{r+1}}$ of distinct vertices in G such that the pair $v_{i_j} v_{i_{j+1}}$ is an edge in G for $j = 1, \dots, r$.

Example 1.2. Let G be the following tree

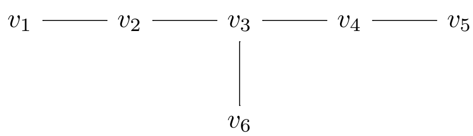

and consider the case $r = 3$. Then G has four distinct 3-paths, namely $v_1v_2v_3v_4$, $v_1v_2v_3v_6$, $v_2v_3v_4v_5$, and $v_6v_3v_4v_5$.

The next notion is key for Theorem A and the rest of the paper.

Definition 1.3. An *r-path vertex cover* of G is a subset $W \subseteq V$ such that for any path $v_{i_1} \dots v_{i_{r+1}}$ of length r in G we have $v_{i_j} \in W$ for some j. In this case, we write that v_{i_j} "covers" the path.

An r-path vertex cover of G is *minimal* if it is minimal with respect to containment, that is, it does not properly contain another r-path vertex cover of G.

For instance, consider the tree G from Example 1.2 with $r = 3$. Then the singleton $\{v_3\}$ is a 3-path vertex cover, since each 3-path in G is covered by v_3. We represent this diagrammatically, as follows.

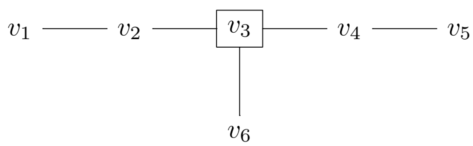

Moreover, this is a minimal 3-path vertex cover of G since $\emptyset$ is not a 3-path vertex cover. On the other hand, no other singleton is a 3-path vertex cover. (For instance, the vertex v_1 does not cover the path $v_6v_3v_4v_5$.) However, the set $\{v_1, v_5\}$ is another minimal 3-path vertex cover of G.

For graphs represented diagrammatically, we use the diagram for a visual representation of the weight function ω by decorating each edge $v_i v_j$ with the weight $\omega(v_i v_j)$, as follows.

Example 1.4. A particular weight function ω on the tree G from Example 1.2 is represented in the following diagram.

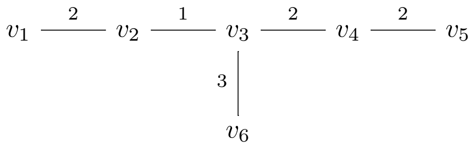

For instance, this means that $\omega(v_3 v_6) = 3$.

As one may expect, the following definition provides a combinatorial description of decompositions of ideals constructed from G_ω. See Section 2.

Definition 1.5. Set $\Lambda = \{(W, \sigma) \mid W \subseteq V \text{ and } \sigma : W \to \mathbb{N}\}$. For each $(W, \sigma) \in \Lambda$, we set $|(W, \sigma)| = |W|$.

An *f-weighted r-path vertex cover* of a weighted graph G_ω is an ordered pair $(W, \sigma) \in \Lambda$ such that for every path $v_{i_1} \dots v_{i_{r+1}}$ of length r in G, there exists an index j such that $v_{i_j} \in W$ and one of the following holds:

(1) if $j = 1$, then $\sigma(v_{i_j}) \leqslant \omega(v_{i_1} v_{i_2})$;
(2) if $j = r + 1$, then $\sigma(v_{i_j}) \leqslant \omega(v_{i_r} v_{i_{r+1}})$; or
(3) if $1 < j \leqslant r$, then $\sigma(v_{i_j}) \leqslant f\{\omega(v_{i_{j-1}} v_{i_j}), \omega(v_{i_j} v_{i_{j+1}})\}$.

(In particular, when (W, σ) is an f-weighted r-path vertex cover of G_ω, the set W is an r-path vertex cover of the unweighted graph G.) The number $\sigma(v_{i_j})$ is the *weight* of v_{i_j}. When v_{i_j} satisfies one of the above conditions, we write that it *covers* the path $v_{i_1} \dots v_{i_{r+1}}$. When $f = \max$, we write that (W, σ) is a *weighted r-path vertex cover* of G_ω.

We represent f-weighted r-path vertex covers algebraically and diagrammatically, as follows.

Example 1.6. Consider the weighted tree G_ω from Example 1.4 with $r = 3$ and with $f = \max$. The set $\{v_3\}$ is a 3-path vertex cover of G, and the function $\sigma : \{v_3\} \to \mathbb{N}$ given by $\sigma(v_3) = 2$ yields a weighted 3-path vertex cover of G_ω. We represent this algebraically and diagrammatically, by decorating the vertex v_3 with the weight $\sigma(v_3) = 2$, as follows.

$(W, \sigma) = \{v_3^2\}$ and

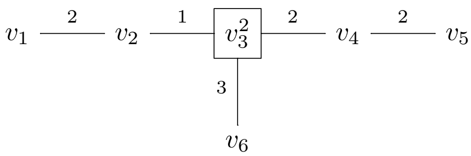

By definition, a function $\sigma' : \{v_3\} \to \mathbb{N}$ yields a weighted 3-path vertex cover of G_ω if and only if $\sigma'(v_3) \leqslant 2$. Similarly, a decorated set $\{v_1^{d_1}, v_5^{d_5}\}$ describes a weighted 3-path vertex cover of G_ω if and only if $d_1, d_5 \leqslant 2$.

Definition 1.7. Given $(W, \sigma), (W', \sigma') \in \Lambda$, we write $(W', \sigma') \leqslant (W, \sigma)$ if $W' \subseteq W$ and for all $v_i \in W'$ we have $\sigma(v_i) \leqslant \sigma'(v_i)$. Naturally, we write $(W', \sigma') < (W, \sigma)$ whenever we have $(W', \sigma') \leqslant (W, \sigma)$ and $(W', \sigma') \neq (W, \sigma)$. An f-weighted r-path vertex cover (W, σ) is *minimal* if it is minimal with respect to this ordering, that is, if there does not exist another f-weighted r-path vertex cover (W', σ') such that $(W', \sigma') < (W, \sigma)$.

Example 1.8. Consider the weighted tree G_ω from Example 1.4 with $r = 3$ and with $f = \max$. The decorated sets $\{v_3^2\}$ and $\{v_1^2, v_5^2\}$ are minimal weighted 3-path vertex covers of G_ω.

Example 1.9. Given an r-path vertex cover W of G, it is straightforward to show that the constant function $\sigma: W \to \mathbb{N}$ with $\sigma(v) = 1$ provides an f-weighted r-path vertex cover (W, σ).

The next two results are for use in the proof of Theorem A.

Lemma 1.10. *Assume that for all $j \in \mathbb{N}$ we have an f-weighted r-path vertex cover $(W_j, \sigma_j) = \{v_{i_1}^{a_1}, \dots, v_{i_p}^{a_p}, v_{i_{p+1}}^{b_j}\}$ of G_ω. If the sequence $\{b_1, b_2, \dots\}$ is unbounded, then $(W, \sigma) = \{v_{i_1}^{a_1}, \dots, v_{i_p}^{a_p}\}$ is also an f-weighted r-path vertex cover of G_ω.*

Proof. By assumption, there exists an index j such that b_j is greater than each of the following numbers: $\omega(v_p v_q)$ for each edge $v_p v_q$ in G, and $f(\omega(v_i v_j), \omega(v_j v_k))$ for each 2-path $v_i v_j v_k$ in G. It follows that the weighted vertex $v_{i_{p+1}}^{b_j}$ does not cover any f-weighted path in G_ω. Since (W_j, σ_j) is an f-weighted r-path vertex cover of G_ω, it follows that (W, σ) is an f-weighted r-path vertex cover of G_ω. □

Lemma 1.11. *For every f-weighted r-path vertex cover (W, σ) of G_ω there is a minimal f-weighted r-path vertex cover (W'', σ'') of G_ω with $(W'', \sigma'') \leqslant (W, \sigma)$.*

Proof. If (W, σ) is a minimal f-weighted r-path vertex cover then we are done. If (W, σ) is not minimal, then either there is a $v_i \in W$ that can be removed or for some $v_i \in W$ the function $\sigma(v_i)$ can be increased. In the first case, remove vertices from W until the removal of one more vertex creates a path without a vertex to cover it. This process must terminate in finitely many steps because W is finite. Let us denote our new f-weighted r-path vertex cover as (W', σ'). If no vertices are removed, then $(W, \sigma) = (W', \sigma')$.

Lemma 1.10 shows that each vertex $v_i \in W'$ has a bound beyond which one cannot increase the weight on v_i without losing the f-weighted r-path vertex covering property, assuming the weights on the other vertices are held constant. In sequence, increase the weight of each vertex to such a bound. Denote the new ordered pair (W'', σ''). Then, by construction, (W'', σ'') is a minimal f-weighted r-path vertex cover such that $(W'', \sigma'') \leqslant (W, \sigma)$, and we are done. □

The next result uses $f = \max$.

Lemma 1.12. *Every minimal weighted r-path vertex cover of G_ω has cardinality at most $n - 1$.*

Proof. In the case $n \leqslant r$, the graph G has no r-paths, so the empty set describes the unique minimal weighted r-path vertex cover of G_ω. This has cardinality $0 < n$, as desired. Thus, for the remainder of the proof, we assume that $n > r$.

Let (W, σ) be a weighted r-path vertex cover of G_ω. We show that, if $|W| = n$, then (W, σ) is not minimal.

Assume that $|W| = n$, and write $(W, \sigma) = \{v_1^{e_1}, \dots, v_n^{e_n}\}$. Reorder the v_i if necessary to assume that $e_1 \leqslant e_2 \leqslant \cdots \leqslant e_n$. We show that $v_n^{e_n}$ is superfluous in the vertex cover.

Suppose by way of contradiction that $v_n^{e_n}$ cannot be removed from (W, σ). This implies that one of the r-paths p in G can only be covered by the weighted vertex $v_n^{e_n}$. In particular, p must pass through v_n, so assume that p uses the vertices $v_{i_1}, \dots, v_{i_r}, v_n$ with $i_1, \dots, i_r < n$.

As a special case, assume that p has the following form.

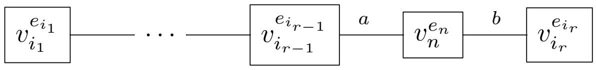

By assumption, the weighted vertices $v_{i_{r-1}}^{e_{i_{r-1}}}$ and $v_{i_r}^{e_{i_r}}$ do not cover this path, so we have $e_{i_{r-1}} > a$ and $e_{i_r} > b$. Also, the weighted vertex $v_n^{e_n}$ does cover this path, so we have $e_n \leqslant a < e_{i_{r-1}} \leqslant e_n$ or $e_n \leqslant b < e_{i_r} \leqslant e_n$, a contradiction.

The general case where v_n is not an endpoint of p is handled similarly. The remaining case where v_n is an endpoint of p is similar, but easier. □

Definition 1.13. A weighted graph G_ω is *r-path unmixed with respect to f* if all minimal f-weighted r-path vertex covers have the same cardinality; G_ω is *r-path mixed with respect to f* is if it is not r-path unmixed. We write that the unweighted graph G is "r-path (un)mixed" when the trivially weighted graph (with $\omega(e) = 1$ for all $e \in E$) is so.

2. Weighted path ideals and their decompositions

In this section, we introduce and study weighted path ideals. In particular, we prove Theorem A from the introduction here.

Definition 2.1. The *f-weighted r-path ideal* associated to G_ω is the ideal $I_{r,f}(G_\omega)$ of S that is "generated by the weighted r-paths in G_ω".

$$I_{r,f}(G_\omega) = \left(X_{i_1}^{e_{i_1}} \cdots X_{i_{r+1}}^{e_{i_{r+1}}} \;\middle|\; \begin{array}{l} v_{i_1} \ldots v_{i_{r+1}} \text{ is a path in } G \text{ with } e_{i_1} = \omega(v_{i_1} v_{i_2}), \\ e_{i_j} = f\{\omega(v_{i_{j-1}} v_{i_j}), \omega(v_{i_j} v_{i_{j+1}})\} \text{ for } 1 < j \leqslant r, \\ \text{and } e_{i_{r+1}} = \omega(v_{i_r} v_{i_r+1}) \end{array} \right) S$$

See Remark 2.4 for some justification for this definition.

Example 2.2. Consider the weighted tree G_ω from Example 1.4 with $r = 3$ and with $f = \max$. The 3-path $v_1 v_2 v_3 v_6$ provides one generator of $I_{3,\max}(G_\omega)$, namely

$$X_1^{\omega(v_1 v_2)} X_2^{\max\{\omega(v_1 v_2), \omega(v_2 v_3)\}} X_3^{\max\{\omega(v_2 v_3), \omega(v_3 v_6)\}} X_6^{\omega(v_3 v_6)} = X_1^2 X_2^2 X_3^3 X_6^3.$$

From the remaining 3-paths, we find that

$$I_{3,\max}(G_\omega) = \left(X_1^2 X_2^2 X_3^3 X_6^3, X_1^2 X_2^2 X_3^2 X_4^2, X_2 X_3^2 X_4^2 X_5^2, X_3^3 X_4^2 X_5^2 X_6^3\right) S.$$

Remark 2.3. In the case $r = 1$, the ideal $I_{1,f}(G_\omega)$ is the "weighted edge ideal" of [11]. (Note that this is independent of the choice of f.) When $\omega(e) = 1$ for all $e \in E$ and $f = \max$, we recover the "path ideal" $I_r(G)$ of [1,2]. Also, the special case $f = \max$ yields the ideal $I_r(G_\omega)$ from the introduction.

Remark 2.4. Our definition of $I_{r,f}(G_\omega)$ probably deserves some justification. Our purpose is to have this definition satisfy the conclusions of Remark 2.3. In order to recover the path ideal of [1,2], the generators should correspond to the r-paths in G. To recover the weighted edge ideal of [11] in the case $r = 1$, the generator corresponding to a path $\zeta = v_{i_1} \ldots v_{i_{r+1}}$ should be of the form $X_{i_1}^{e_{i_1}} \cdots X_{i_{r+1}}^{e_{i_{r+1}}}$ where the exponent e_{i_j} depends on the weights of the edges in ζ that are adjacent to the vertex v_{i_j}. For the endpoints v_{i_1} and $v_{i_{r+1}}$, it seems reasonable to simply use the weight of the only relevant edges, namely, $\omega(v_{i_1} v_{i_2})$ and

$\omega(v_{i_r}v_{i_{r+1}})$. However, when $1 < j \leqslant r$, the value of e_{i_j} should depend on both weights $\omega(v_{i_{j-1}}v_{i_j})$ and $\omega(v_{i_j}v_{i_{j+1}})$. We entertained several ideas about the "best" way to combine these two weights to define e_{i_j}, including max, min, gcd, and lcm.

Theorem 2.7 shows that, from the point of view of decomposing $I_{r,f}(G_\omega)$ (e.g., computing a primary decomposition of $I_{r,f}(G_\omega)$, determining unmixedness, etc. when A is a field) there is no "best" choice for f. In other words, every choice for f yields an ideal that we can explicitly decompose. (In principle, this explains our choice of condition (3) in Definition 1.5. While this condition may seem a little strange, it is the exact condition that works for our decomposition result.) On the other hand, our results on Cohen–Macaulayness in Sections 3 and 4 indicate that the choice $f = \max$ is somewhat nicer than others we considered, in that it seems more difficult to characterize Cohen–Macaulayness of $I_{r,f}(G_\omega)$ when $f \neq \max$.

In the next definition, recall the notation Λ from 1.5.

Definition 2.5. For all $(W, \sigma) \in \Lambda$ we write $P_{(W,\sigma)} = (X_i^{\sigma(v_i)} | v_i \in W)S$.

One advantage for the algebraic notation from Example 1.6 for elements of Λ, is that it explicitly provides generators for the ideal $P_{(W,\sigma)}$. For instance, with $(W, \sigma) = \{v_1^2, v_5^2\}$, we have

$$P_{(W,\sigma)} = P_{\{v_1^2, v_5^2\}} = \big(X_1^2, X_5^2\big)S.$$

Remark 2.6. It is straightforward to show that the ideals in S of the form $P_{(W,\sigma)}$ are precisely the indecomposable elements of the set of monomial ideals of S. In other words, a monomial ideal I of S is of the form $P_{(W,\sigma)}$ if and only if it satisfies the following: for all monomial ideals J_1, J_2 such that $I = J_1 \cap J_2$, one has $I = J_i$ for some $j \in \{1, 2\}$. (In the language of [12], these are the "m-irreducible" monomial ideals of S.) When the coefficient ring A is a field, the ideal $P_{(W,\sigma)}$ is primary with $\operatorname{rad}(P_{(W,\sigma)}) = (X_i \mid v_i \in W)S$. Hence, when we are working over a field, Theorem 2.7(b) below gives an irredundant primary decomposition of $I_{r,f}(G_\omega)$. In general, this is the "m-irreducible decomposition" of [12].

It is straightforward to show that every monomial ideal I of S admits a unique irredundant m-irreducible decomposition $I = P_{(W_1,\sigma_1)} \cap \cdots \cap P_{(W_t,\sigma_t)}$; uniqueness here is up to reordering of the ideals in the decomposition, and "irredundant" means that no ideal in this decomposition is contained in any other ideal in the decomposition. We write that I is *m-unmixed* provided that all the W_i in this decomposition have the same cardinality. We write that I is *m-mixed* provided that it is not m-unmixed. When we are working over a field, these are equivalent to I being unmixed or mixed, respectively.

The next result contains Theorem A from the introduction.

Theorem 2.7.

(a) *Given $(W, \sigma) \in \Lambda$, one has $I_{r,f}(G_\omega) \subseteq P_{(W,\sigma)}$ if and only if (W, σ) is an f-weighted r-path vertex cover of G_ω.*

(b) *One has decompositions*

$$I_{r,f}(G_\omega) = \bigcap_{(W,\sigma)} P_{(W,\sigma)} = \bigcap_{(W,\sigma)\min} P_{(W,\sigma)}$$

where the first intersection is taken over all f-weighted r-path vertex covers of G_ω, and the second intersection is taken over all minimal f-weighted r-path vertex covers of G_ω. Moreover, the second intersection is irredundant.

Proof. (a) First assume that (W,σ) is an f-weighted r-path vertex cover of G_ω, and let $v_{i_1}\cdots v_{i_{r+1}}$ be an r-path in G. By definition, there exists a $j\in\{1,\dots,r+1\}$ such that $v_{i_j}\in W$ and one of the following holds:

$j=1$: we have $\sigma(v_{i_1})\leqslant\omega(v_{i_1}v_{i_2})=e_{i_1}$;
$j=r+1$: we have $\sigma(v_{i_{r+1}})\leqslant\omega(v_{i_r}v_{i_{r+1}})=e_{i_{r+1}}$; or
$1<j\leqslant r$: we have $\sigma(v_{i_j})\leqslant f\{\omega(v_{i_{j-1}}v_{i_j}),\omega(v_{i_j}v_{i_{j+1}})\}=e_{i_j}$.

In each case we have $v_{i_j}\in W$ and $\sigma(v_{i_j})\leqslant e_{i_j}$. Thus, $X_{i_j}^{\sigma(v_{i_j})}$ divides $X_{i_j}^{e_{i_j}}$, and hence the generator $X_{i_1}^{e_{i_1}}\cdots X_{i_{r+1}}^{e_{i_{r+1}}}$ of $I_{r,f}(G_\omega)$ is in $P_{(W,\sigma)}$. Since this is true for each r-path in G, we conclude that $I_{r,f}(G_\omega)\subseteq P_{(W,\sigma)}$.

Conversely, assume that $I_{r,f}(G_\omega)\subseteq P_{(W,\sigma)}$ and let $v_{i_1}\cdots v_{i_{r+1}}$ be an r-path in G. By assumption we have $X_{i_1}^{e_{i_1}}\cdots X_{i_{r+1}}^{e_{i_{r+1}}}\in I_{r,f}(G_\omega)\subseteq P_{(W,\sigma)}=(X_i^{\sigma(v_i)}|v_i\in W)$. Hence there exists an i such that v_i is in W and the associated generator $X_i^{\sigma(v_i)}$ divides $X_{i_1}^{e_{i_1}}\cdots X_{i_{r+1}}^{e_{i_{r+1}}}$. Since $\sigma(v_i)\geqslant 1$, there exists a j such that $i_j=i$ and $\sigma(v_i)\leqslant e_{i_j}$. That is, there exists a j such that $v_{i_j}=v_i\in W$ and $\sigma(v_{i_j})\leqslant e_{i_j}$. Since this is true for each r-path in G, we conclude that (W,σ) is an f-weighted r-path vertex cover of G_ω.

(b) This follows from Lemma 1.11 and part (a), as in [11, Theorem 3.5]. □

Corollary 2.8. *We have* $\operatorname{depth}(S/I_{r,f}(G_\omega))\geqslant 1$.

Proof. Lemma 1.12 and Theorem 2.7 imply that the maximal ideal $(X_1,\dots,X_n)S$ is not associated to $I_{2,f}(K^n_\omega)$, hence the desired conclusion. □

Remark 2.9. Remark 2.6 and Theorem 2.7 imply that $I_{r,f}(G_\omega)$ is m-unmixed if and only if G_ω is r-path unmixed. In particular, the r-path ideal $I_r(G)$ of [1,2] is m-unmixed if and only if the unweighted graph G is r-path unmixed.

Example 2.10. Consider the weighted tree G_ω from Example 1.4 with $r=3$ and with $f=\max$. The ideal $I_{3,\max}(G_\omega)$, computed in Example 2.2, decomposes irredundantly as follows:

$$\begin{aligned}I_{3,\max}(G_\omega)=\big(X_3^2\big)S\cap\big(X_1^2,X_4^2\big)S\cap\big(X_1^2,X_5^2\big)S\cap\big(X_2^2,X_4^2\big)S\cap\big(X_2^2,X_5^2\big)S\\ \cap\big(X_3^3,X_4^2\big)S\cap\big(X_4^2,X_6^3\big)S\cap\big(X_2,X_3^3\big)S\cap\big(X_2,X_6^3\big)S.\end{aligned}$$

If one computes this algebraically (as we did), one can identify all of the minimal weighted r-path vertex covers of G_ω. (For instance, the minimal weighted r-path vertex covers $\{v_3^2\}$ and $\{v_1^2,v_5^2\}$ from Example 1.8 are visible via the ideals $(X_3^2)S$ and $(X_1^2,X_5^2)S$ in the decomposition.) On the other hand, if one is combinatorially inclined, one can first identify all of the minimal f-weighted r-path vertex covers of G_ω, and then obtain the decomposition from Theorem 2.7.

The next lemma is for use in the proof of Theorem B.

Lemma 2.11. *If $I_{r,f}(G_\omega)$ is m-unmixed, then $I_r(G)$ is also m-unmixed.*

Proof. Assume that $I_{r,f}(G_\omega)$ is m-unmixed. Then there exists an integer k such that every minimal f-weighted r-path vertex cover (W,σ) of G_ω has $|W|=k$. Let W' be a minimal r-path vertex cover of G. We show that $|W'|=k$.

As we observed in Example 1.9, the constant function $\sigma':W'\to\mathbb{N}$ given by $\sigma'(v_i)=1$ yields an f-weighted r-path vertex cover (W',σ') of G_ω. Lemma 1.11 implies that there exists a minimal f-weighted r-path vertex cover (W'',σ'') of G_ω such that $(W'',\sigma'')\leqslant(W',\sigma')$. By assumption, we have $|W''|=k$. By the minimality of W', we have $W''=W'$, hence $|W'|=|W''|=k$. □

We conclude this section with two lemmas used in the proof of Theorem C.

Lemma 2.12. *Let $G'_{\omega'}$ denote the weighted subgraph of G induced by $V \smallsetminus \{v_n\}$. Set $S' = A[X_1, \ldots, X_{n-1}]$. Then the natural isomorphism $S/(X_n)S \cong S'$ induces an isomorphism*

$$S/\big(I_{r,f}(G_\omega) + (X_n)S\big) \cong S'/I_{r,f}\big(G'_{\omega'}\big).$$

Proof. Let $\tau: S/(X_n)S \to S'/I_{r,f}(G'_{\omega'})$ denote the composition of the natural maps $S/(X_n)S \xrightarrow{\cong} S' \twoheadrightarrow S'/I_{r,f}(G'_{\omega'})$. To show that τ induces a well-defined epimorphism $\pi: S/(I_{r,f}(G_\omega) + (X_n)S) \twoheadrightarrow S'/I_{r,f}(G'_{\omega'})$, it suffices to show that each generator of $I_{r,f}(G_\omega)(S/(X_n)S)$ is in $\operatorname{Ker}(\tau)$. Note that the generators of $I_{r,f}(G_\omega)(S/(X_n)S)$ correspond to the r-paths in G that do not pass through v_n. That is, they correspond to the r-paths in G'. Since $\omega'(e) = \omega(e)$ for each edge in G', it follows that the generators of $I_{r,f}(G_\omega)(S/(X_n)S)$ and $I_{r,f}(G'_{\omega'})$ corresponding to such a path are equal. This gives the desired result about $\operatorname{Ker}(\tau)$. A similar argument shows that $\operatorname{Ker}(\tau) = I_{r,f}(G_\omega)(S/(X_n)S)$, so the induced map π is an isomorphism. □

Lemma 2.13. *The ideal $I_{r,f}(G_\omega)$ can be written as*

$$I_{r,f}(G_\omega) = \sum I_{r,f}\big(G'_{\omega'}\big)S$$

where the sum is taken over all weighted subgraphs $G'_{\omega'}$ of G_ω induced by $r+1$ vertices. (If $G'_{\omega'}$ is induced by $v_{i_1}, \ldots, v_{i_{r+1}}$ with $i_1 < \cdots < i_{r+1}$, then we consider $I_{r,f}(G'_{\omega'})$ in the polynomial subring $A[X_{i_1}, \ldots, X_{i_{r+1}}] \subseteq S$.)

Proof. For the containment $\supseteq$, note that each generator g of $I_{r,f}(G'_{\omega'})S$ is determined by an r-path in G', which is an r-path in G with the same weights; hence g is also a generator of $I_{r,f}(G_\omega)$. For the reverse containment, note that each generator h of $I_{r,f}(G_\omega)$ comes from an r-path in G_ω, and this r-path lives in a (unique) induced weighted subgraph $G'_{\omega'}$ of G_ω on $r+1$ vertices; thus, h is also a generator of $I_{r,f}(G'_{\omega'})S$. □

3. Cohen–Macaulay weighted trees

Assumption. Throughout this section, A is a field.

The point of this section is to prove Theorem B from the introduction characterizing Cohen–Macaulayness of trees in the context of weighted path ideals for the function $f = \max$.

Definition 3.1. Assume that v_i is a vertex of degree 1 in G that is not a part of any r-path in G. We write that v_i is an *r-pathless leaf* of G_ω. Let H_λ be the weighted subgraph of G_ω induced by the vertex subset $V \smallsetminus \{v_i\}$. We write that H_λ is obtained by *pruning an r-pathless leaf* from G_ω. A weighted subgraph Γ_μ of G_ω is obtained by *pruning a sequence of r-pathless leaves* from G_ω if there exists a sequence of weighted graphs $G_\omega = G^{(0)}_{\omega^{(0)}}, G^{(1)}_{\omega^{(1)}}, \ldots, G^{(l)}_{\omega^{(l)}} = \Gamma_\mu$ such that each $G^{(i+1)}_{\omega^{(i+1)}}$ is obtained by pruning an r-pathless leaf from $G^{(i)}_{\omega^{(i)}}$.

Example 3.2. In the weighted tree G_ω from Example 1.4, the vertex v_6 is a 4-pathless leaf. Pruning this leaf yields the following weighted path H_λ.

$$v_1 \overset{2}{\text{———}} v_2 \overset{1}{\text{———}} v_3 \overset{2}{\text{———}} v_4 \overset{2}{\text{———}} v_5$$

Next, we state some consequences of the existence of an r-pathless leaf in G_ω.

Lemma 3.3. *Let H_λ be a weighted graph obtained by pruning a single r-pathless leaf v_i from G_ω.*

(a) *The set of r-paths in G is the same as the set of r-paths in H.*
(b) *Assume that (W, σ) is an f-weighted r-path vertex cover of G_ω such that $v_i \in W$. Set $W' = W \smallsetminus \{v_i\}$ and $\sigma' = \sigma|_{W'}$. Then (W', σ') is an f-weighted r-path vertex cover of G_ω.*
(c) *The minimal f-weighted r-path vertex covers of G_ω are the same as the minimal f-weighted r-path vertex covers of H_λ, so G_ω is r-path unmixed with respect to f if and only if H_λ is so.*

Proof. (a) This follows by definition of H since no r-paths in G pass through v_i.

(b) Since no r-paths pass through v_i, this vertex does not cover any r-paths, so it can be removed.

(c) Combining parts (a) and (b), we conclude that the f-weighted r-path vertex covers of H_λ are exactly the f-weighted r-path vertex covers (W, σ) of G_ω such that $v_i \notin W$. The desired conclusion about minimal elements now follows. □

The next definition is key for Theorem B.

Definition 3.4. Let Γ_μ be a weighted graph. The *r-path suspension* of the unweighted graph Γ is the graph obtained by adding a new path of length r to each vertex of Γ. The new r-paths are called *r-whiskers*. A weighted graph H_λ is a *weighted r-path suspension* of Γ_μ provided that the unweighted graph H is an r-path suspension of Γ.

Example 3.5. The weighted tree G_ω from Example 1.4 is a weighted 1-path suspension of the following weighted path.

$$v_2 \;\overset{1}{\rule{2em}{0.4pt}}\; v_3 \;\overset{2}{\rule{2em}{0.4pt}}\; v_4$$

Examples of weighted r-path suspensions of G_ω itself are given by the following, where the edges of G are drawn double for emphasis.

$r = 1$

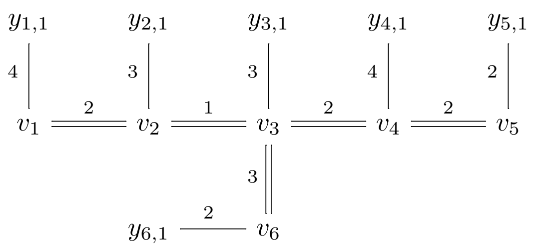

$(G'_{\omega'})$

$r = 2$

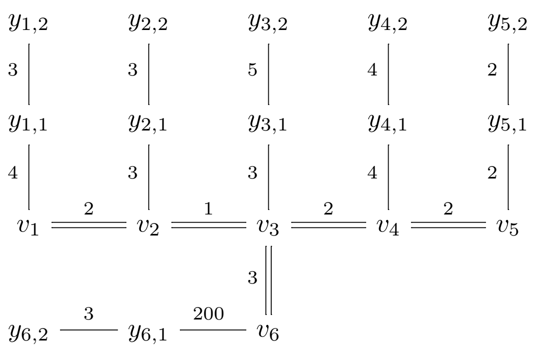

$(G''_{\omega''})$

Remark 3.6. A weighted graph H_λ is an r-path suspension of another weighted graph Γ_μ if and only if H has a sequence of pair-wise disjoint paths $p_1, p_2, \dots, p_\beta$ of length r such that (after appropriately renaming the vertices of H) the vertices of each p_i can be ordered as $v_i, y_{i,1}, \dots, y_{i,r}$ where $\deg(y_{i,k}) = 2$ for $k = 1, \dots, r-1$, and $\deg(y_{i,r}) = 1$, such that $V(H) = \{v_1, y_{1,1}, \dots, y_{1,r}, \dots, v_\beta, y_{\beta,1}, \dots, y_{\beta,r}\}$. In this case, Γ is the induced subgraph of H associated to the subset $\{v_1, \dots, v_\beta\} \subseteq V$. When this is the case, we write $S = A[X_1, Y_{1,1}, \dots, Y_{1,r}, \dots, X_\beta, Y_{\beta,1}, \dots, Y_{\beta,r}]$ instead of $A[X_1, \dots, X_n]$ for the polynomial ring containing $I_{r,\max}(H_\lambda)$.

The following proposition contains one implication of Theorem B.

Proposition 3.7. *Let H_λ be an r-path suspension of the weighted graph Γ_μ, with notation as in Remark 3.6, such that for all $v_iv_j \in E(\Gamma)$ we have $\omega(v_iv_j) \leqslant \min\{\omega(v_iy_{i,1}), \omega(v_jy_{j,1})\}$. Then $I_{r,\max}(H_\lambda)$ is Cohen–Macaulay.*

Proof. As in the proof of [11, Lemma 5.3], we polarize the ideal $I := I_{r,\max}(H_\lambda)$ to obtain a new ideal $\widetilde{I}$ in a new polynomial ring $\widetilde{S}$. We then show that $\widetilde{I}$ is the polarization of another monomial ideal J in another polynomial ring T such that T/J is artinian. In particular, T/J is Cohen–Macaulay. Since T/J and S/I are graded specializations of $\widetilde{S}/\widetilde{I}$, it follows that $\widetilde{S}/\widetilde{I}$ and S/I are also Cohen–Macaulay.

In preparation, we set some notation

$$\begin{aligned}
&a_i := \omega(v_iy_{i,1}) && \text{for } i = 1, \dots, \beta\\
&a_{i,1} := \max\{\omega(v_iy_{i,1}), \omega(y_{i,1}y_{i,2})\} && \text{for } i = 1, \dots, \beta\\
&a_{i,j} := \max\{\omega(y_{i,j-1}y_{i,j}), \omega(y_{i,j}y_{i,j+1})\} && \text{for } i = 1, \dots, \beta \text{ and } j = 2, \dots, r-1\\
&a_{i,r} := \omega(y_{i,r-1}y_{i,r}) && \text{for } i = 1, \dots, \beta\\
&t_{i,j} = \omega(y_{i,j-1}y_{i,j}) && \text{for } i = 1, \dots, \beta \text{ and } j = 2, \dots, r-1\\
&b_{p,q,r} = \max\{\omega(v_pv_q), \omega(v_qv_r)\} && \text{for all 2-paths } v_pv_qv_r \text{ in } \Gamma\\
&c_{i,j} = \omega(v_iv_j) && \text{for all edges } v_iv_j \text{ in } \Gamma
\end{aligned}$$

The polynomial ring $\widetilde{S}$ has coefficients in A with the following list of variables.

$$\begin{aligned}
&X_{1,1}, \dots, X_{1,a_1}, Y_{1,1,1}, \dots, Y_{1,1,a_{1,1}}, Y_{1,2,1}, \dots, Y_{1,2,a_{1,2}}, \dots, Y_{1,r,1}, \dots, Y_{1,r,a_{1,r}},\\
&X_{2,1}, \dots, X_{2,a_2}, Y_{2,1,1}, \dots, Y_{2,1,a_{2,1}}, Y_{2,2,1}, \dots, Y_{2,2,a_{2,2}}, \dots, Y_{2,r,1} \dots Y_{2,r,a_{2,r}}, \dots,\\
&X_{\beta,1}, \dots, X_{\beta,a_\beta}, Y_{\beta,1,1}, \dots, Y_{\beta,1,a_{\beta,1}}, Y_{\beta,2,1}, \dots, Y_{\beta,2,a_{\beta,2}}, \dots, Y_{\beta,r,1} \dots Y_{\beta,r,a_{\beta,r}}
\end{aligned}$$

To polarize the ideal I, we need to polarize the generators, which correspond to the r-paths in H. There are four types of r-paths in H: paths completely contained in an r-whisker (that is, exactly an r-whisker); paths partially in a r-whisker and partially in Γ; paths that start in a r-whisker, run through part of Γ, then end in another r-whisker; and paths that are completely in Γ.

First, consider an r-whisker $v_iy_{i,1}\dots y_{i,r}$. The generator associated to this path in I is $X_i^{a_i}Y_{i,1}^{a_{i,1}}Y_{i,2}^{a_{i,2}}\cdots Y_{i,r}^{a_{i,r}}$. When we polarize this generator, we obtain the following generator of $\widetilde{I}$.

$$X_{i,1}\cdots X_{i,a_i}Y_{i,1,1}\cdots Y_{i,1,a_{i,1}}Y_{i,2,1}\cdots Y_{i,2,a_{i,2}}\cdots Y_{i,r,1}\cdots Y_{i,r,a_{i,r}} \tag{3.7.1}$$

Next, consider an r-path $v_{i_1}v_{i_2}\cdots v_{i_p}v_jy_{j,1}\cdots y_{j,k}$ that starts in Γ and ends in an r-whisker. Note that here we have $p+k = r$. The generator of I associated to this path is $X_{i_1}^{c_{i_1,i_2}}X_{i_2}^{b_{i_1,i_2,i_3}}\cdots X_{i_p}^{b_{i_{p-1},i_p,j}}X_j^{a_j}Y_{j,1}^{a_{j,1}}\cdots Y_{j,k-1}^{a_{j,k-1}}Y_{j,k}^{t_{j,k}}$. When we polarize this generator for I, we obtain the next generator for $\widetilde{I}$.

$$X_{i_1,1}\cdots X_{i_1,c_{i_1,i_2}}X_{i_2,1}\cdots X_{i_2,b_{i_1,i_2,i_3}}\cdots X_{i_p,1}\cdots X_{i_p,b_{i_{p-1},i_p,j}} \cdot X_{j,1}\cdots X_{j,a_j}Y_{j,1,1}\cdots Y_{j,1,a_{j,1}}\cdots Y_{j,k-1,1}\cdots Y_{j,k-1,a_{j,k-1}}Y_{j,k,1}\cdots Y_{j,k,t_{j,k}} \tag{3.7.2}$$

Observe that the assumption $\omega(v_iv_j) \leqslant \min\{\omega(v_iy_{i,1}), \omega(v_jy_{j,1})\}$ for all $v_iv_j \in E(\Gamma)$ implies that we have $c_{i_1,i_2} \leqslant a_{i_1}$. Similarly, we have $b_{i_1,i_2,i_3} \leqslant a_{i_2}$, and the inequality $t_{j,k} \leqslant a_{j,k}$ is by construction. Thus, the generator (3.7.2) is in $\widetilde{S}$.

Next, consider an r-path $y_{j,q}\dots y_{j,1}v_jv_{m_1}\dots v_{m_l}v_iy_{i,1}\dots y_{i,p}$ that starts in an r-whisker, runs through part of Γ, and ends in another r-whisker. Note that we have $l \geqslant 0$ and $q+l+p+1 = r$. The generator in I associated to this type of path is the following.

$$Y_{j,q}^{t_{j,q}}Y_{j,q-1}^{a_{j,q-1}}\cdots Y_{j,1}^{a_{j,1}}X_j^{a_j}X_{m_1}^{b_{j,m_1,m_2}}\cdots X_{m_l}^{b_{m_{l-1},m_l,i}}X_i^{a_i}Y_{i,1}^{a_{i,1}}\cdots Y_{i,p-1}^{a_{i,p-1}}Y_{i,p}^{t_{i,p}}$$

When we polarize this generator we obtain the next generator for $\widetilde{I}$.

$$Y_{j,q,1}\cdots Y_{j,q,t_{j,q}}Y_{j,q-1,1}\cdots Y_{j,q-1,a_{j,q-1}}\cdots Y_{j,1,1}\cdots Y_{j,1,a_{j,1}} \cdot X_{j,1}\cdots X_{j,a_j}X_{m_1,1}\cdots X_{m_1,b_{j,m_1,m_2}}\cdots X_{m_l,1}\cdots X_{m_l,b_{m_{l-1},m_l,i}}X_{i,1}\cdots X_{i,a_i} \cdot Y_{i,1,1}\cdots Y_{i,1,a_{i,1}}\cdots Y_{i,p-1,1}\cdots Y_{i,p-1,a_{i,p-1}}Y_{i,p,1}\cdots Y_{i,p,t_{i,p}} \tag{3.7.3}$$

For the last type of generator, consider an r-path $v_{i_1}\dots v_{i_{r+1}}$ entirely in Γ. The generator in I associated to this path is the following.

$$X_{i_1}^{c_{i_1,i_2}}X_{i_2}^{b_{i_1,i_2,i_3}}\cdots X_{i_r}^{b_{i_{r-1},i_r,i_{r+1}}}X_{i_{r+1}}^{c_{i_r,i_{r+1}}}$$

When we polarize this generator we obtain the next generator for $\widetilde{I}$.

$$X_{i_1,1}\cdots X_{i_1,c_{i_1,i_2}}X_{i_2,1}\cdots X_{i_2,b_{i_1,i_2,i_3}}\cdots X_{i_{r+1},1}\cdots X_{i_{r+1},c_{i_r,i_{r+1}}} \tag{3.7.4}$$

Set $T = A[X_{1,1},\dots,X_{\beta,1}]$, and let J be the monomial ideal of T with the following generators. For each r-whisker $v_iy_{i,1}\dots y_{i,r}$, include the following generator.

$$X_{i,1}^{a_i+a_{i,1}+\cdots+a_{i,r}} \tag{3.7.5}$$

For each r-path $v_{i_1}v_{i_2}\cdots v_{i_p}v_jy_{j,1}\cdots y_{j,k}$ that starts in Γ and ends in an r-whisker, include the next generator.

$$X_{i_1,1}^{c_{i_1,i_2}}X_{i_2,1}^{b_{i_1,i_2,i_3}}\cdots X_{i_p,1}^{b_{i_{p-1},i_p,j}}X_{j,1}^{a_j+a_{j,1}+\cdots+a_{j,k-1}+t_{j,k}} \tag{3.7.6}$$

For each r-path $y_{j,q}\dots y_{j,1}v_jv_{m_1}\dots v_{m_l}v_iy_{i,1}\dots y_{i,p}$ that starts in an r-whisker, runs through part of Γ, and ends in another r-whisker, include the next generator.

$$X_{j,1}^{t_{j,q}+a_{j,q-1}+\cdots+a_{j,1}+a_j}X_{m_1,1}^{b_{j,m_1,m_2}}\cdots X_{m_l,1}^{b_{m_{l-1},m_l,i}}X_{i,1}^{a_i+a_{i,1}+\cdots+a_{i,p-1}+t_{i,p}} \tag{3.7.7}$$

For each r-path $v_{i_1}\dots v_{i_{r+1}}$ entirely in Γ, include the next generator.

$$X_{i_1,1}^{c_{i_1,i_2}}X_{i_2,1}^{b_{i_1,i_2,i_3}}\cdots X_{i_{r+1},1}^{c_{i_r,i_{r+1}}} \tag{3.7.8}$$

It is straightforward to show that the polarization of J is exactly $\widetilde{I}$: for $n = 1,2,3,4$, the polarization of the generator (3.7.$n+4$) of J is exactly the generator (3.7.n) of $\widetilde{I}$. Since J contains a power of each of the variables in T, namely (3.7.5), we conclude that T/J is artinian. Thus, the first paragraph of this proof implies that S/I is Cohen–Macaulay. □

Example 3.8. For the weighted graph G_ω in Example 1.4, Proposition 3.7 shows that $I_{1,\max}(G_\omega)$ is Cohen–Macaulay, and similarly for $I_{2,\max}(G''_{\omega''})$ in Example 3.5. See also Examples 3.12 and 3.13.

Note that the ideals $I_{r,\max}(G_\omega)$ and $I_{r,\max}(H_\lambda)$ in the next result live in different polynomial rings.

Lemma 3.9. *Let H_λ be a weighted graph obtained by pruning a sequence of r-pathless leaves from G_ω.*

(a) *The ideals $I_{r,\max}(G_\omega)$ and $I_{r,\max}(H_\lambda)$ have the same generators.*
(b) *The ideal $I_{r,\max}(G_\omega)$ is m-unmixed if and only if $I_{r,\max}(H_\lambda)$ is so.*
(c) *The ideal $I_{r,\max}(G_\omega)$ is Cohen–Macaulay if and only if $I_{r,\max}(H_\lambda)$ is so.*

Proof. Arguing by induction on the number of r-pathless leaves being pruned from G_ω, we assume that H_λ is obtained by pruning a single r-pathless leaf v_i from G_ω.

(a) By Lemma 3.3(a), the set of r-paths in G is the same as the set of r-paths in H, and $\lambda(e) = \omega(e)$ for each edge $e \in E(H) \subseteq E(G)$. The claim about the generators now follows directly.

(b) This follows from Theorem 2.7(b) and Lemma 3.3(c).

(c) Part (a) implies that $(S'/I_{r,\max}(H_\lambda))[X] \cong S/I_{r,\max}(G_\omega)$, where $S' := A[X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n]$. It follows that $S/I_{r,\max}(G_\omega)$ is Cohen–Macaulay if and only if $S'/I_{r,\max}(H_\lambda)$ is Cohen–Macaulay, as desired. □

The next result compares directly to Theorem B from the introduction, though it does not assume that G is a tree.

Proposition 3.10. *Assume that H_λ is obtained by pruning a sequence of r-pathless leaves from G_ω and that H_λ is an r-path suspension of a weighted graph Γ_μ. With notation as in Remark 3.6, the following conditions are equivalent:*

(i) *$I_{r,\max}(G_\omega)$ is Cohen–Macaulay;*
(ii) *$I_{r,\max}(G_\omega)$ is m-unmixed; and*
(iii) *for all $v_iv_j \in E(\Gamma_\mu)$ we have $\omega(v_iv_j) \leqslant \min\{\omega(v_iy_{i,1}), \omega(v_jy_{j,1})\}$.*

Proof. The case $r = 1$ is handled in [11, Theorem 5.7], so we assume that $r \geqslant 2$ for the remainder of the proof. The implication (i) $\Longrightarrow$ (ii) always holds.

(ii) $\Longrightarrow$ (iii) Assume that $I_{r,\max}(G_\omega)$ is m-unmixed. It follows from Lemma 3.9(b) that $I_{r,\max}(H_\lambda)$ is also unmixed. From an analysis of the r-paths of H as in the proof of Proposition 3.7, it is straightforward to show that $V(\Gamma_\mu)$ is a minimal r-path vertex cover of H. (It covers all the paths, and the r-whiskers show that it is minimal.) Let $\tau: V(\Gamma_\mu) \to \mathbb{N}$ be the constant function $\tau(v_i) = 1$. Lemma 1.11 implies that there is a minimal weighted r-path vertex cover (W'', σ'') of H_λ such that $(W'', \sigma'') \leqslant (V(\Gamma_\mu), \tau)$. The minimality of $V(\Gamma_\mu)$ implies that $W'' = V(\Gamma_\mu)$, so $(V(\Gamma_\mu), \sigma'')$ is a minimal weighted r-path vertex cover of H_λ. The unmixedness condition implies that every minimal weighted r-path vertex cover of H_λ has cardinality $|V(\Gamma_\mu)|$.

We proceed by contradiction. Suppose that there is an edge $v_iv_j \in E(\Gamma_\mu)$ such that $\omega(v_iv_j) > \min\{\omega(v_iy_{i,1}), \omega(v_jy_{j,1})\}$. We produce a contradiction by showing that there exists a minimal weighted r-path vertex cover (W, σ) of H_λ such that $|W| > |V(\Gamma_\mu)|$. Assume by symmetry that

$$a := \omega(v_iy_{i,1}) = \min\{\omega(v_iy_{i,1}), \omega(v_jy_{j,1})\} < \omega(v_iv_j) =: b.$$

Set $c = \omega(v_j y_{j,1})$ and $a' := \omega(y_{i,r-1}y_{i,r})$ and $c' := \omega(y_{j,r-1}y_{j,r})$. The following diagram (where the column represents Γ, and the rows represent the r-whiskers in H) is our guide for constructing an approximation of (W, σ).

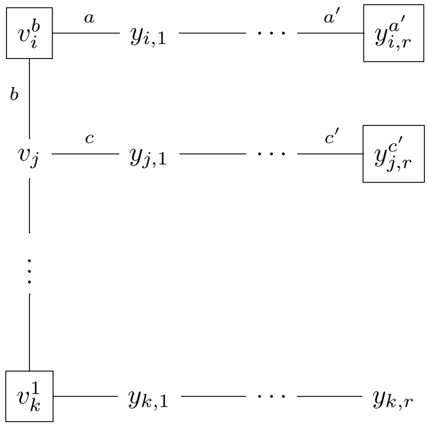

Set $W = \{v_k \mid k \neq j\} \cup \{y_{i,r}, y_{j,r}\}$ and define $\sigma : W \to \mathbb{N}$ by

$$\sigma(v_k) = \begin{cases} 1 & \text{if } k \neq i \\ b & \text{if } k = i \end{cases}$$
$$\sigma(y_{i,r}) = a'$$
$$\sigma(y_{j,r}) = c'.$$

It is straightforward to show that (W, σ) is a weighted r-path vertex cover of H_λ. Lemma 1.11 provides a minimal weighted r-path vertex cover (W', σ') of G_ω such that $(W', \sigma') \leqslant (W, \sigma)$.

We claim that $W' = W$. (This then yields the promised contradiction, completing the proof of this implication.) To this end, first note that we have $W' \subseteq W$, by assumption. So, we need to show that $W' \supseteq W$. We cannot remove the vertex $y_{j,r}$ from W, since that would leave the r-path $v_j y_{j,1} \dots y_{j,r}$ uncovered. Thus, we have $y_{j,r} \in W'$. Similarly, for $k \neq i, j$ the vertex v_k cannot be removed, so $v_k \in W'$. If we remove the vertex v_i, the r-path $v_j v_i y_{i,1} \dots y_{i,r-1}$ is not covered, so $v_i \in W'$. Since $\sigma(v_i) = b > a$, the vertex v_i does not cover the r-path $v_i y_{i,1} \dots y_{i,r}$. It follows that the vertex $y_{i,r}$ cannot be removed. Thus, we have $y_{i,r} \in W'$, and it follows that $W' = W$, as claimed.

(iii) $\Longrightarrow$ (i) Assuming condition (iii), Proposition 3.7 implies that $I_{r,\max}(H_\lambda)$ is Cohen–Macaulay, so Lemma 3.9(c) implies that $I_{r,\max}(G_\omega)$ is as well. □

The next result contains Theorem B from the introduction.

Theorem 3.11. *Assume that G_ω is a weighted tree. Then the following conditions are equivalent:*

(i) $I_{r,\max}(G_\omega)$ *is Cohen–Macaulay;*
(ii) $I_{r,\max}(G_\omega)$ *is m-unmixed; and*
(iii) *there is a weighted tree Γ_μ and an r-path suspension H_λ of Γ_μ such that H_λ is obtained by pruning a sequence of r-pathless leaves from G_ω and for all $v_i v_j \in E(\Gamma_\mu)$ we have $\omega(v_i v_j) \leqslant \min\{\omega(v_i y_{i,1}), \omega(v_j y_{j,1})\}$.*

When G_ω satisfies the above equivalent conditions, the graph H can be constructed by pruning r-pathless leaves from G until no more r-pathless leaves remain.

Proof. The implications (iii) $\Longrightarrow$ (i) $\Longrightarrow$ (ii) are from Proposition 3.10. For the implication (ii) $\Longrightarrow$ (iii), assume that $I_{r,\max}(G_\omega)$ is m-unmixed. Since G is finite, prune a sequence of r-pathless leaves from G_ω to obtain a weighted subgraph H_λ that has no r-pathless leaves. Lemma 3.9(b) implies that $I_{r,\max}(H_\lambda)$ is m-unmixed, so Lemma 2.11 implies that $I_r(H)$ is m-unmixed. Thus, H is an r-path suspension of a tree Γ by [1, Theorem 3.8 and Remark 3.9]. Finally, Proposition 3.10 implies that $\omega(v_iv_j) \leqslant \min\{\omega(v_iy_{i,1}), \omega(v_jy_{j,1})\}$ for all $v_iv_j \in E(\Gamma_\mu)$. $\square$

Example 3.12. Consider the weighted graph G_ω in Example 1.4. Then $I_{r,\max}(G_\omega)$ is Cohen–Macaulay if and only if $r \neq 2, 3$, as follows. Example 3.8 deals with the case $r = 1$.

For $r > 5$, the ideal $I_{r,\max}(G_\omega)$ is trivially Cohen–Macaulay since G has no r-paths. (One can also deduce this from Lemma 3.9 since every leaf is r-pathless.)

This graph has a single 4-path, so $S/I_{4,\max}(G_\omega)$ is a hypersurface, hence Cohen–Macaulay. One can also deduce this from Theorem 3.11 by pruning the 4-pathless leaf v_6 to obtain the weighted 4-path H_λ in Example 3.2. Since H_λ is a 4-path suspension of the trivial graph v_1, the desired conclusion follows from Theorem 3.11.

For $r = 2, 3$, the ideal $I_{r,\max}(G_\omega)$ is not Cohen–Macaulay by Theorem 3.11. To see this, observe that G does not have any r-pathless leaves and is not an r-suspension for $r = 2, 3$.

Example 3.13. Arguing as in Example 3.12, we have the following for the weighted graphs $G'_{\omega'}$ and $G''_{\omega''}$ of Example 3.5. The ideal $I_{r,\max}(G'_{\omega'})$ is Cohen–Macaulay if and only if $r \geqslant 6$, and $I_{r,\max}(G''_{\omega''})$ is Cohen–Macaulay if and only if $r \neq 1, 3, 4, 5, 6, 7$.

4. Cohen–Macaulay weighted complete graphs when $r = 2$

Assumption. Throughout this section, K^n_ω is a weighted n-clique, and A is a field.

In this section, we prove Theorem C from the introduction characterizing Cohen–Macaulayness of n-cliques in the context of weighted path ideals for the function $f = \max$ with $r = 2$. We begin with two results about arbitrary f and r. Note that the assumption $r < n$ causes no loss of information since, when $r \geqslant n$, we have $I_{r,f}(G_\omega) = 0$.

Lemma 4.1. *If (W, σ) is an f-weighted r-path vertex cover for K^n_ω where $r < n$, then $|W| \geqslant n - r$.*

Proof. Suppose that $|W| < n - r$ and assume that $v_{i_1}, \ldots, v_{i_{r+1}} \notin W$. Then the path $v_{i_1} \ldots v_{i_{r+1}}$ in K^n_ω is not covered by (W, σ). $\square$

Lemma 4.2. *Assume that $r < n$, and consider an arbitrary subset $W \subseteq V$ with $|W| = n - r$. Then there is a function $\sigma'' : W \to \mathbb{N}$ such that (W, σ'') is a minimal f-weighted r-path vertex cover for K^n_ω.*

Proof. Using the inclusion–exclusion principle, it is straightforward to show that W is an r-path vertex cover of K^n. The trivial weight $\sigma : W \to \mathbb{N}$ with $\sigma(v) = 1$ for all $v \in W$ makes (W, σ) into an f-weighted r-path vertex cover of K^n_ω. Lemma 1.11 yields a minimal f-weighted r-path vertex cover (W'', σ'') of G_ω such that $(W'', \sigma'') \leqslant (W, \sigma)$. Lemma 4.1 shows that $|W''| \geqslant n - r = |W|$. Since $W'' \subseteq W$, we must have $W = W''$, as desired. $\square$

For the remainder of this section, we focus on the case $f = \max$.

Proposition 4.3. *If $r \leqslant n$, then* $\dim(S/I_{r,\max}(K^n_\omega)) = r$.

Proof. If $r = n$, then $I_{r,\max}(K^n_\omega) = 0$ and therefore $\dim(S/I_{r,\max}(K^n_\omega)) = \dim(S) = n = r$, as claimed. Assume for the rest of the proof that $r < n$. Lemma 4.1 implies that for every weighted r-path vertex cover (W, σ) we have $|W| \geqslant n - r$. Furthermore, Lemma 4.2 implies that there is a minimal weighted r-path vertex cover (W, σ) with $|W| = n - r$. Thus, the desired conclusion follows from Theorem 2.7(b). □

For the rest of the section, we focus on the case $r = 2$.

The next result characterizes the weighted 3-cliques K^3_ω such that $I_{2,\max}(K^3_\omega)$ is Cohen–Macaulay. Note that these cliques are key for the characterization of Cohen–Macaulayness of larger n-cliques in Theorem C. Also, smaller n-cliques are very small trees that always give Cohen–Macaulay ideals; argue as in Example 3.12.

Proposition 4.4. *Consider a weighted 3-clique K^3_ω, which we assume by symmetry to be of the following form*

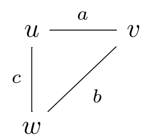

with weights a, b, and c such that $a \leqslant b \leqslant c$. Then the following conditions are equivalent:

(i) *The ideal $I_{2,\max}(K^3_\omega)$ is Cohen–Macaulay;*
(ii) *The ideal $I_{2,\max}(K^3_\omega)$ is unmixed; and*
(iii) *We have $a = b$, that is, $a = b \leqslant c$.*

Proof. First, we note that

$$I_{2,\max}(K^3_\omega) = (X^aY^bZ^b, X^cY^bZ^c, X^cY^aZ^c)S = (X^aY^bZ^b, X^cY^aZ^c)S.$$

The implication (i) $\Longrightarrow$ (ii) is standard.

(ii) $\Longrightarrow$ (iii) We argue by contrapositive. Assume that $a < b$. If $a < b = c$, then it is straightforward to show that the weighted 2-path ideal decomposes irredundantly as follows:

$$I_{2,\max}(K^3_\omega) = (X^aY^bZ^b, X^bY^aZ^b)S = (X^a)S \cap (Y^a)S \cap (Z^b)S \cap (X^b, Y^b).$$

In particular, this ideal is mixed. When $a < b < c$, the weighted 2-path ideal is also mixed because of the following irredundant decomposition:

$$\begin{aligned} I_{2,\max}(K^3_\omega) &= (X^aY^bZ^b, X^cY^aZ^c)S \\ &= (X^a)S \cap (Y^a)S \cap (Z^b)S \cap (X^c, Y^b)S \cap (Y^b, Z^c)S. \end{aligned}$$

(iii) $\Longrightarrow$ (i) If $a = b$, then we have

$$I_{2,\max}(K^3_\omega) = (X^aY^aZ^a)S \tag{4.4.1}$$

which is generated by a regular element and is therefore Cohen–Macaulay. □

Remark 4.5. The first display in the proof of Proposition 4.4 shows that the generating sequence used to define $I_{r,f}(G_\omega)$ can be redundant, i.e., non-minimal.

Our next result uses the following information about colon ideals.

Remark 4.6. Let I be a monomial ideal in S, that is an ideal of S generated by a list $g_1, \dots, g_t$ of monomials in the variables $X_1, \dots, X_n$. Given another monomial $h \in S$, it is straightforward to show that the colon ideal $(I :_S h)$ is generated by the following list of monomials: $g_1/\gcd(g_1, h), \dots, g_t/\gcd(g_t, h)$.

The next result contains one implication of Theorem C from the introduction. Note that the 2-path Cohen–Macaulay weighted 3-cliques are characterized in Proposition 4.4.

Theorem 4.7. *Let $n \geqslant 3$. Assume that every induced weighted sub-3-clique $K^3_{\omega'}$ of K^n_ω has $I_{2,\max}(K^3_{\omega'})$ Cohen–Macaulay. Then $I_{2,\max}(K^n_\omega)$ is also Cohen–Macaulay.*

Proof. Set $I := I_{2,\max}(K^n_\omega)$. Note that our hypothesis on the induced weighted sub-3-cliques of K^n_ω imply that I is generated by the following set of monomials:

$$\left\{X_i^{a_{i,j,k}} X_j^{a_{i,j,k}} X_k^{a_{i,j,k}} \;\middle|\; i < j < k \text{ and } a_{i,j,k} = \min\bigl(\omega(e_ie_j), \omega(e_ie_k), \omega(e_je_k)\bigr)\right\}.$$

Indeed, this follows from Lemma 2.13 and the description of $I_{2,\max}(K^3_{\omega'})$ from Eq. (4.4.1) in the proof of Proposition 4.4. In particular, the generators of I are determined by the induced weighted sub-3-cliques of K^n_ω.

We proceed by induction on n. The base case $n = 3$ is trivial.

For the inductive step, assume that $n \geqslant 4$ and the following: for every weighted $(n-1)$-clique K^{n-1}_μ, if every induced weighted sub-3-clique $K^3_{\mu'}$ of K^{n-1}_μ has $I_{2,\max}(K^3_{\mu'})$ Cohen–Macaulay, then $I_{2,\max}(K^{n-1}_\mu)$ is also Cohen–Macaulay. Set $R := S/I_{2,\max}(K^n_\omega)$ and $a := \min\{\omega(v_iv_j) \text{ over all } i \text{ and } j\}$. Assume by symmetry that $\omega(v_1v_2) = a$. Let $K^{n-1}_{\omega'}$ denote the weighted sub-clique of K^n_ω induced by $V \smallsetminus \{v_1\}$. Set $S' = A[X_2, \dots, X_n]$. Lemma 2.12 implies that $R' := R/(X_1)R \cong S'/I_{2,\max}(K^{n-1}_{\omega'})$. Since $K^{n-1}_{\omega'}$ has the same condition on the induced weighted sub-3-cliques, R' is Cohen–Macaulay by the inductive hypothesis. Note that Proposition 4.3 says that $\dim(R') = 2$. We consider the following short exact sequence:

$$0 \to X_1^a R \to R \to R/X_1^a R \to 0. \tag{4.7.1}$$

Since a is the smallest edge weight on K^n_ω, we have $R/X_1^aR \cong R'[T_1]/(T_1^a)$, which is Cohen–Macaulay of dimension 2. As $\dim(R) = 2$, in order to show that R is Cohen–Macaulay, it suffices to show that $\operatorname{depth}(R) \geqslant 2$. Applying the Depth Lemma to the sequence (4.7.1), we see that it suffices to show that $\operatorname{depth}_S(X_1^aR) = 2$.

Case 1. Assume that $\omega(v_1v_i) = a$ for all $i = 2, \dots, n$.

Claim 1. $(I :_S X_1^a) = (X_i^aX_j^a \mid 1 < i < j \leqslant n)S$. *For the containment $\supseteq$, let $1 < i < j \leqslant n$. Our assumptions on a imply that the generator of I corresponding to the sub-clique induced by v_1, v_i, v_j is $X_1^aX_i^aX_j^a$. It follows that the element $X_i^aX_j^a$ is in $(I :_S X_1^a)$, as desired. For the reverse containment, note that the generators for I are of the form $X_p^\alpha X_q^\alpha X_r^\alpha$ such that $p < q < r$ and $\alpha \geqslant a$. The corresponding generator of $(I :_S X_1^a)$ when $p = 1$ is $X_1^{\alpha-a}X_q^\alpha X_r^\alpha \in (X_q^aX_r^a)$. When $p \neq 1$ we have $X_p^\alpha X_q^\alpha X_r^\alpha \in (X_q^aX_r^a)$. Therefore the claim holds.*

Also, we have

$$X_1^aR \cong R/\operatorname{Ann}_R(X_1^a) \cong S/(I :_S X_1^a) \cong \bigl(S'/I_{1,\max}(K^{n-1}_a)\bigr)[X_1]$$

where the graph K_a^{n-1} has constant weight a on each edge; this is by Claim 1. The proof of [11, Proposition 5.2] shows that $S'/I_{1,\max}(K_a^{n-1})$ is Cohen–Macaulay of dimension 1. Therefore $X_1^a R \cong (S'/I_{1,\max}(K_a^{n-1}))[X_1]$ is Cohen–Macaulay of dimension 2.

Case 2. Assume that $\omega(v_1v_2) = a < \omega(v_1v_i)$ for some $i > 2$. This assumption implies that there exists a subset $W \subseteq V$ such that $v_1, v_i \in W$ and for each $v_j, v_k \in W$ we have $\omega(v_jv_k) > a$. By the finiteness of the graph K^n, there exists a maximal such set W. Note that $|W| \geqslant 2$.

Claim 2. *for all $v_p \in V \smallsetminus W$ and all $v_j \in W$, we have $\omega(v_jv_p) = a$. Suppose by way of contradiction that $\omega(v_jv_p) > a$. Let $v_k \in W$ such that $v_k \neq v_j$. By assumption, we have $\omega(v_jv_k) > a$ and $\omega(v_jv_p) > a$. Let $K_{\omega'}^3$ be the weighted sub-3-clique of K_ω^n induced by v_j, v_k, v_p. By assumption, the ideal $I_{2,\max}(K_{\omega'}^3)$ is Cohen–Macaulay, so Proposition 4.4 implies that either $\omega(v_kv_p) \geqslant \omega(v_jv_p) > a$ or $\omega(v_kv_p) \geqslant \omega(v_jv_k) > a$. Since v_k was chosen arbitrarily, the set $W \cup \{v_p\}$ satisfies the condition for W, contradicting the maximality of W.*

Let λ be a new weight on K^n such that

$$\lambda(v_\alpha v_\beta) = \begin{cases} \omega(v_\alpha v_\beta) & \text{if } v_\alpha, v_\beta \in W \\ a & \text{if } v_\alpha \notin W \text{ or } v_\beta \notin W. \end{cases}$$

Observe that this implies for $v_j, v_k \in W$ and $v_p, v_q \notin W$ we have

$$\lambda(v_jv_k) = \omega(v_jv_k)$$
$$\lambda(v_jv_p) = a = \omega(v_jv_p)$$
$$\lambda(v_pv_q) \text{ may be different from } \omega(v_pv_q).$$

Hence the graph K_λ^n satisfies the induced weighted sub-3-clique assumption. (The four types of induced weighted sub-3-cliques are displayed next, with $v_j, v_k, v_l \in W$ and $v_p, v_q, v_r \notin W$.)

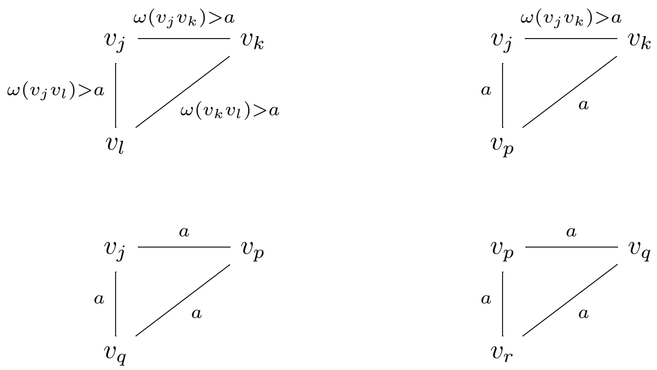

Since $\omega(v_1v_2) = a$, we have $v_2 \notin W$. Thus $\lambda(v_2v_l) = a$ for all $l \neq 2$. Hence the ideal $J := I_{2,\max}(K_\lambda^n)$ is Cohen–Macaulay by Case 1. Note that the condition $\omega(e) \geqslant \lambda(e)$ for each edge e implies that $I \subseteq J$.

Claim 3. *We have the equality $(I :_S X_1^a) = (J :_S X_1^a)$. The containment $\subseteq$ follows from the fact that $I \subseteq J$. For the reverse containment, recall that the generators for the ideals I and J are determined by the induced sub-3-cliques of K^n. For the first three sub-3-cliques displayed above, the corresponding generators of I and J are the same. Therefore, the generators in the colon ideals produced by these generators are the same; see Remark 4.6. In the case of the fourth induced sub-3-clique, the associated generator for J is*

$X_p^a X_q^a X_r^a$. Since $p, q, r \neq 1$, the associated generator for $(J :_S X_1^a)$ is $X_p^a X_q^a X_r^a \in (X_p^a X_q^a)S \subseteq (I :_S X_1^a)$; the last containment is explained as follows. The existence of distinct elements $v_p, v_q, v_r \in V \setminus W$ provides a sub-3-clique induced by v_1, v_p, v_q, which is of the third type, with corresponding generator for the colon ideals being $X_p^a X_q^a$. This establishes Claim 3.

Lastly, Case 1 shows that $\operatorname{depth}_S((X_1^a)S/J) = 2$. Claim 3 implies that

$$(X_1^a)S/J \cong S/(J :_S X_1^a) = S/(I :_S X_1^a) \cong (X_1^a)S/I = X_1^a R.$$

Therefore $\operatorname{depth}_S(X_1^a R) = 2$, as desired. □

The converse of Theorem 4.7 is more complicated. We break the proof into (hopefully) manageable pieces, culminating in Theorem 4.12.

Proposition 4.8. *Let $n \geqslant 3$ and assume that K_ω^n contains an induced weighted sub-3-clique of the form*

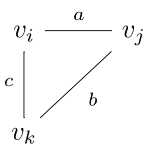

with weights a, b, and c such that $a < b < c$. Then $I_{2,\max}(K_\omega^n)$ is mixed. In particular, $I_{2,\max}(K_\omega^n)$ is not Cohen–Macaulay.

Proof. By symmetry, assume without loss of generality that $i = 1$, $j = 2$, and $k = 3$. By Theorem 2.7(b), it suffices to exhibit two minimal weighted 2-path vertex covers for K_ω^n whose cardinalities are not equal. Since $\dim(S/I_{2,\max}(K_\omega^n)) = 2$ by Proposition 4.3, we know that K_ω^n has a minimal weighted 2-path vertex cover of size $n - 2$. Thus, it suffices to find a minimal weighted 2-path vertex cover of size $n - 1$.

Consider the weighted set $\{v_2^b, v_3^c, v_4^1, \dots, v_n^1\}$. In light of the assumptions on a, b, and c, it is straightforward to show that this is a weighted 2-path vertex cover for K_ω^n. We show that it gives rise to a minimal one of the form $\{v_2^b, v_3^c, v_4^{r_4}, \dots, v_n^{r_n}\}$. Since $c > b$, the weighted path

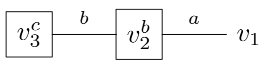

is covered only by the weighted vertex v_2^b. If the weight b on this vertex were increased, then this weighted path would no longer be covered. Thus, the vertex v_2 cannot be removed from the cover, and its weight cannot be increased. Similarly, the weighted path $v_3 v_1 v_2$ shows that the vertex v_3 cannot be removed from the cover, and its weight cannot be increased. Lastly, for $j \geqslant 4$ the weighted path $v_2 v_1 v_j$ is only covered by v_j^1. Thus, the vertex v_j cannot be removed from the cover; however, its weight can be increased. □

Remark 4.9. The weighted 2-path vertex cover $\{v_2^b, v_3^c, v_4^1, \dots, v_n^1\}$ in the previous proof is not incredibly mysterious. Indeed, the induced weighted sub-3-clique

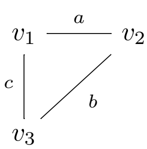

has $\{v_2^b, v_3^c\}$ as a minimal weighted 2-path vertex cover. (This can be checked readily as in the previous proof. Alternately, it follows from the proof of Proposition 4.4; see the discussion in Example 2.10.) The given cover for K_ω^n is built from this one.

When $a < b = c$, one might guess that the vertex cover $\{v_1^b, v_2^b, v_4^1, \dots, v_n^1\}$ can be used to show that $I_{2,\max}(K_\omega^n)$ is mixed in this case as well. However, the next example shows that this is not the case.

Example 4.10. Consider the following weighted 4-clique.

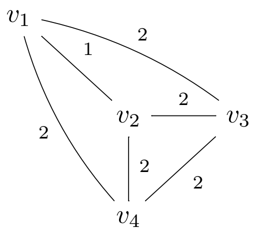

It is straightforward to show that we have the following.

$$\begin{aligned} I_{2,\max}(K_\omega^4) &= (X_1X_2^2X_3^2, X_1^2X_2^2X_3, X_1X_3^2X_4^2, X_1^2X_3X_4^2, X_1^2X_2^2X_4^2, X_2^2X_3^2X_4^2)S \\ &= (X_1, X_2^2)S \cap (X_1^2, X_3^2)S \cap (X_1, X_4^2)S \\ &\quad \cap (X_2^2, X_3)S \cap (X_2^2, X_4^2)S \cap (X_3, X_4^2)S \end{aligned}$$

The decomposition here shows that $I_{2,\max}(K_\omega^4)$ is unmixed. However, Theorem 4.12 below shows that it is not Cohen–Macaulay because the weighted sub-3-clique induced by v_1, v_2, v_3 is not Cohen–Macaulay; see Proposition 4.4.

Proposition 4.11. *Assume that $I_{2,\max}(K_\omega^n)$ is unmixed, and that K_ω^n has an induced weighted sub-3-clique $K_{\omega'}^3$ such that $I_{2,\max}(K_{\omega'}^3)$ is not Cohen–Macaulay. Then K_ω^n has an induced weighted sub-4-clique of the form*

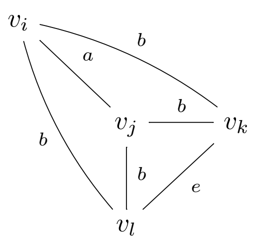

such that $a < b$.

Proof. Without loss of generality, assume that the non-Cohen–Macaulay induced weighted sub-3-clique is on the vertices v_1, v_2, v_3 as follows

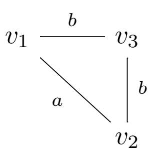

with $a < b$. Note that it must have this form by Propositions 4.4 and 4.8, because of our unmixedness assumption. Assume without loss of generality that b is maximal among all weights occurring in a non-Cohen–Macaulay induced sub-3-clique.

It is readily shown that the set $\{v_1^b, v_2^b, v_4^1, \ldots, v_n^1\}$ is a weighted 2-path vertex cover. As in the proof of Proposition 4.8, the path $v_3v_1v_2$ shows that the vertex v_1^b cannot be removed from this cover, and its weight cannot be increased. Similarly, the path $v_1v_2v_3$ shows that the vertex v_2^b cannot be removed from this cover, and its weight cannot be increased. Because of our unmixedness assumption, Theorem 2.7(b) and Proposition 4.3 imply that every minimal weighted 2-path vertex cover of K_ω^n has cardinality $n-2$. Since the given cover has size $n-1$, one of the vertices v_4 through v_n can be removed to create a weighted 2-path vertex cover. Reorder the vertices if necessary so that v_4 is the vertex that can be removed. Lemma 1.11 shows that this gives rise to a minimal weighted 2-path vertex cover of the form $\{v_1^b, v_2^b, v_5^{r_5}, \ldots, v_n^{r_n}\}$.

Label the induced weighted subgraph with vertices v_1, v_2, v_3, v_4 as follows.

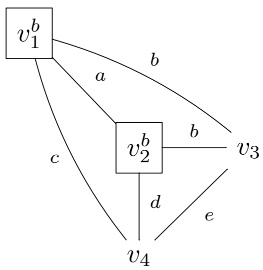

Since $a < b$, the path $v_1v_2v_4$ must be covered by v_2^b. Thus $b \leqslant d$. Similarly, the vertex v_1^b must cover the path $v_2v_1v_4$, so $b \leqslant c$. Thus, we have $a < b \leqslant c, d$, so the weighted sub-3-clique induced by v_1, v_2, v_4 is not Cohen–Macaulay. Proposition 4.8 implies that $c = d$, and the maximality of b implies that $c \leqslant b$, that is $c = b$. Thus, the above sub-4-clique has the desired form. □

The next result contains the remainder of Theorem C from the introduction.

Theorem 4.12. *Assume that K_ω^n contains at least one induced weighted sub-3-clique $K_{\omega'}^3$ such that $I_{2,\max}(K_{\omega'}^3)$ is not Cohen–Macaulay. Then $I_{2,\max}(K_\omega^n)$ is not Cohen–Macaulay.*

Proof. If $I := I_{2,\max}(K_\omega^n)$ is mixed, then we are done. So, we assume that I is unmixed. Theorem 2.7(b) and Proposition 4.3 imply that every minimal weighted 2-path vertex cover of K_ω^n has cardinality $n-2$. Also, Lemma 4.2 shows that every subset of V of cardinality $n-2$ occurs as a minimal weighted 2-path vertex cover.

Every induced weighted sub-3-clique of K_ω^n has the form

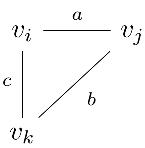

with $a \leqslant b \leqslant c$. By assumption, K_ω^n contains at least one such sub-clique with $a < b \leqslant c$; see Proposition 4.4. Furthermore, Proposition 4.8 implies that every such sub-clique has $a < b = c$.

Using Proposition 4.11 and reordering the vertices if necessary, we obtain an induced weighted subgraph of the following form

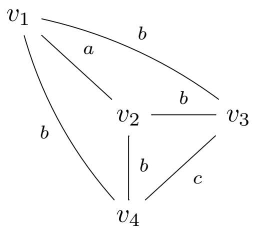

(4.12.1)

with $a < b$.

Using Theorem 2.7(b) we have a minimal m-irreducible decomposition

$$I = \bigcap (X_{j_1}^{\beta_1}, X_{j_2}^{\beta_2}, \dots, X_{j_{n-2}}^{\beta_{n-2}})S \tag{4.12.2}$$

where the intersection is taken over all minimal weighted 2-path vertex covers $\{v_{j_1}^{\beta_1}, v_{j_2}^{\beta_2}, \dots, v_{j_{n-2}}^{\beta_{n-2}}\}$ of K_ω^n. We set

$$I_1 := \bigcap (X_1^{\alpha_1}, X_{k_1}^{\alpha_{k_1}}, \dots, X_{k_{n-3}}^{\alpha_{k_{n-3}}})S \tag{4.12.3}$$

where the intersection is taken over all minimal weighted 2-path vertex covers of K_ω^n that contain the vertex v_1. Next, set

$$I_* := \bigcap_{j_i \neq 1} (X_{j_1}^{\beta_1}, X_{j_2}^{\beta_2}, \dots, X_{j_{n-2}}^{\beta_{n-2}})S \tag{4.12.4}$$

where the intersection is taken over all minimal weighted 2-path vertex covers that do not contain the vertex v_1. By definition, this yields $I = I_1 \cap I_*$. Moreover, the first paragraph of this proof implies that each of these intersections is taken over a non-empty index set.

Note that the irredundancy of the intersection in (4.12.2) implies that the two subsequence intersections are also irredundant. It follows that the maximal ideal $\mathfrak{m} = (X_1, \dots, X_n)S$ is not associated to I_1 and is not associated to I_*. Thus, we have $1 \leqslant \operatorname{depth}(S/I_1) \leqslant \dim(S/I_1) = 2$ and $1 \leqslant \operatorname{depth}(S/I_*) \leqslant \dim(S/I_*) = 2$. Since we have $\dim(S/I) = 2$, it remains to show that $\operatorname{depth}(S/I) = 1$.

Consider the short exact sequence

$$0 \to S/I \to S/I_1 \oplus S/I_* \to S/(I_1 + I_*) \to 0.$$

By the Depth Lemma (or a routine long-exact-sequence argument), in order to show that $\operatorname{depth}(S/I) = 1$, it suffices to show that $\operatorname{depth}(S/(I_1 + I_*)) = 0$, that is, that $\mathfrak{m}$ is associated to $I_1 + I_*$.

From the decompositions (4.12.3) and (4.12.4), we have

$$I_1 + I_* = \bigcap \bigcap_{j_i \neq 1} \left[(X_1^{\alpha_1}, X_{k_1}^{\alpha_{k_1}}, \dots, X_{k_{n-3}}^{\alpha_{k_{n-3}}})S + (X_{j_1}^{\beta_1}, X_{j_2}^{\beta_2}, \dots, X_{j_{n-2}}^{\beta_{n-2}})S \right] \tag{4.12.5}$$

where the first intersection is taken over all minimal weighted 2-path vertex covers that contain the vertex v_1, and the second intersection is taken over all minimal weighted 2-path vertex covers that do not contain the vertex v_1; see, e.g., [7, Lemma 2.7]. Note that this is an m-irreducible decomposition, though it may be redundant. We need to show that there is an ideal in this intersection of the form $(X_1^{\delta_1}, X_2^{\delta_2}, \dots, X_n^{\delta_n})S$ that is irredundant in the intersection.

Given the sub-clique (4.12.1), it is straightforward to show that there are minimal weighted 2-path vertex covers of K_ω^n of the form $\{v_1^b, v_2^b, v_5^{\alpha_5}, \dots, v_n^{\alpha_n}\}$ and $\{v_3^b, v_4^b, v_5^{\beta_5}, \dots, v_n^{\beta_n}\}$. In particular, the ideal $P_1 :=$

$(X_1^b, X_2^b, X_5^{\alpha_5}, \dots, X_n^{\alpha_n})S$ occurs in the decomposition (4.12.3), and the ideal $P_* := (X_3^b, X_4^b, X_5^{\beta_5}, \dots, X_n^{\beta_n})S$ occurs in the decomposition (4.12.4). Thus, the ideal

$$P_1 + P_* = \big(X_1^b, X_2^b, X_3^b, X_4^b, X_5^{\gamma_5}, X_6^{\gamma_6}, \dots, X_n^{\gamma_n}\big)S$$

is in the intersection (4.12.5), where $\gamma_i = \min\{\alpha_i, \beta_i\}$.

Let Q_1 be an ideal occurring in the intersection (4.12.3), and let Q_* be an ideal occurring in the intersection (4.12.4). Suppose that

$$\big(X_{t_1}^{\zeta_1}, X_{t_2}^{\zeta_2}, \dots, X_{t_g}^{\zeta_g}\big)S = Q_1 + Q_* \subseteq P_1 + P_* \quad \text{with } g \leqslant n-1. \tag{4.12.6}$$

Claim 1. *We have $Q_* = (X_3^{\eta_3}, X_4^{\eta_4}, X_5^{\eta_5}, X_6^{\eta_6}, \dots, X_n^{\eta_n})S$ for some $\eta_3, \dots, \eta_n$. By assumption, we have $Q_* = (X_{j_1}^{\eta_1}, X_{j_2}^{\eta_2}, \dots, X_{j_{n-2}}^{\eta_{n-2}})S$ with $j_i > 1$ for $i = 1, \dots, n-2$. It suffices to show that $j_i \neq 2$ for all i. Suppose that $j_i = 2$ for some i. Given the conditions on the generators of Q_*, there must be an index $k \neq 1$ such that $j_i \neq k$ for all i. Then $v_2^{\eta_2}$ must cover the path $v_2v_1v_k$. This implies that $\eta_2 \leqslant a$. On the other hand, since*

$$X_2^{\eta_2} \in Q_1 + Q_* \subseteq P_1 + P_* = \big(X_1^b, X_2^b, X_3^b, X_4^b, X_5^{\gamma_5}, X_6^{\gamma_6}, \dots, X_n^{\gamma_n}\big)S$$

we have $\eta_2 \geqslant b > a \geqslant \eta_2$, a contradiction. This establishes Claim 1.

Claim 2. *We have $Q_1 = (X_1^{\mu_1}, X_{m_1}^{\mu_{m_1}}, \dots, X_{m_{n-3}}^{\mu_{m_{n-3}}})S$ for some $\mu_1, \mu_{m_1}, \dots, \mu_{m_{n-3}}$ with $m_i > 2$ for all i. By assumption, we have $Q_1 = (X_1^{\mu_1}, X_{m_1}^{\mu_{m_1}}, \dots, X_{m_{n-3}}^{\mu_{m_{n-3}}})S$ with $m_i \geqslant 2$. From the equality in (4.12.6), we have*

$$\{t_1, \dots, t_g\} = \{3, \dots, n\} \cup \{1, m_1, \dots, m_{n-3}\}.$$

Since $g \leqslant n-1$, the inclusion–exclusion principle implies that

$$\big|\{3, \dots, n\} \cap \{1, m_1, \dots, m_{n-3}\}\big| \geqslant n-3.$$

Since $1 \notin \{3, \dots, n\}$ it follows that $m_1, \dots, m_{n-3} \in \{3, \dots, n\}$, that is, that $m_i > 2$ for all i. This establishes Claim 2.

Claim 2 says that X_2 does not appear to any power in the list of generators of Q_1. Given the form and number of the generators of Q_1, it follows that there is another variable, say X_p with $p \geqslant 3$, that has no power occurring in this list. By assumption, the set $\{v_1^{\mu_1}, v_{m_1}^{\mu_{m_1}}, \dots, v_{m_{n-3}}^{\mu_{m_{n-3}}}\}$ is a minimal weighted 2-path vertex cover of K_ω^n. It follows that the path $v_1v_2v_p$ is covered by $v_1^{\mu_1}$, which implies that $\mu_1 \leqslant a$. However, we have $X_1^{\mu_1} \in Q_1 + Q_* \subset P_1 + P_*$; as in the proof of Claim 1, this implies that $\mu_1 \geqslant b > a \geqslant \mu_1$, contradiction. We conclude that the supposition (4.12.6) is impossible.

From this, we deduce that the only way one can have $Q_1 + Q_* \subseteq P_1 + P_*$ is with

$$Q_1 + Q_* = \big(X_1^{\delta_1}, X_2^{\delta_2}, \dots, X_n^{\delta_n}\big)S$$

for some δ_i. It follows that at least one ideal of this form is irredundant in the intersection (4.12.5), as desired. □

We end with a question motivated by the results of this section.

Questions 4.13. Is there a similar characterization of the Cohen–Macaulayness of $I_{r,\max}(K^n_\omega)$ when $r \geqslant 3$? For instance, must the ideal $I_{r,\max}(K^n_\omega)$ be Cohen–Macaulay if and only if every induced weighted sub-$(r+1)$-clique $K^{r+1}_{\omega'}$ of K^n_ω has $I_{r,\max}(K^{r+1}_{\omega'})$ Cohen–Macaulay?

Acknowledgements

We are grateful to Susan Morey for helpful conversations about this material.

References

[1] D. Campos, R. Gunderson, S. Morey, C. Paulsen, T. Polstra, Depths and Cohen–Macaulay properties of path ideals, J. Pure Appl. Algebra 218 (8) (2014) 1537–1543, MR 3175038.
[2] A. Conca, E. De Negri, M-sequences, graph ideals, and ladder ideals of linear type, J. Algebra 211 (2) (1999) 599–624, MR 1666661 (2000d:13020).
[3] C.A. Francisco, H.T. Hà, Whiskers and sequentially Cohen–Macaulay graphs, J. Comb. Theory, Ser. A 115 (2) (2008) 304–316, MR 2382518 (2008j:13050).
[4] C.A. Francisco, H.T. Hà, A. Van Tuyl, Colorings of hypergraphs, perfect graphs, and associated primes of powers of monomial ideals, J. Algebra 331 (2011) 224–242, MR 2774655 (2012b:13054).
[5] C.A. Francisco, A. Hoefel, A. Van Tuyl, Edge Ideals: a package for (hyper)graphs, J. Softw. Algebra Geom. 1 (2009) 1–4, MR 2878668 (2012m:05243).
[6] C.A. Francisco, A. Van Tuyl, Sequentially Cohen–Macaulay edge ideals, Proc. Am. Math. Soc. 135 (8) (2007) 2327–2337 (electronic), MR 2302553 (2008a:13030).
[7] D. Ingebretson, S. Sather-Wagstaff, Decompositions of monomial ideals in real semigroup rings, Commun. Algebra 41 (11) (2013) 4363–4377, MR 3169524.
[8] J. Martínez-Bernal, S. Morey, R.H. Villarreal, Associated primes of powers of edge ideals, Collect. Math. 63 (3) (2012) 361–374, MR 2957976.
[9] J. Martínez-Bernal, C. Rentería-Márquez, R.H. Villarreal, Combinatorics of symbolic Rees algebras of edge ideals of clutters, in: Commutative Algebra and Its Connections to Geometry, in: Contemp. Math., vol. 555, Amer. Math. Soc., Providence, RI, 2011, pp. 151–164, MR 2882681.
[10] S. Morey, R.H. Villarreal, Edge Ideals: Algebraic and Combinatorial Properties, Prog. Commut. Algebra, vol. 1, de Gruyter, Berlin, 2012, pp. 85–126, MR 2932582.
[11] C. Paulsen, S. Sather-Wagstaff, Edge ideals of weighted graphs, J. Algebra Appl. 12 (5) (2013), 24 pp., MR 3055580.
[12] M. Rogers, S. Sather-Wagstaff, Monomial ideals and their decompositions, draft, http://www.ndsu.edu/pubweb/~ssatherw/DOCS/monomial.pdf, 2011.
[13] A. Van Tuyl, A beginner's guide to edge and cover ideals, in: Monomial Ideals, Computations and Applications, in: Lecture Notes in Math., vol. 2083, Springer, Heidelberg, 2013, pp. 63–94, MR 3184120.
[14] A. Van Tuyl, Edge ideals using Macaulay2, in: Monomial Ideals, Computations and Applications, in: Lecture Notes in Math., vol. 2083, Springer, Heidelberg, 2013, pp. 95–105, MR 3184121.
[15] K.N. Vander Meulen, A. Van Tuyl, C. Watt, Cohen–Macaulay circulant graphs, Commun. Algebra 42 (5) (2014) 1896–1910, MR 3169680.
[16] R.H. Villarreal, Cohen–Macaulay graphs, Manuscr. Math. 66 (3) (1990) 277–293, MR 1031197 (91b:13031).
[17] R.H. Villarreal, Monomial algebras, in: Monographs and Textbooks in Pure and Applied Mathematics, vol. 238, Marcel Dekker Inc., New York, 2001, MR 1800904 (2002c:13001).