AN EXPLORATION OF ALGEBRAIC APPROACHES TO GRAPH THEORY

MAYA URBSCHAT

Abstract. In a 2010 paper, De Loera *et al.* explore the use of polynomial ideals to determine properties of graphs. Here we present an exposition of the results in said paper on k-colorability and Hamiltonicity of graphs, as well as the improvements on these results made by Li *et al.* in 2015. We provide detail and background necessary for an undergraduate reader. Additionally, we provide an explicit formula for the Nullstellensatz certificate of non-2-colorability of a graph, and find the graph with smallest order that has a linear Nullstellensatz certificate of non-3-colorability but does not contain an odd wheel.

Date: April 21, 2016.

An Exploration of Algebraic Approaches to Graph Theory

Maya Urbschat

A thesis presented for the degree of

Bachelor of Arts

Department of Mathematics

Mount Holyoke College

April 21, 2016

Acknowledgements

My thanks to Professor Jessica Sidman for a rigorous and exciting year of advising, to my committee for the generous contribution of their time and expertise, and to the Department for the many opportunities it has offered me.

Contents

List of Figures

1. Introduction

Graph theory is a valuable tool for solving a range of real-world problems. Social networks, class schedules of students, genetic links between species, and linguistic patterns can all be described with graphs. We will use the following notation for graphs.

Definition 1.1. An **undirected graph** $G = \{V(G), E(G)\}$ consists of a set of vertices $V(G)$ and a set of two-element subsets of $V(G)$, the edge set $E(G)$.

A **directed graph** $G = \{V(G), A(G)\}$ consists of a set of vertices $V(G)$ and a set of ordered two-element subsets of $V(G)$, the set of arcs (or directed edges) $A(G)$.

In either case, we say that G has **order** $|V(G)|$.

When there is no confusion about which graph we are discussing, we will abbreviate $V(G)$ and $E(G)$ as V and E, respectively.

Example 1.2. Figure 1 displays a graph G where $V(G) = \{v_1, v_2, v_3, v_4, v_5\}$ and $E(G) = \{\{v_1, v_2\}, \{v_1, v_3\}, \{v_1, v_4\}, \{v_2, v_3\}, \{v_3, v_4\}, \{v_3, v_5\}, \{v_4, v_5\}\}$.

Definition 1.3. A graph G is **k-colorable** if we can assign each vertex in G a (not necessarily unique) label from the set $\{1, \ldots, k\}$ such that no pair of adjacent vertices have the same label.

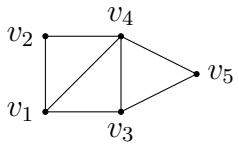

FIGURE 1. A graph of order five.

The **chromatic number** of G is the smallest number k such that G is k-colorable.

Example 1.4. The graph in Figure 1 is 3-colorable. For example, we can let v_1 and v_5 have color 1, v_2 and v_3 have color 2, and v_4 have color 3. This graph is not 2-colorable, since v_1, v_2, and v_4 must all be different colors.

Graph coloring problems arise in many settings. For example, to coordinate the class schedules of a group of students, we could assign a vertex of a graph to each class, and include an edge between every pair of vertices where there is a student taking both classes. Then the chromatic number of the graph is the number of class periods necessary to accommodate every student. While stating a problem in this way is simple and intuitive, determining an arbitrary graph's chromatic number remains a hard question.

We can use systems of polynomials to characterize the properties of many graphs. It is quite computationally intensive to determine whether a given graph satisfies the algebraic characterization of certain properties, but algorithms such as the Nullstellensatz Linear Algebra algorithm (NulLA) of De Loera *et al.* [5] provide relatively efficient tools for such problems.

The following is an exposition of the results of a paper of De Loera *et al.* [4] and an improvement on one such result by Li *et al.* [9]. Section 2 offers a brief introduction of the algebraic concepts that will come up, in subsections on Elementary Algebra, Complex Roots of Unity, Monomials, and Polynomial Rings and Ideals. In the last subsection of this section, we present a proof of the Weak Nullstellensatz, which we will later use in Section 3.

With this background in hand, we proceed to a discussion of the content of [4]. In Section 3 we explore how Section 2 of [4] relates the 3-colorability of a graph to a set of polynomials, and show how the feasibility of the system–and thus the 3-colorability of the graph–can be determined using the Nullstellensatz. We then present a combinatorial characterization of some non-3-colorable graphs as laid out in [9]. Among these are graphs with odd wheels as subgraphs, and we include examples of these, as well as of graphs that satisfy the characterization but do not contain an odd wheel. The smallest of these is shown in Figure 10, and we prove that there are no smaller such graphs.

Section 4 is original work considering how the methods in Section 2 of [4] can be applied to the simpler case of 2-coloring graphs. We find that all non-2-colorable graphs have a Nullstellensatz certificate of degree one, and state the aforementioned certificate explicitly.

In Section 5, we discuss Section 3 of [4] and its algebraic methods of determining whether a graph is uniquely Hamiltonian. We go through the proof, and consider the merits of their encodings and definitions,

noting subtle requirements of the central polynomial system that the authors left for the reader.

2. Algebra

Our work centers on algebraic solutions to various problems in graph theory, and these solutions all involve systems of polynomials with coefficients in a field. We will now define the algebraic concepts necessary to our later arguments.

2.1. **Elementary Algebra.** In this section, we define some basic objects in Abstract Algebra, and give a few of their properties. These objects will be used to prove Theorem 2.36, and throughout Sections 3, 4, and 5.

Definition 2.1. A **ring** R is a non-empty set with two binary operations (here we use addition and multiplication) where for all $a, b, c \in R$,

(1) Addition is commutative; $a + b = b + a$.
(2) Addition is associative; $(a + b) + c = a + (b + c)$.
(3) There exists an additive identity $0 \in R$ such that $a + 0 = 0 + a = 0$.
(4) Every element has an additive inverse; there exists $-a \in R$ such that $a + (-a) = (-a) + a = 0$.
(5) Multiplication is associative; $a(bc) = (ab)c$.
(6) The distributive property holds; $a(b + c) = ab + ac$ and $(a + b)c = ac + bc$.

A **commutative ring** R is a ring in which multiplication is commutative, i.e. where $ab = ba$ for all $a, b \in R$

If a ring R has **unity**, then there exists a multiplicative identity $1 \in R$ where $1 \cdot a = a \cdot 1 = a$ for all a.

Definition 2.2. A **field** is a commutative ring with unity in which every nonzero element has a multiplicative inverse. [6]

Example 2.3. The rational numbers $\mathbb{Q}$ are an infinite field.

The arguments below often center on finite fields, such as $\mathbb{F}_2 = \{0, 1\}$ and $\mathbb{F}_3 = \{0, 1, 2\}$. The binary operations in $\mathbb{F}_2$ and $\mathbb{F}_3$ are addition and multiplication modulo 2 and 3, respectively. For instance, $1+1 = 0$ and $1 \cdot 1 = 1$ in $\mathbb{F}_2$, and $2 + 2 = 1$ and $2 \cdot 2 = 1$ in $\mathbb{F}_3$.

Definition 2.4. The **characteristic** of a field is the smallest positive number of times 1 can be added to itself in which the result is zero. If there is no such number, we say the field has characteristic zero.

Example 2.5. Here are the characteristics of some familiar fields.

(1) $\mathbb{F}_2$ has characteristic 2.

(2) $\mathbb{F}_3$ has characteristic 3.

(3) $\mathbb{Q}$ has characteristic 0.

Definition 2.6. A field is **algebraically closed** if every nonconstant single-variable polynomial with coefficients in the field has a root in the field.

Lemma 2.7. *Every algebraically closed field is infinite.*

Proof. We will prove this by contradiction. Suppose $\mathbb{K}$ is an algebraically closed field, but is finite, containing n elements $a_1, ..., a_n$. Then

$$f(x) = (\prod_{i=1}^{n}(x - a_i)) + 1$$

gives $f(a_i) = 1$ for all $1 \leq i \leq n$. Then f is a single-variable polynomial with coefficients in $\mathbb{K}$, but f has no roots in $\mathbb{K}$, and so $\mathbb{K}$ cannot be algebraically closed.

Thus, if $\mathbb{K}$ is algebraically closed, $\mathbb{K}$ has an infinite number of elements. □

2.2. **Roots of unity.** Many of the polynomial systems used in [4] involve numbers known as roots of unity. This section provides background on these numbers, which are useful in constructing polynomial systems because a system involving roots of unity usually restricts the form of solutions to be roots of unity as well. This gives us a finite, and often small, set of possible solutions to check. Systems using roots of unity appear in Propositions 3.1 and 5.3.

Definition 2.8. Let k be a positive integer. A **k-th root of unity** is a number ω in a field $\mathbb{K}$ such that $\omega^k = 1$.

We say that ω is a **primitive k-th root of unity** if ω is a k-th root of unity but is not a j-th root of unity for any $j < k$.

Example 2.9. Here are some examples of roots of unity.

(1) 1 is a k-th root of unity for any positive integer k.

(2) 1 and -1 are the square roots of unity ($k = 2$).

(3) 1, -1, i, and -i are the fourth roots of unity.

If ω is a primitive k-th root of unity, we can express any k-th root of unity in the form ω^j where $0 \leq j < k$. Each of the distinct k-th roots of unity have different values of j when written in this manner, which supports the following result: In an algebraically closed field whose characteristic does not divide k, there are exactly k k-th roots of unity [8]. These form a cyclic group generated by a primitive k-th root of unity, so we know that a primitive k-th root of unity exists in any algebraically closed field whose characteristic does not divide k.

Theorem 2.10. *For $k > 1$, the sum of all k-th roots of unity is zero.*

Proof. Let $S = 1 + \omega + \omega^2 + \cdots + \omega^{k-1}$ be the sum of the k-th roots of unity, where ω is a primitive k-th root of unity.

Then $\omega S = \omega(1 + \omega + \omega^2 + \cdots + \omega^{k-1}) = \omega + \omega^2 + \cdots + \omega^{k-1} + \omega^k$ However, since $\omega^k = \omega^0 = 1$,

$$\omega + \omega^2 + \cdots + \omega^{k-1} + \omega^k = 1 + \omega + \omega^2 + \cdots + \omega^{k-1}.$$

In other words, $\omega S = S$. Since $k > 1$ and ω is a primitive k-th root of unity, $\omega \neq 1$. Thus, if $\omega S = S$, $S = 0$. □

2.3. **Monomials.** In this section we define monomials, and some monomial orderings. This will prepare us for the discussion of polynomials in the following subsection.

Definition 2.11. A **monomial** is a product of some number of (not necessarily distinct) variables.

The **degree** of a monomial is the number of factors in the product.

Example 2.12. The monomial $x_1^2 x_2 x_3^4$ has degree 7.

Polynomials are linear combinations of monomials, where each monomial is multiplied by some coefficient from the field $\mathbb{K}$.

We often need to choose a way to order a set of monomials. In the arguments below we use two monomial orderings, lexicographic (lex) order, and graded lexicographic (grlex) order.

Definition 2.13. Let $a = (a_1, ..., a_n)$ and $b = (b_1, ..., b_n) \in \mathbb{Z}^n_{\geq 0}$. We say a is greater than b in lex order, or equivalently $a >_{lex} b$ if, in the vector difference $a - b \in \mathbb{Z}^n_{\geq 0}$, the leftmost nonzero entry is positive. We will write $x^a >_{lex} x^b$ if $a >_{lex} b$, where x^a is shorthand for $x_1^{a_1} x_2^{a_2} \cdots x_n^{a_n}$.

Definition 2.14. Let x^a and x^b be monomials. We say x^a is greater than x^b in grlex order, or equivalently $x^a >_{grlex} x^b$, if x^a has a higher degree than x^b, or if the two monomials have equal degree and $x^a >_{lex} x^b$.

Example 2.15. In grlex order, $x_1 x_2^4 > x_1^3 > x_1^2 x_2 > x_1^2 > x_2^2$.

Given a monomial ordering, we can define some key parts of a polynomial, which we will use to prove Proposition 2.39 and Lemma 5.14.

Definition 2.16. Fix a monomial order, and let f be a nonzero polynomial.

- The **leading monomial** of f, $LM(f)$, is the largest monomial in f under the fixed monomial ordering.
- The **leading coefficient** of f, $LC(f)$, is the coefficient that appears with $LM(f)$ in f.

- The **leading term** of f, $LT(f)$, is the product $LC(f) \cdot LM(f)$.

Example 2.17. Let $f = 3x_1x_2^4 + 2x_1^2 + 5x_2$.

In grlex order, $LM(f) = x_1x_2^4$, $LC(f) = 3$, and $LT(f) = 3x_1x_2^4$. In lex order, $LM(f) = x_1^2$, $LC(f) = 2$, and $LT(f) = 2x_1^2$.

2.4. **Polynomial Rings and Ideals.** In this section, we discuss polynomial ideals and varieties, building understanding toward our use of these objects in proving Theorem 2.36, and throughout the discussions in Section 3 and beyond.

Definition 2.18. Let $\mathbb{K}$ be a field. We denote by $\mathbb{K}[x_1, \ldots, x_n]$ the ring of polynomials in n variables with coefficients in $\mathbb{K}$.

Example 2.19. The ring of polynomials in two variables with coefficient 0 or 1 is written $\mathbb{F}_2[x, y]$. The nonzero elements of this ring all have the form $\sum x^p y^q$ where the exponents on each x or y are nonnegative (possibly zero) integers.

The set of monomials in n variables is a subset of $\mathbb{K}[x_1, \ldots, x_n]$.

Definition 2.20. Let P be a system of polynomials in $\mathbb{K}[x_1, \ldots, x_n]$. The **variety** of P, abbreviated $\mathcal{V}(P)$, is the set of solutions in $\mathbb{K}$ for which each polynomial in the system is equal to zero. Elements of the variety have the form $(a_1, a_2, \ldots, a_n)$ for $a_i \in \mathbb{K}$.

Example 2.21. Varieties in any field $\mathbb{K}$ can have a finite number of elements, an infinite number of elements, or be empty.

- If $P = \{x_1 + x_2, 2x_1\}$, then $\mathcal{V}(P) = \{(0, 0)\}$.

- If $Q = \{x_1 + x_2, 2x_1 + 2x_2\}$, then $\mathcal{V}(Q) = \{(x, -x) | x \in \mathbb{K}\}$.
- If $S = \{x_1 + x_2, 1\}$, then $\mathcal{V}(S) = \emptyset$.

Definition 2.22. An **ideal** I of a commutative ring R is a subset of the ring with the following properties:

- $0 \in I$.
- if $f \in I$ and $g \in I$, then $f + g \in I$.
- if $f \in I$ and $r \in R$, then $fr \in I$.

Example 2.23. The even integers $2\mathbb{Z}$ are an ideal of the integers $\mathbb{Z}$:

- $0 \in 2\mathbb{Z}$
- If $f, g \in 2\mathbb{Z}$ then $f = 2p$ and $g = 2q$ for some integers p and q. Then $f + g = 2p + 2q = 2(p + q) \in 2\mathbb{Z}$.
- If $f \in 2\mathbb{Z}$ and $r \in \mathbb{Z}$, then $f = 2p$ for some integer p, and $fr = 2pr = 2(pr) \in 2\mathbb{Z}$.

Example 2.24. The zero ideal $\{0\}$ is an ideal of any ring R.

- $0 \in \{0\}$
- There is only one element in $\{0\}$, and $0 + 0 = 0 \in \{0\}$.
- For any $r \in R$, by definition $r \cdot 0 = 0 \in \{0\}$.

Lemma 2.25. *Let I be an ideal of a commutative ring R (with unity). We have the unity element $1 \in I$ if and only if $I = R$.*

Proof. If $1 \in I$ and $g \in R$, then since ideals absorb products in the ring, $1 \cdot g = g \in I$. Therefore, $R \subseteq I$, and by definition of an ideal $I \subseteq R$. Therefore, $I = R$.

If $I = R$, it is clear that as $1 \in R$, $1 \in I$ as well. □

Lemma 2.26. *Let $\mathbb{K}$ be a field. The only ideals of $\mathbb{K}$ are $\{0\}$ and $\mathbb{K}$*

Proof. Let I be an ideal of $\mathbb{K}$, and suppose we have $a \neq 0 \in I$. As $\mathbb{K}$ is a field, there exists $a^{-1} \in \mathbb{K}$, so by product absorbtion, $a^{-1}a = 1 \in I$, and therefore $I = \mathbb{K}$. Thus if an ideal of $\mathbb{K}$ has nonzero elements, it is $\mathbb{K}$ itself, so by Example 2.24, $\{0\}$ and $\mathbb{K}$ are the only ideals of $\mathbb{K}$. $\square$

Definition 2.27. We say that an ideal of $\mathbb{K}[x_1, \ldots, x_n]$ is **generated** by a set of polynomials $f_1, \ldots, f_s$ in $\mathbb{K}[x_1, \ldots, x_n]$ if every element in I can be written as $h_1f_1 + h_2f_2 + \cdots + h_sf_s$ for some $h_1, \ldots, h_s$ in $\mathbb{K}[x_1, \ldots, x_n]$.

We can then write $I = \langle f_1, \ldots, f_s \rangle$, and we say that $f_1, \ldots, f_s$ is a **basis** of I.

Example 2.28. We would like to be sure that $I = \langle f_1, \ldots, f_s \rangle$ is an ideal.

Since $0 \in \mathbb{K}[x_1, \ldots, x_n]$, $0 \cdot f_1 = 0 \in I$.

Suppose $g_1 = h_1f_1 + h_2f_2 + \cdots + h_sf_s$ and $g_2 = m_1f_1 + m_2f_2 + \cdots + m_sf_s$ are in I, where $h_i, m_i \in \mathbb{K}[x_1, \ldots, x_n]$. Then

$$\begin{aligned} g_1 + g_2 &= (h_1f_1 + m_1f_1) + \cdots + (h_sf_s + m_sf_s) \\ &= (h_1 + m_1)f_1 + \cdots + (h_s + m_s)f_s. \end{aligned}$$

Since the ring $\mathbb{K}[x_1, \ldots, x_n]$ is closed under addition, each sum $h_i + m_i \in \mathbb{K}[x_1, \ldots, x_n]$, so $g_1 + g_2 \in I$.

Suppose $g = h_1f_1 + h_2f_2 + \cdots + h_sf_s$ is in I and $r \in \mathbb{K}[x_1, \ldots, x_n]$. Then $gr = rh_1f_1 + rh_2f_2 + \cdots + rh_sf_s$. Since the ring $\mathbb{K}[x_1, \ldots, x_n]$

is closed under multiplication, each product $rh_i \in \mathbb{K}[x_1, \ldots, x_n]$, so $gr \in I$.

Hence, $I = \langle f_1, \ldots, f_s \rangle$ is an ideal.

The generating set of an ideal in $\mathbb{K}[x_1, \ldots, x_n]$ is not unique, as we can add elements of the ideal to an existing generating set to make a larger generating set. Further, an ideal does not have a unique minimal generating set. This can cause problems when we try to determine if a given polynomial is in an ideal.

One way to determine ideal membership is that if a polynomial f has remainder zero upon division by a generating set of I, then f is in I. However, the inverse of this statement, "if f does not have remainder zero upon division by a given generating set of I, then f is not in I," does not hold. It is possible for f to have a nonzero remainder on division by some generating set of I, but still be a member of I. This is because f might have different remainders on division by different generating sets of the ideal.

The following definition is therefore quite useful, as the generating sets in this category all produce the same remainder from a given polynomial.

Definition 2.29. Fix a monomial order. A finite subset $G = \{g_1, \ldots, g_t\}$ of an ideal I is a **Gröbner basis** of I if

$$\langle LT(I) \rangle = \langle LT(g_1), \ldots, LT(g_t) \rangle$$

where $\langle LT(I) \rangle$ is the ideal generated by the leading terms of all polynomials in I. [2]

Example 2.30. Let $I = \langle x^3 - 2xy, x^2y - 2y^2 + x \rangle$, and use grlex order.

$\{x^3 - 2xy, x^2y - 2y^2 + x\}$ is a basis of I, but not a Gröbner basis of I because $x^2 = -y(x^3 - 2xy) + x(x^2y - 2y^2 + x)$ is in I, but the leading terms of these functions are x^3 and x^2y, and $x^2 \notin \langle x^3, x^2y \rangle$.

However, $\{x^3 - 2xy, x^2y - 2y^2 + x, x^2, xy, 2y^2 + x\}$ is a Gröbner basis of I. (This may not be intuitively clear, but we will discuss a means of checking that it is true below.)

Definition 2.31. A Gröbner basis $G = \{g_1, \ldots, g_t\}$ of a polynomial ideal I is a **reduced Gröbner basis** of I if the following hold for all $1 \leq i \leq t$

(1) $LC(g_i) = 1$

(2) No monomial of g_i lies in $\langle LT(G \backslash \{g_i\}) \rangle$.

Example 2.32. Let $I = \langle x^2, y^2 + x \rangle$, and use grlex order. The set $G = \{x^2, y^2 + x\}$ is a reduced Gröbner basis of I.

Unlike Gröbner bases in general, reduced Gröbner bases are unique.

Definition 2.33. Let f and g be nonzero polynomials in $\mathbb{K}[x_1, \ldots, x_n]$. The **s-polynomial** of f and g is

$$s(f, g) = \frac{x^\gamma}{LT(f)} f - \frac{x^\gamma}{LT(g)} g$$

where x^γ is the least common multiple of $LM(f)$ and $LM(g)$.

Example 2.34. We see that

$$s(x^2, y^2 + x) = \frac{x^2y^2}{x^2} x^2 - \frac{x^2y^2}{y^2}(y^2 + x) = x^2y^2 - x^2y^2 - x^3 = x^3.$$

Theorem 2.35. *(Buchberger's Criterion) A generating set* $G = \{g_1, \ldots, g_t\}$ *of an ideal* I *is a Gröbner basis of* I *if and only if the s-polynomial* $s(g_i, g_j)$ *for all pairs* $i \neq j$ *has remainder zero on division by the polynomials in* G*.*

A proof of this theorem can be found in 2.6, Theorem 6 of Cox *et al.* [2].

2.5. **The Nullstellensatz.** We now have the background necessary to prove an important theorem in algebraic geometry, which we will apply in Section 3.

Theorem 2.36. *(Weak Nullstellensatz)* [2] *Let* $\mathbb{K}$ *be an algebraically closed field and let* $I \subset \mathbb{K}[x_1, \ldots, x_n]$ *be an ideal.* I *will satisfy* $\mathcal{V}(I) = \emptyset$ *if and only if* $I = \mathbb{K}[x_1, \ldots, x_n]$*.*

We note that by Lemma 2.25, $1 \in I$ if and only if $I = \mathbb{K}[x_1, \ldots, x_n]$.

First assume that $I = \mathbb{K}[x_1, \ldots, x_n]$. Then $1 \in I$, and it follows that, since $1 \neq 0$, $\mathcal{V}(I) = \emptyset$.

We will now prove the other part of the if and only if statement: that $I \subsetneq \mathbb{K}[x_1, \ldots, x_n]$ implies that $\mathcal{V}(I) \neq \emptyset$. This proof follows the one provided in [2], with added detail and all exercises worked out.

The body of this proof centers on the following operation on polynomials in the ideal I.

Definition 2.37. Given a number $a \in \mathbb{K}$ and a polynomial $f \in \mathbb{K}[x_1, \ldots, x_n]$, define $\overline{f} = f(x_1, \ldots, x_{n-1}, a)$. In other words, $\overline{f}$ is f evaluated at $x_n = a$, and doing this gives us a polynomial in the subring $\mathbb{K}[x_1, \ldots, x_{n-1}]$.

Lemma 2.38. *The set*

$$I_{x_n=a} = \{\overline{f} | f \in I\}$$

is an ideal of $\mathbb{K}[x_1, \ldots, x_{n-1}]$.

Proof. Let $z = 0$ be the zero polynomial. As I is an ideal, $z \in I$. Evaluating z at $x_n = a$ still gives 0, so $\overline{z} = 0 \in I_{x_n=a}$.

Let $f, g \in I$. Then $f + g \in I$, and $\overline{f}, \overline{g}, \overline{f+g} \in I_{x_n=a}$. We can see that $f(x_1, \ldots, x_{n-1}, x_n) + g(x_1, \ldots, x_{n-1}, x_n) = (f+g)(x_1, \ldots, x_{n-1}, x_n)$ by the usual polynomial properties. From this we see that $f(x_1, \ldots, x_{n-1}, a) + g(x_1, \ldots, x_{n-1}, a) = (f + g)(x_1, \ldots, x_{n-1}, a)$, or equivalently $\overline{f} + \overline{g} = \overline{f+g}$. We know $\overline{f+g} \in I_{x_n=a}$, so $\overline{f} + \overline{g} \in I_{x_n=a}$ and we have closure.

Let $f \in I$, $h \in \mathbb{K}[x_1, \ldots, x_n]$. Then $hf \in I$, so $\overline{f}, \overline{hf} \in I_{x_n=a}$. We know that $\overline{hf} = hf(x_1, \ldots, x_{n-1}, a) = h(x_1, \ldots, x_{n-1}, a) f(x_1, \ldots, x_{n-1}, a)$ $= h\overline{f}$ since h is not a function of x_n. Thus $h\overline{f} \in I_{x_n=a}$ and $I_{x_n=a}$ absorbs products. Therefore, $I_{x_n=a}$ is an ideal of $\mathbb{K}[x_1, \ldots, x_{n-1}]$. □

With this ideal in hand, we present the following claim:

Proposition 2.39. *Given that* $\mathbb{K}$ *is algebraically closed and that* $I \subsetneq \mathbb{K}[x_1, \ldots, x_n]$, *there exists* $a \in \mathbb{K}$ *such that* $I_{x_n=a} \subsetneq \mathbb{K}[x_1, \ldots, x_{n-1}]$.

Now we will prove the claim using two cases.

Proof. We first look at the case where $I \cap \mathbb{K}[x_n] \neq \{0\}$; in other words, where I contains single-variable polynomials in x_n. Let $f \in I \cap \mathbb{K}[x_n]$ be a nonzero polynomial of degree r. We also know that f is not

constant, because if it were, it would have an inverse in $\mathbb{K}$ such that $f \cdot f^{-1} = 1 \in I \cap \mathbb{K}[x_n] \subseteq I$, and we know that as $I \subsetneq \mathbb{K}[x_1, \ldots, x_n]$, $1 \notin I$.

Since $\mathbb{K}$ is algebraically closed, we can write $f = c\prod_{i=1}^{r}(x_n - b_i)^{m_i}$, where $c, b_i \in \mathbb{K}$ and $m_i \in \mathbb{N}$ for $1 \leq i \leq r$, and taking $c \neq 0$.

Suppose that $I_{x_n=b_i} = \mathbb{K}[x_1, \ldots, x_{n-1}]$ for $1 \leq i \leq r$. Then $1 \in I_{x_n=b_i}$ for all i, so for every i there exists some $B_i \in I$ such that $B_i(x_1, \ldots, x_{n-1}, b_i) = 1$. Then

$$1 = B_i(x_1, \ldots, x_{n-1}, b_i) = B_i(x_1, \ldots, x_{n-1}, x_n - (x_n - b_i))$$

by a simple substitution, which we can rewrite as

$$1 = B_i(x_1, \ldots, x_n) + A_i(x_n - b_i)$$

for some $A_i \in \mathbb{K}[x_1, \ldots, x_n]$.

To see that this is true, consider an arbitrary monomial in $B_i(x_1, \ldots, x_{n-1}, b_i)$. If it contains no factors of $b_i = x_n - (x_n - b_i)$, then it is already a monomial in B_i. If instead the monomial contains a factor $b_i^j = (x_n - (x_n - b_i))^j$, then using the binomial theorem we can expand the monomial into factors of $x_1, \ldots, x_{n-1}$ multiplied by $(x_n^j +$ (terms divisible by $(x_n - b_i)))$. Then the term containing x_n^j above is a monomial in B_i, and the latter terms, when divided by $(x_n - b_i)$, are monomials in A_i. Applying this method to all monomials in B_i, we will be left with the correct A_i.

We can then multiply the equations $1 = B_i + A_i(x_n - b_i)$ over $1 \leq i \leq r$, repeating the i-th factor m_i times, to obtain the following:

$$1 = \prod_{i=1}^{r}(A_i(x_n - b_i) + B_i)^{m_i} = A\prod_{i=1}^{r}(x_n - b_i)^{m_i} + B$$

with $A = \prod_{i=1}^{r} A_i^{m_i} \in \mathbb{K}[x_1, \ldots, x_n]$ and B being the sum of all terms but the first in the expansion of $\prod_{i=1}^{r}(A_i(x_n - b_i) + B_i)^{m_i}$. We know that $B \in I$ because each term in B is a multiple of some $B_i \in I$.

We see that $\prod_{i=1}^{r}(x_n - b_i)^{m_i} = c^{-1}f$ from our definition of f above, and since $c, c^{-1} \in \mathbb{K}[x_1, \ldots, x_n]$, we have $c^{-1}f \in I$. Then by properties of ideals, $Ac^{-1}f + B = 1 \in I$.

This contradicts $I \subsetneq \mathbb{K}[x_1, \ldots, x_n]$. Therefore, it must be true that $I_{x_n=b_i} \neq \mathbb{K}[x_1, \ldots, x_{n-1}]$ for some i. That b_i is the desired a for the claim.

In the second case, $I \cap \mathbb{K}[x_n] = \{0\}$, in other words, there are no single-variable polynomials in x_n contained in I.

Let $\{g_1, \ldots, g_t\}$ be a Gröbner basis of I for lex order with $x_1 > \cdots > x_n$, and write

$$g_i = c_i(x_n)x^{\alpha_i} + (\text{terms less than } x^{\alpha_i})$$

where $c_i(x_n) \in \mathbb{K}[x_n]$ is nonzero (but possibly constant), and x^{α_i} is a monomial in $x_1, \ldots, x_{n-1}$. This means that $LM(c_i(x_n)x^{\alpha_i})$ is the leading monomial of g_i.

By Lemma 2.7, $\mathbb{K}$ is infinite. Since we have a finite number t of polynomials $c_i(x_n)$, each having a finite number of roots, we know we

can find some $a \in \mathbb{K}$ that is not a root of any $c_i(x_n)$. This means $c_i(a) \neq 0$ for all i.

We will now show that the polynomials $\overline{g_i} = g_i(x_1, \ldots, x_{n-1}, a)$, $(1 \leq i \leq t)$ form a basis of $I_{x_n=a}$.

Let $f \in I$. Then $\overline{f} \in I_{x_n=a}$. Further, as $\{g_i\}$ is a Gröbner basis of I, $f = \sum_{i=1}^{t} h_i g_i$ for some $h_i \in \mathbb{K}[x_1, \ldots, x_n]$. Then

$$\overline{f} = \sum_{i=1}^{t} \overline{h_i}\,\overline{g_i}$$

by polynomial properties. Thus every $\overline{f}$ in $I_{x_n=a}$ can be written as a combination of $\overline{g_i}$, so $\{\overline{g_i}\}$ is a basis of $I_{x_n=a}$.

Substituting $x_n = a$ in $g_i = c_i(x_n)x^{\alpha_i} + (\text{terms less than } x^{\alpha_i})$, it is clear that $LT(\overline{g_i}) = c_i(a)x^{\alpha_i}$ since $c_i(a) \neq 0$. We also note that $x^{\alpha_i} \neq 1$, otherwise $g_i = c_i(x_n) \in I \cap \mathbb{K}[x_n] = \{0\}$, which is a contradiction because $c_i(x_n)$ is nonzero. This shows that the $LT(\overline{g_i})$ are nonconstant for all $1 \leq i \leq t$.

We will show in Lemma 2.40 that $\{\overline{g_i}\}$ is a Gröbner basis of $I_{x_n=a}$. It follows that $1 \notin I_{x_n=a}$ since no $LT(\overline{g_i})$ can divide 1, as the leading terms are nonconstant. Thus $I_{x_n=a} \neq \mathbb{K}[x_1, \ldots, x_{n-1}]$, which is what we want to show. □

Lemma 2.40. *Let I be an ideal with Gröbner basis $\{g_1, \ldots, g_t\}$. Then $\{\overline{g_i}\}$ is a Gröbner basis of $I_{x_n=a}$.*

Proof. Let us take a pair g_i, g_j in the Gröbner basis of I, and consider

$$S = c_j(x_n)\frac{x^\gamma}{x^{\alpha_i}}g_i - c_i(x_n)\frac{x^\gamma}{x^{\alpha_j}}g_j$$

where $x^\gamma = lcm(x^{\alpha_i}, x^{\alpha_j})$. Note that $S \in I$.

We can see that $x^\gamma > LT(S)$ because the terms with $LT(g_i)$ and $LT(g_j)$ cancel, so $LT(S)$ is the maximum of

$$LT(c_j(x_n))\frac{x^\gamma}{x^{\alpha_i}} \cdot \text{a monomial less than } x^{\alpha_i}$$

and

$$LT(c_i(x_n))\frac{x^\gamma}{x^{\alpha_j}} \cdot \text{a monomial less than } x^{\alpha_j}.$$

It is intuitively clear that dividing x^γ by a monomial and then multiplying it by a smaller monomial will yield a result smaller than x^γ. We also know that x^{α_i} and x^{α_j} are nonconstant, so they have some factor of $x_1, \ldots, x_{n-1}$ to a power $m \geq 1$. Since $LT(c_i(x_n)), LT(c_j(x_n))$ are both functions of only x_n, multiplying this term into the result we had before to form $LT(S)$ will not increase the product in lex order enough to change the inequality. Therefore, $x^\gamma > LT(S)$.

Since $S \in I$, we can write $S = \sum_{l=1}^{t} h_l g_l$ for some $h_l \in \mathbb{K}[x_1, \ldots, x_n]$. Then evaluating at $x_n = a$ gives

$$c_j(a)\frac{x^\gamma}{x^{\alpha_i}}\overline{g_i} - c_i(a)\frac{x^\gamma}{x^{\alpha_j}}\overline{g_j} = \overline{S} = \sum_{l=1}^{t} \overline{h_l}\overline{g_l}.$$

Since $LT(\overline{g_i}) = c_i(a)x^{\alpha_i}$, $\overline{S} = s(\overline{g_i}, \overline{g_j})$, the s-polynomial for $\overline{g_i}, \overline{g_j}$, up to the nonzero constant multiple $c_i(a)c_j(a)$.

We want to find an lcm representation $s(\overline{g_i}, \overline{g_j}) = \sum_{l=1}^{t} \overline{h_l}\overline{g_l}$ where $lcm(LM(\overline{g_i}), LM(\overline{g_j}) > LT(\overline{h_l}\overline{g_l})$ for all $\overline{h_l}\overline{g_l} \neq 0$. If we can express all s-pairs in this manner, by the extension of Buchberger's Criterion in Chapter 2.9 Theorem 6 of [2], $\{\overline{g_i}\}$ is a Gröbner basis of $I_{x_n=a}$.

Consider an arbitrary polynomial $f \in I$. We know that $LT(\overline{f})$ corresponds to some term ϕ in f. Either ϕ contained some factor of x_n, in which case that factor becomes a constant in $\overline{f}$ but nothing else is changed, and $\phi > LT(\overline{f})$ in lex order, or ϕ contained no factor of x_n, and $\phi = LT(\overline{f})$. Then we have $\phi \geq LT(\overline{f})$, and by the definition of the leading term, $LT(f) \geq \phi \geq LT(\overline{f})$.

Then $x^\gamma > LT(S)$ and $LT(S) \geq LT(\overline{S})$ imply

$$x^\gamma > LT(\overline{S}) = LT(\overline{h_l}\overline{g_l})$$

for all $\overline{h_l}\overline{g_l} \neq 0$.

Since $x^\gamma = lcm(x^{\alpha_i}, x^{\alpha_j}) = lcm(LM(\overline{g_i}), LM(\overline{g_j}))$, we have an lcm representation $s(\overline{g_i}, \overline{g_j}) = \sum_{l=1}^{t} \overline{h_l}\overline{g_l}$ where

$$lcm(LM(\overline{g_i}), LM(\overline{g_j})) > LT(\overline{h_l}\overline{g_l})$$

for all $\overline{h_l}\overline{g_l} \neq 0$, as desired. □

Proof. (Second direction of Theorem 2.36, the Weak Nullstellensatz)

By Proposition 2.39, we can use induction to find a set of elements $a_1, \ldots, a_n \in \mathbb{K}$ for which $I_{x_n=a_n,\ldots,x_1=a_1} \subsetneq \mathbb{K}$.

By Lemma 2.26, the only ideals of $\mathbb{K}$ are $\{0\}$ and $\mathbb{K}$, so $I_{x_n=a_n,\ldots,x_1=a_1} = \{0\}$. It is equivalent to say that $(a_1, \ldots, a_n) \in \mathcal{V}(I)$, since evaluating all elements of the ideal at $(a_1, \ldots, a_n)$ makes all elements equal to zero. Therefore, $\mathcal{V}(I) \neq \emptyset$ and the theorem is true. □

3. Graph Coloring

This section looks at graph coloring from an algebraic perspective, specifically focusing on ways to determine if a graph is 3-colorable.

Throughout this section, we will let $\mathbb{K}$ be a field with characteristic relatively prime to k, the number of colors used.

In Section 3.1, we consider a polynomial system that determines if a graph is k-colorable. In Section 3.2, we define a Nullstellensatz certificate of the infeasability of a system. In Section 3.3, we present a combinatorial characterization of the non-3-colorability of a graph. Finally, in Section 3.4, we consider examples of graphs that satisfy this characterization.

3.1. **Bayer's Formulation.**

To begin, we will present a polynomial system related to the chromatic number of a graph [1].

Proposition 3.1. *(Bayer's Formulation) Let $G = (V, E)$ be an undirected simple graph (that is, with at most one edge between any pair of vertices) on vertices $V = \{1, ..., n\}$. Fix a positive integer k and let $\mathbb{K}$ be a field with characteristic relatively prime to k. The polynomial system*

$$J_G = \{x_a^k - 1 = 0, x_a^{k-1} + x_a^{k-2}x_b + ... + x_b^{k-1} = 0 \mid a \in V, \{a, b\} \in E\}$$

has a common zero over the algebraic closure of $\mathbb{K}$ *if and only if the graph* G *is* k*-colorable.*

The following proof was written independently.

Proof. First we will prove the forwards direction. If J_G has a common zero, we have $x_a^k = 1$. Then x_a is a kth root of unity, so we can write $x_a = \omega^{k_a}$ for some $k_a \in \{0, ..., k-1\}$, where ω is a primitive k-th root of unity. Substituting for x_a and x_b in the equation $x_a^{k-1} + x_a^{k-2}x_b + ... + x_b^{k-1} = 0$, we have

$$\sum_{j=1}^{k} \omega^{((k-j)k_a+(j-1)k_b)} = 0$$

Suppose that $x_a = x_b$ for some $\{a, b\} \in E$, in other words that the vertices a and b in V have the same color. Then $k_a = k_b$, since there is only one $k_a \in \{0, ..., k-1\}$ satisfying $x_a = \omega^{k_a}$ for a given kth root of unity x_a. If $k_a = k_b$, then $(k-j)k_a + (j-1)k_b = (k-1)k_a$ for all $j \in \{1, ..., k\}$. This gives us

$$\sum_{j=1}^{k} \omega^{((k-j)k_a+(j-1)k_b)} = \sum_{j=1}^{k} \omega^{(k-1)k_a} = k\omega^{(k-1)k_a} \neq 0.$$

This contradicts the fact that J_G has a common zero, so $x_a \neq x_b$ if $\{a, b\} \in E$, and thus G is k-colorable.

Now we will prove the backwards direction. If G is k-colorable, for any $\{a, b\} \in E$, we can color a and b with different kth roots of unity $x_a = \omega^{k_a}, x_b = \omega^{k_b}$. Then $x_a^k - 1 = 0$ for all $a \in V$. As above, we can

write $x_a^{k-1} + x_a^{k-2}x_b + ... + x_b^{k-1}$ as

$$\sum_{j=1}^{k} \omega^{((k-j)k_a+(j-1)k_b)}.$$

If $\{a, b\} \in E$, then $x_a \neq x_b$, so $k_a \neq k_b$. We will show by contradiction that every term in the sum above is unique. Let $m, n \in \{1, ..., k\}$ such that $m \neq n$, and suppose that

$$(k-m)k_a + (m-1)k_b = (k-n)k_a + (n-1)k_b.$$

We can rewrite this equation as
$((k-m)-(k-n))k_a = ((n-1)-(m-1))k_b$ or $(n-m)k_a = (n-m)k_b$, which gives us $k_a = k_b$. This is a contradiction, as we know $k_a \neq k_b$. Therefore $(k-m)k_a + (m-1)k_b \neq (k-n)k_a + (n-1)k_b$ if $m \neq n$, so each term in the sum
$\sum_{j=1}^{k} \omega^{((k-j)k_a+(j-1)k_b)}$ is a different kth root of unity. There are k terms in the sum, so all kth roots of unity must be included. Since the sum of all kth roots of unity is zero, we have

$$x_a^{k-1} + x_a^{k-2}x_b + ... + x_b^{k-1} = \sum_{j=1}^{k} \omega^{((k-j)k_a+(j-1)k_b)} = 0$$

if $\{a, b\} \in E$. Thus, if G is k-colorable, J_G has a common zero. Hence, J_G has a common zero if and only if G is k-colorable. □

3.2. **Nullstellensatz Certificates of Infeasability.** Now that we have established this correspondence, we would like a method of determining if a system such as J_G has a common zero. To find one, let us begin with a definition.

Definition 3.2. Let $S = \{f_1 = 0, \ldots, f_r = 0\}$ be a system of polynomial equations with coefficients in $\mathbb{K}$. The **Nullstellensatz certificate** of infeasability of S is an equation

$$1 = \sum_{i=1}^{r} \beta_i f_i$$

for some polynomials $\beta_1, \ldots, \beta_r \in \mathbb{K}[x_1, \ldots, x_n]$. [4]

Corollary 3.3. *(Corollary to the Weak Nullstellensatz) A system of polynomials with coefficients in $\mathbb{K}$ has no solution if and only if the system has a Nullstellensatz certificate of infeasability.* [4]

Proof. Let $S = \{f_1 = 0, \ldots, f_r = 0\}$ be a system of polynomial equations with coefficients in $\mathbb{K}$. We will first consider the backwards implication. If S has a Nullstellensatz certificate of infeasability, there is a set of polynomials $\beta_1, \ldots, \beta_r \in \mathbb{K}[x_1, \ldots, x_n]$ such that $1 = \sum_{i=1}^{r} \beta_i f_i$. In other words, $1 \in \langle S \rangle$. Then by Lemma 2.25, we see that $\langle S \rangle = \mathbb{K}[x_1, \ldots, x_n]$. By the Weak Nullstellensatz, the variety $\mathcal{V}(S) = \emptyset$.

Now consider the forwards implication. If S has no solution, then the variety $\mathcal{V}(S) = \emptyset$. By the Weak Nullstellensatz, this means $\langle S \rangle = \mathbb{K}[x_1, \ldots, x_n]$. Then $1 \in \langle S \rangle$, so by definition of the ideal generated by S, there are some polynomials $\beta_1, \ldots, \beta_r \in \mathbb{K}[x_1, \ldots, x_n]$ such that $1 = \sum_{i=1}^{r} \beta_i f_i$. □

Thus, we can show that a graph G is not k-colorable by finding a Nullstellensatz certificate of infeasability for the system J_G.

3.3. **Characterizing Non-3-Colorable Graphs.** The first main section of [4] presents a combinatorial characterization of graphs that have a linear Nullstellensatz certificate of non-3-colorability (that is, all of the polynomials in the certificate have degree less than or equal to 1). While their characterization is sound, and the proof relatively straightforward, it requires covering an undirected graph with directed edges in one of two ways. Building on this work, the authors in [9] give a simpler characterization, which is the one we shall present here.

To understand the characterization in [9], we need the following definitions.

Definition 3.4. A **path** of length k in a graph G is a list of vertices $v_1, v_2, \ldots, v_{k+1}$ where each $v_i \in V(G)$ is distinct and $\{v_i, v_{i+1}\} \in E(G)$ for each $1 \leq i \leq k$. If the above are the only edges in G, we say that G itself is a path.

A list of vertices $v_1, v_2, \ldots, v_{k+1}$ where $\{v_i, v_{i+1}\} \in E(G)$ for each $1 \leq i \leq k$, but the $v_i \in V(G)$ are not required to be distinct, is called a **walk**.

Example 3.5. In Figure 2, v_1, v_4, v_5, v_3 is a path of length three in the graph. On the other hand, v_1, v_4, v_3, v_1 is not a path, because a vertex is repeated.

The graph in Figure 3 is a path of length three.

Note: A path of length zero is a single vertex.

Definition 3.6. A **cycle** of length k in a graph G is a list of vertices $v_1, v_2, \ldots, v_k, v_1$ where each $v_i \in V(G)$ is distinct and $\{v_i, v_{i+1}\} \in E(G)$

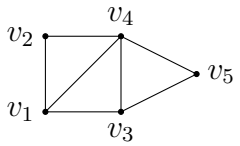

FIGURE 2. The same graph of order five, reprinted for convenience.

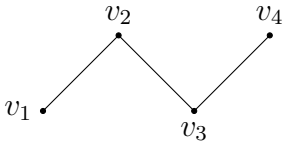

FIGURE 3. A path of length three.

for each $1 \leq i \leq k-1$, and $\{v_k, v_1\} \in E(G)$. If the above are the only edges in G, we say that G itself is a cycle.

Example 3.7. In Figure 2, v_1, v_4, v_5, v_3, v_1 is a cycle of length four in the graph. On the other hand, v_1, v_4, v_3, v_5 is not a cycle, because it is not closed.

The graph in Figure 4 is a cycle of length four.

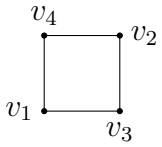

FIGURE 4. A cycle of length four.

Definition 3.8. If two vertices in a graph are connected by an edge, we say that the vertices are **adjacent**. Further, we call each of these vertices **incident** to the edge between them.

It is occasionally expedient to stretch this definition and refer to an edge as incident to a vertex rather than the reverse.

Example 3.9. In Figure 2, the vertices v_1 and v_2 are adjacent, and both are incident to the edge $\{v_1, v_2\}$.

In addition to these elementary definitions, we will use a definition introduced in [9].

Definition 3.10. [9] A graph G with vertex set $V = \{v_1, \ldots, v_n\}$ and edge set E is **covered by length 2 paths** if there exists a set C of length 2 paths in G such that

(1) each edge in E appears in an even number of paths in C,

(2) the number of paths $v_i v_k v_j$ in C in which $k < i, j$ or $k > i, j$ is odd, and

(3) if $v_i, v_j \in V$ but $\{v_i, v_j\} \notin E$, then the number of paths in C with v_i and v_j as endpoints is even.

Note that the term "covered" is misleading, as not all edges nor vertices of a graph need be included in the set of length 2 paths that "cover" it.

Example 3.11. We can cover the complete graph (where all vertices are adjacent) of order four, K_4, with the following set of length 2 paths:

$$C = \{v_2 v_1 v_3, v_3 v_1 v_4, v_2 v_1 v_4\}$$

(see Figure 5). To be certain we have satisfied Definition 3.10, consider:

The edges $\{v_1, v_2\}, \{v_1, v_3\}$ and $\{v_1, v_4\}$ appear twice in C, while the edges $\{v_2, v_3\}, \{v_3, v_4\}$ and $\{v_2, v_4\}$ appear zero times in C. This satisfies condition (1) of Definition 3.10.

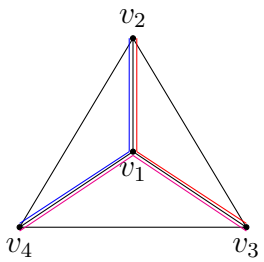

FIGURE 5. K_4 covered by length 2 paths.

All three paths in C are of the form $v_i v_j v_k$ in which $j < i, k$, so there is an odd number of them. This satisfies condition (2) of Definition 3.10.

As all vertices in K_4 are adjacent, we need not consider condition (3), so this covering of K_4 satisfies Definition 3.10.

We will now define a class of graphs that can be covered by length 2 paths.

Definition 3.12. A graph G of order n is a **wheel** if its edge set is as follows:

$$E(G) = \{\{v_1, v_2\}, \{v_2, v_3\}, \ldots, \{v_{n-2}, v_{n-1}\}\{v_{n-1}, v_1\},$$
$$\{v_1, v_n\}, \{v_2, v_n\}, \ldots, \{v_{n-1}, v_n\}\}$$

We say that G is an odd wheel if the length of the outer cycle $v_1, \ldots, v_{n-1}, v_1$ is odd, that is if n is even. If a wheel is not odd, it is even.

Example 3.13. Some small wheels appear in Figure 6. The complete graph K_4 can also be referred to as the wheel W_3. We note that the wheels are indexed by the number of vertices on the outer cycle.

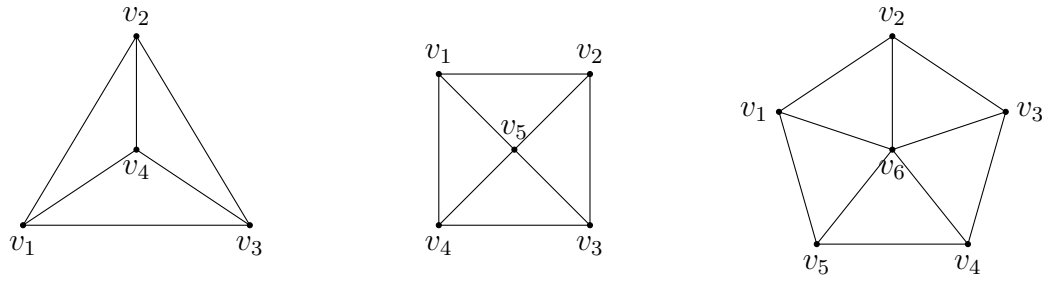

FIGURE 6. $K_4 = W_3$, W_4, and W_5.

This proposition and its proof are new work.

Proposition 3.14. *An odd wheel can be covered by paths of length two in accordance with Definition 3.10 for any ordering of its vertices.*

Proof. Let W be the odd wheel of order $2n$. Let v_{2n} be the center vertex of the wheel, and consider the following set of paths of length two:

$$C = \{v_i, v_{2n}, v_{i+1} | 1 \leq i \leq 2n-2\} \bigcup \{v_{2n-1}, v_{2n}, v_1\}.$$

There are $2n-1$ paths in C. Each edge not incident to v_{2n} is covered by no paths in C, while each edge incident to v_{2n} is covered by two paths in C. Hence, C satisfies condition i) of Definition 3.10. Additionally, all pairs of non-adjacent vertices in W are connected by no paths in C, so C satisfies condition iii) of Definition 3.10.

Now, we define a new ordering by assigning each vertex of W a unique label from the set $\{1, 2, ..., 2n\}$. Let j be the label assigned to v_{2n}. We will consider the set S_j^+ of vertices whose labels are greater than j, and the set S_j^- of vertices whose labels are less than j. Note that all non-center vertices of W are in exactly one of S_j^+ or S_j^-.

We can then partition C into the subset A of paths where both endpoints are in either S_j^+ or S_j^-, and the subset B of paths where one endpoint is in each of S_j^+ and S_j^-. Since $|C| = |A| + |B|$ is odd, $|A|$ and $|B|$ have different parity. To show that C satisfies part ii) of Definition 3.10, we must show that $|A|$, the number of paths in C where the endpoints are either both greater than or both less than the midpoint, is odd.

All paths in C have endpoints that are adjacent, so to find $|B|$ we need only consider the number of edges along the outer cycle of W that connect vertices in S_j^+ and S_j^-. Either one of S_j^+ or S_j^- is empty, in which case B is empty, or there is some number m of disjoint paths in the outer cycle of W where all vertices in the path are in S_j^+. The two endpoints of each of these disjoint paths are both incident to edges that connect vertices in S_j^+ and S_j^-. As the paths are disjoint, each endpoint of a path contributes exactly one to $|B|$, and we conclude that $|B| = 2m$, which is even. Then $|A|$ is odd, and C satisfies part ii) of Definition 3.10.

Therefore, W can be covered by paths of length two regardless of how its vertices are ordered. □

Having now solidified our understanding of Definition 3.10, we can approach the combinatorial characterization in [9] directly.

For ease, however, we will prove a short lemma first.

Lemma 3.15. *Let G be a graph with $V(G) = \{v_1, \ldots, v_n\}$ that is covered by a set C of length 2 paths. The following statements are equivalent.*[9]

(1) The number of paths $v_iv_kv_j$ in C in which $k < i, j$ or $k > i, j$ is odd.

(2) Let A be the sum over all pairs $i < j$ of the number of length 2 paths in C containing v_i and having v_j as an endpoint. Then A is odd.

(3) Let B be the number of pairs $v_i, v_j \in V(G)$ with $i < j$ such that the number of length 2 paths in C containing v_i and having v_j as an endpoint is odd. Then B is odd.

Li *et al.* left the proof of this lemma to the reader. We include it here.

Proof. Consider a path $v_iv_kv_j \in C$. We will assume without loss of generality that $i < j$, since we consider $v_iv_kv_j$ and $v_jv_kv_i$ to be the same path in the undirected graph.

In our first case, assume $i < k < j$. Then when calculating A, $v_iv_kv_j$ is counted both as a path in C containing v_k and having v_j as an endpoint and as a path in C containing v_i and having v_j as an endpoint. Since $v_iv_kv_j$ is counted exactly twice, there are an even number of terms from this kind of path in A.

Next, consider a path $v_iv_kv_j$ where $i < j < k$. Then when calculating A, $v_iv_kv_j$ is counted only once, as a path in C containing v_i and having v_j as an endpoint. Then there are an odd number of terms from this kind of path in A.

Finally, consider a path $v_iv_kv_j$ where $k < i < j$. Then when calculating A, $v_iv_kv_j$ is counted both as a path in C containing v_k and having v_j as an endpoint and as a path in C containing v_i and having v_j as an

endpoint, but further as a path in C containing v_k and having v_i as an endpoint. In other words, paths of this form are counted three times, and there are an odd number of terms from this kind of path in A.

If the number of paths $v_iv_kv_j$ in C in which $k < i, j$ or $k > i, j$ is odd, there are an odd number of paths in C that contribute an odd number to the sum A, and all the rest of the paths contribute an even number. Hence, A is odd. If the number of paths $v_iv_kv_j$ in C in which $k < i, j$ or $k > i, j$ is even, there are an even number of paths in C that contribute an odd number to the sum A, and all the rest of the paths contribute an even number. Hence, A is even, and (1) and (2) are equivalent statements.

We can break the sum A into two parts. One is the sum of the number of paths in C containing v_i and having v_j as an endpoint over such pairs $i < j$ where this number is even–in other words, the sum over the pairs not counted in B. This first sum is clearly even, as all of its terms are even. The other part of A is the sum of the number of paths in C containing v_i and having v_j as an endpoint over the B pairs $i < j$ where this number is even. This second sum has B terms, each of which is odd.

Thus if B is even, the two parts of A are both even, and A is itself even. If B is odd, one part of A is even and the other is odd, and A is odd. Therefore, (2) and (3) are equivalent statements. □

Note that if G is covered by length 2 paths, by Definition 3.10 the three statements are true as well as equivalent.

Now we present the combinatorial characterization from [9].

Theorem 3.16. *A graph G has a linear Nullstellensatz certificate of non-3-colorability over $\mathbb{F}_2$ if and only if G can be covered by length 2 paths.* [9]

Before we prove this theorem, we will provide some setup, and then discuss an error in the proof in [9]. Following this discussion, we provide a proof which corrects the error.

We are looking for a Nullstellensatz certificate of infeasability of the system

$$J_G = \{x_a^3 - 1 = 0, x_a^2 + x_a x_b + x_b^2 = 0 \mid a \in V, \{a, b\} \in E\}$$

where we take our polynomials to have coefficients in $\mathbb{F}_2$ (We can use this finite field because we are considering 3-colorability, and 3 is relatively prime to 2). To begin the proof in [9], the authors state (and prove in their Appendix A) that J_G and the system

$$F = \{x_i^2 x_j + x_i x_j^2 + 1 = 0, x_i^2 x_k + x_j^2 x_k + x_i x_j^2 + x_i x_k^2 = 0 | \{v_i, v_j\}, \{v_j, v_k\} \in E\}$$

have the same solution set, so if 1 is a degree one combination of polynomials in J_G, 1 is also a degree one combination of polynomials in F. Since we know that G has a linear Nullstellensatz certificate of non-3-colorability if and only if 1 is a degree one combination of polynomials in J_G, it is sufficient to show that 1 is a degree one combination of polynomials in F.

First, suppose G is covered by a set C of length 2 paths. We will consider a subset H of F that is determined by the paths in C. Let H be the set containing the polynomials

(1) $x_i^2x_k + x_j^2x_k + x_ix_j^2 + x_ix_k^2$ for each path $v_iv_jv_k \in C$.

(2) $x_i^2x_j + x_ix_j^2 + 1$ for each pair $v_i, v_j \in V(G)$ with $i < j$ such that the number of length 2 paths in C containing v_i and having v_j as an endpoint is odd.

To show that 1 is a degree one combination of polynomials in H, and therefore in F, consider the non-constant monomials in H. All of them have the form $x_r^2x_s$, where v_r and v_s can be any vertices in G. We will show that each of these monomials appears an even number of times in polynomials in H. This will imply that all non-constant terms in $\sum_{h\in H} h$ vanish, because our coefficients are in $\mathbb{F}_2$.

Remark 3.17. It is at this point that an error appears in the proof presented in [9].

The authors claim there are four ways the monomial $x_r^2x_s$ can appear, counting:

(a) one for each path in C with v_r and v_s as endpoints.

(b) one for each path in C with v_r as the middle vertex and v_s as an endpoint, as well as

(c) one if the number of paths in C containing v_r with v_s as an endpoint is odd, and

(d) one if the number of paths in C containing v_s with v_r as an endpoint is odd.

If this were true, we would not be able to cancel the non-constant monomials in all cases. To see this, consider the following.

Since we are working in $\mathbb{F}_2$, the combined contribution of (a) and (b) is the parity of the number of paths in C of the form $v_r v_i v_s$ or $v_i v_r v_s$ for any $v_i \in V(G)$. We then see that (c) also contributes the parity of the number of paths in C of the form $v_r v_i v_s$ or $v_i v_r v_s$, since these are the two forms a path containing v_r with v_s as an endpoint could take. By Definition 3.10, the number of paths containing the edge $\{v_r, v_s\}$–that is, paths of the form $v_i v_r v_s$ or $v_i v_s v_r$–is even. Therefore the number of paths in C of the form $v_i v_r v_s$ has the same parity as the number of paths of the form $v_i v_s v_r$. This means that, as (d) contributes the parity of the number of paths in C of the form $v_r v_i v_s$ or $v_i v_s v_r$, the combined contribution of (a) and (b) is equal to the contribution of (c) and to the contribution of (d). Then if the number of paths containing v_r with v_s as an endpoint is odd, $x_r^2 x_s$ is left with a coefficient of 1 in $\mathbb{F}_2$, so the non-constant monomials would not necessarily cancel.

Fortunately, the monomial $x_r^2 x_s$ does not appear exactly as claimed above. We will now present a proof that fixes the error.

Proof. First, consider the backwards direction. Recall that to show that 1 is a degree one combination of polynomials in J_G, it is sufficient to show that 1 is a degree one combination of polynomials in F, for instance of the polynomials in $H \subset F$:

(1) $x_i^2 x_k + x_j^2 x_k + x_i x_j^2 + x_i x_k^2$ for each path $v_i v_j v_k \in C$.

(2) $x_i^2 x_j + x_i x_j^2 + 1$ for each pair $v_i, v_j \in V(G)$ with $i < j$ such that the number of length 2 paths in C containing v_i and having v_j as an endpoint is odd.

Examining the polynomials in H, we find that the monomial $x_r^2 x_s$ appears in one of these three ways (the third depending on whether $r > s$): for each vertex $v_i \in V(G)$, there is

(a) one for each path $v_r v_i v_s \in C$ (from the first or last term in a polynomial of the form (1)).

(b) one for each path $v_i v_r v_s \in C$ (from one of the middle terms in a polynomial of the form (1)), and either

(c) one if the number of paths $v_i v_r v_s$ and $v_r v_i v_s \in C$ over all i is odd and $r < s$, or one if the number of paths $v_i v_s v_r$ and $v_r v_i v_s \in C$ over all i is odd and $r > s$ (from one of the first two terms of polynomial of the form (2)).

By Definition 3.10, the number of paths in C of the form $v_i v_r v_s$ has the same parity as the number of paths of the form $v_i v_s v_r$, so the contribution of (c) is the same whichever of r and s is greater. It is clear that when $r < s$, the number of paths that contribute to the parity of (c) is equal to the number of paths that count in the combined contribution of (a) and (b), since they are exactly the same set of paths. Then whichever of r and s is greater, in $\mathbb{F}_2$ the combined contribution of (a) and (b) is equal to the contribution of (c), so their sum will be 0 in $\mathbb{F}_2$, and all non-constant monomials will cancel.

Now that we know that all non-constant monomials in $\sum_{h \in H} h$ vanish, we must only determine the constant term in the sum. Only the polynomials of the form (2) above contribute a constant to $\sum_{h \in H} h$, and each of these polynomials contributes exactly one to the constant term. So the constant term is simply the parity of the number of pairs

$v_i, v_j \in V(G)$ $(i < j)$ such that the number of length 2 paths in C containing v_i and having v_j as an endpoint is odd.

Since G is covered by the set C of length 2 paths, by Lemma 3.15 the number of pairs of this form is odd, so the constant term in $\sum_{h\in H} h$ is 1 in $\mathbb{F}_2$.

Thus 1 is a combination of polynomials in H, and therefore also of polynomials in J_G.

Now we will consider the forwards direction. Suppose that J_G has a linear Nullstellensatz certificate of infeasability. Then 1 is a degree one combination of polynomials in F, so there is some set $H \subseteq F$ such that $\sum_{h\in H} h = 1$.

We then take $J \subseteq H$ to be the polynomials in H of the form $x_i^2x_k + x_j^2x_k + x_ix_j^2 + x_ix_k^2$, and construct C to consist of the paths $v_iv_jv_k$ where the polynomial of the form above has a nonzero coefficient in $\sum_{h\in J} h$. We will show that C is a covering of G according to Definition 3.10.

Suppose $\{v_r, v_s\} \in E$, and let $S_{r,s}$ be the sum of the coefficients of the monomials $x_r^2x_s$ and $x_rx_s^2$ that appear in $\sum_{h\in J} h$. We know that $\sum_{h\in H} h = 1$, and the only summand of $S_{r,s}$ in $\sum_{h\in H} h$ and not in $\sum_{h\in J} h$ is $x_r^2x_s + x_rx_s^2 + 1$, so $S_{r,s}$ is 0 in $\mathbb{F}_2$.

The contribution of a single polynomial $x_i^2x_k + x_j^2x_k + x_ix_j^2 + x_ix_k^2$ in $\sum_{h\in J} h$ to $S_{r,s}$ is 1 when $\{v_r, v_s\}$ is an edge on the path $v_iv_jv_k$, 2 if v_r and v_s are the endpoints of $v_iv_jv_k$, and 0 otherwise. As $S_{r,s}$ is 0 in $\mathbb{F}_2$, the edge $\{v_r, v_s\}$ lies on an even number of paths in C, and condition i) of Definition 3.10 holds.

Each edge $\{v_i, v_j\}$ with $i < j$ contributes a 1 to the sum $\sum_{h\in H} h$. Because the monomial $x_i^2 x_j$ appears an even number of times in H, and once in $H\backslash J$, it appears an odd number of times in J. We know that $x_i^2 x_j$ appears once in J for each path in C containing v_i and having v_j as an endpoint.

Since the number of 1s appearing in $\sum_{h\in H} h$ is odd, there are an odd number of monomials $x_i^2 x_j$ in $H\backslash J$ and in J to cancel. This means there are an odd number of pairs $i < j$ such that the number of length 2 paths in C containing v_i and having v_j as an endpoint is odd. By Lemma 3.15, we have condition ii) of Definition 3.10.

If $v_r, v_s \in V$ but $\{v_r, v_s\} \notin E$, then any $x_r^2 x_s$ term in $\sum_{h\in H} h$ is also in $\sum_{h\in J} h$. Thus the coefficient of $x_r^2 x_s$ in $\sum_{h\in J} h$ is 0 in $\mathbb{F}_2$. As $x_r^2 x_s$ only appears once in the polynomial in J corresponding to the path $v_i v_j v_k$ when v_r and v_s are endpoints of the path. Hence the number of paths whose endpoints are v_r and v_s is even, and condition iii) of Definition 3.10 holds. □

Since the combinatorial characterization relies on an ordering of the vertices of the graph, the following corollary is useful in ensuring that any ordering will satisfy the characterization.

Corollary 3.18. *If a graph can be covered by paths of length two in accordance with Definition 3.10 for some ordering of its vertices, it can be covered by paths of length two for all orderings of its vertices.*

Proof. Let G be a graph. By Theorem 3.16, if G can be covered by paths of length two, G has a degree one Nullstellensatz certificate of

non-3-colorability. Then we can find this certificate explicitly. Reordering the vertices of G will not alter the Nullstellensatz certificate, as changing the labels on the vertices does not change the correspondence between vertices and variables in J_G. Then the reordered graph has a degree one Nullstellensatz certificate of non-3-colorability, and by Theorem 3.16, it too can be covered by paths of length two. □

3.4. **Graphs That Can Be Covered by Length 2 Paths.** Now that we have a combinatorial characterization of graphs that have linear Nullstellensatz certificates of non-3-colorability, we will consider some broad examples of graphs that can be covered by length 2 paths.

We saw in Proposition 3.14 that all odd wheels can be covered by length 2 paths, so by Theorem 3.16, all odd wheels have linear Nullstellensatz certificates of non-3-colorability.

With the following definition, we can extend this fact to find more examples of graphs that can be covered by length 2 paths.

Definition 3.19. A graph H is a **subgraph** of a graph G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$.

A graph H is an **induced subgraph** of a graph G if $V(H) \subseteq V(G)$ and $E(H) = \{\{u, v\} | u, v \in V(H) \text{ and } \{u, v\} \in E(G)\}$.

Any graph is a subgraph (indeed, an induced subgraph) of itself.

Example 3.20. Figure 8 shows a subgraph of the graph in Figure 7. Figure 9 shows the induced subgraph of the graph in Figure 7 on the vertices v_1, v_2, and v_4.

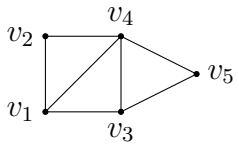

FIGURE 7. The same graph of order five, reprinted again.

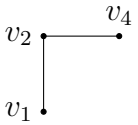

FIGURE 8. A subgraph.

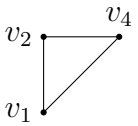

FIGURE 9. The induced subgraph on vertices v_1, v_2, and v_4.

If some subgraph G' of a graph G has a linear Nullstellensatz certificate of non-3-colorability, then G has one as well–in fact, the same one will work. This is true because $J_{G'} \subseteq J_G$, so if we have

$$1 = \sum_{i=1}^{r} \beta_i f_i$$

for some linear polynomials $\beta_1, \ldots, \beta_r \in \mathbb{K}[x_1, \ldots, x_n]$, where all $f_i \in J_{G'}$, then it is clear that

$$1 = \sum_{i=1}^{r} \beta_i g_i$$

for $g_i \in J_G$ and some linear polynomials $\beta_1, \ldots, \beta_r \in \mathbb{K}[x_1, \ldots, x_n]$, because we can set $\beta_i = 0$ for any g_i not in $J_{G'}$.

Therefore, it is certainly true that all graphs containing odd wheels as subgraphs can be covered by length 2 paths, and so have linear Nullstellensatz certificates of non-3-colorability by Theorem 3.16.

However, there are graphs that have linear Nullstellensatz certificates of non-3-colorability but do not contain an odd wheel. We will now present an example of this kind.

Example 3.21. The graph in Figure 10 is not 3-colorable, and has a Nullstellensatz certificate of degree one. Additionally, it does not contain an odd wheel.

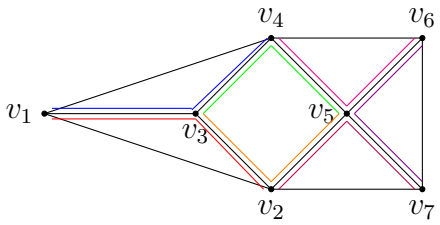

FIGURE 10. A graph that is not 3-colorable and contains no odd wheels, covered by length two paths.

This graph can be covered by a set of length 2 paths as follows:

$C = \{v_1v_3v_2, v_1v_3v_4, v_3v_2v_5, v_3v_4v_5, v_2v_5v_7, v_4v_5v_6, v_6v_5v_7\}$

This satisfies Definition 3.10: there are three paths in C whose midpoint is greater than or less than both endpoints, every edge in the graph is contained in an even number of paths in C, and the only non-adjacent endpoints of a path in C, v_3 and v_5, are the endpoints of two paths in C.

It would be useful to know if this example is the smallest graph that can be covered by length 2 paths (i.e. has a linear Nullstellensatz

certificate of non-3-colorability) but contains no odd wheels. Clearly all graphs of order less than four are 3-colorable, as we can assign each vertex a different color. Further, the only graph of order four that is not 3-colorable is K_4, which is an odd wheel.

Proposition 3.22. *All graphs of order less than seven are either 3-colorable or contain an odd wheel.*

We will prove this proposition in two parts. We have already discussed graphs of order less than five, so we will prove first that the proposition holds for graphs of order five, and then that it holds for graphs of order six.

The following definitions facilitate our notation in the proof.

Definition 3.23. Let v be a vertex of a graph G. The **neighborhood** of v, denoted $\mathcal{N}(v)$, is the set of vertices of G that are adjacent to v.

The **closed neighborhood** of v, denoted $\mathcal{N}[v]$, is $\mathcal{N}(v) \cup \{v\}$.

Example 3.24. In Figure 10, $\mathcal{N}(v_1) = \{v_2, v_3, v_4\}$.

Definition 3.25. The **degree** of a vertex $v \in V(G)$ is the number of vertices in $V(G)$ that are adjacent to v: that is, the size of $\mathcal{N}(v)$.

Example 3.26. In Figure 10, v_1 has degree three, while v_5 has degree four.

Definition 3.27. Let v be a vertex of a graph G. We denote by $G\backslash\{v\}$ the graph with vertex set $V(G)\backslash\{v\}$ and edge set

$$E(G)\backslash\{\{x, v\} \in E(G) | x \in V(G)\}.$$

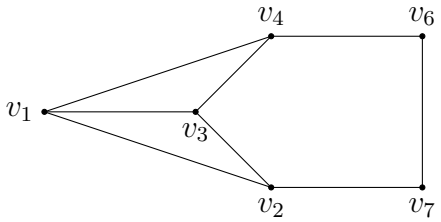

FIGURE 11. Deleting v_5 from Figure 10.

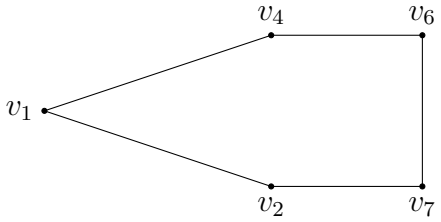

FIGURE 12. Deleting v_3 and v_5 from Figure 10.

In other words, $G\backslash\{v\}$ is the graph obtained from G by removing the vertex v and all edges incident to v. This process is called **deletion**, and can be applied to a single vertex or to a set of vertices.

Example 3.28. Deleting v_5 from the graph in Figure 10 gives us the graph in Figure 11.

Deleting v_3 and v_5 from the graph in Figure 10 gives us the graph in Figure 12.

Definition 3.29. A graph G is **bipartite** if $V(G) = V_1 \cup V_2$ with $V_1 \cap V_2 = \emptyset$, and every edge in G has one endpoint in V_1 and one in V_2.

We can see from this definition that a graph is bipartite if and only if it is 2-colorable.

Example 3.30. The graph in Figure 13 is bipartite: we can partition the vertices into the sets $V_1 = \{v_1, v_3\}$ and $V_2 = \{v_2, v_4\}$.

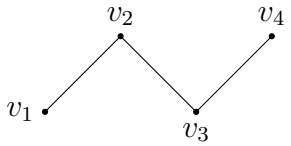

FIGURE 13. The same path of length 3.

Theorem 3.31. *A graph is bipartite if and only if it contains no odd cycles* [7].

Proof. (Proof of Proposition 3.22, Order Five) We know by Theorem 3.16 and Proposition 3.14 that graphs containing an odd wheel are not 3-colorable, and in fact that they have a degree one Nullstellensatz certificate of non-3-colorability. It remains to show that all graphs of order five and six that do not contain odd wheels are 3-colorable.

Suppose we have a connected graph G of order five (we can view chromatic number in terms of connected components, so disconnected graphs of order five can be considered as sets of smaller graphs for our purposes) that does not contain an odd wheel. Deleting an arbitrary vertex v from G gives us a graph $G\backslash\{v\}$ of order four. Since G did not contain an odd wheel, neither does $G\backslash\{v\}$. Because all non-3-colorable graphs of order four contain K_4, it follows that $G\backslash\{v\}$ is 3-colorable. Note that if $G\backslash\{v\}$ is 2-colorable, then G is 3-colorable. Also, since $G\backslash\{v\}$ is 3-colorable, if $|\mathcal{N}(v)| \leq 2$ we clearly have one of the three colors available for v. Therefore, assume that $|\mathcal{N}(v)| \geq 3$.

If $\mathcal{N}(v)$ contains a 3-cycle, then G contains K_4. We know that $\mathcal{N}(v)$ contains at most four vertices, so since G contains no odd wheels, $\mathcal{N}(v)$ contains no odd cycles. By Theorem 3.31, $\mathcal{N}(v)$ is bipartite and hence 2-colorable.

Suppose we have colored $\mathcal{N}(v)$ using colors 1 and 2. We can safely give v color 3. There is at most one vertex in $G\backslash\mathcal{N}[v]$, since v has at least three neighbors, and G has order five. If there are no such vertices, we have already colored G. If there is one, giving it color 3 provides the desired 3-coloring. Thus, G is 3-colorable, so all graphs of order five that do not contain odd wheels are 3-colorable. □

Proof. (Proof of Proposition 3.22, Order Six) Now suppose we have a connected graph H of order six that does not contain an odd wheel. Deleting an arbitrary vertex v from H gives us a graph $H\backslash\{v\}$ of order five. Since H did not contain an odd wheel, neither does $H\backslash\{v\}$. We have shown that all non-3-colorable graphs of order five contain K_4, so it follows that $H\backslash\{v\}$ is 3-colorable, for the same reasons as above. We see again that since $G\backslash\{v\}$ is 3-colorable, if $|\mathcal{N}(v)| \leq 2$ we clearly have one of the three colors available for v, so we assume that $|\mathcal{N}(v)| \geq 3$.

If $\mathcal{N}(v)$ contains a 3-cycle, then H contains K_4, and if $\mathcal{N}(v)$ is a 5-cycle, then H is W_5. We know that $\mathcal{N}(v)$ contains at most five vertices, so if H contains no odd wheels, $\mathcal{N}(v)$ contains no odd cycles. By Theorem 3.31, $\mathcal{N}(v)$ is bipartite and hence 2-colorable.

Suppose we have colored $\mathcal{N}(v)$ using colors 1 and 2. We can safely give v color 3. There are at most two vertices in $H\backslash\mathcal{N}[v]$. If there are no such vertices, we have a 3-coloring of H, and if there is one such vertex, giving it color 3 provides the desired 3-coloring, so suppose there are exactly two vertices in $H\backslash\mathcal{N}[v]$, which we shall call w_1 and w_2. If they are not adjacent, both of them can have color 3. If they are adjacent, and both are adjacent to vertices in $\mathcal{N}(v)$ with different colors, they

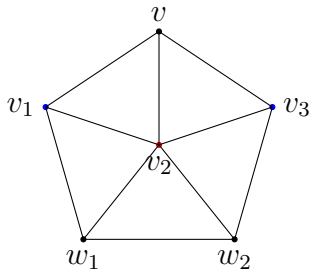

FIGURE 14. When w_1 and w_2 are both are adjacent to vertices in $\mathcal{N}(v)$ with different colors and have one neighbor in common.

must have at least one neighbor in common. (This is because $\mathcal{N}(v)$ has three elements, so if they are colored with two colors there is some vertex that does not share a color with the other neighbors of v.) If w_1 and w_2 are both are adjacent to vertices in $\mathcal{N}(v)$ with different colors and have one neighbor in common, then H is W_5, an odd wheel (see Figure 14). If they have two or more neighbors in common, then H contains K_4, an odd wheel. Therefore, if H does not contain an odd wheel, at least one of w_1 and w_2 is not adjacent to two differently-colored elements of $\mathcal{N}(v)$. Then we can assign that vertex either color 1 or 2, and the other can have color 3, which gives us a 3-coloring of H.

Hence, H is 3-colorable, so all graphs of order six that do not contain odd wheels are 3-colorable. Thus we have shown that our statement holds for all graphs of order less than seven.

□

4. 2-Colorability

In the previous sections, we discussed a combinatorial characterization of graphs that have linear Nullstellensatz certificates of non-3-colorability. It is of interest to consider whether there is a similar characterization of graphs that have linear Nullstellensatz certificates of non-2-colorability. In this section we find a characterization of such graphs, and determine their Nullstellensatz certificates of non-2-colorability explicitly.

Our task is much easier in this case because we can already characterize graphs that are not 2-colorable: by Theorem 3.31 they are exactly those graphs that contain an odd cycle. Thus we only need to consider which of these graphs have linear Nullstellensatz certificates of non-2-colorability.

Proposition 4.1. *Let C_{2k-1} be a cycle of length $2k-1$. For simplicity, label the vertices $1, 2, \ldots, 2k-1$ in order around the cycle. Let $\{f_1, \ldots, f_{2k-1}\}$ be the subset of $J_{C_{2k-1}}$ where $f_i = x_i^2 - 1$, and let $\{g_1, \ldots, g_{2k-1}\}$ be the subset of $J_{C_{2k-1}}$ where $g_i = x_i + x_{i+1}$ for $1 \leq i \leq 2k-2$ and $g_{2k-1} = x_1 + x_{2k-1}$.*

$$N(C_{2k-1}) = 2f_1 + \sum_{i=1}^{k} 2x_1 g_{2i-1} + \sum_{i=1}^{k-1} x_1 g_{2i} = 1$$

is a degree one Nullstellensatz certificate of non-2-colorability of C_{2k-1}.

Proof. First note that C_{2k-1} is 2-colorable if and only if $J_{C_{2k-1}}$ has a common zero over $\mathbb{F}_3$, so we consider $N(C_{2k-1})$ over $\mathbb{F}_3$. In other words, terms in a sum will cancel if their coefficients add up to a multiple of three.

Certainly $N(C_{2k-1})$ has the form $\sum_{h_i \in J_{C_{2k-1}}} \beta_i h_i$ where the β_i are polynomials of degree at most one. So to show that we have a degree one Nullstellensatz certificate, we need only justify our claim that $N(C_{2k-1}) = 1$. To see this, we will examine each monomial in $N(C_{2k-1})$ individually.

The monomial x_1^2 appears in $f_1, x_1 g_1$, and $x_1 g_{2k-1}$, with coefficient 2 in each case. Thus we have $(2 + 2 + 2)x_1^2 = 0$, and x_1^2 cancels in $N(C_{2k-1})$.

Each monomial $x_1 x_i$ where $i \neq 1$ appears in $x_1 g_{i-1}$ and $x_1 g_i$. In all cases, one of $x_1 g_{i-1}$ and $x_1 g_i$ is even and the other is odd, so $x_1 x_i$ appears once with coefficient 1 and once with coefficient 2. Thus $x_1 x_i$ cancels in $N(C_{2k-1})$.

Finally, the term -1 in appears only in f_1, where it has coefficient 2. As $-2 = 1$ in $\mathbb{F}_3$, we are left with $N(C_{2k-1}) = 1$. □

It is clear that if a subgraph of a graph G has a degree one Nullstellensatz certificate of non-2-colorability, then G has such a certificate also. In fact, we can use the same one.

Thus, any graph containing an odd cycle has a degree one Nullstellensatz certificate of non-2-colorability, which can be defined explicitly from the Proposition above. Since all non-2-colorable graphs contain

an odd cycle, we have shown that all non-2-colorable graphs have a Nullstellensatz certificate of degree one.

Example 4.2. We will find a degree one Nullstellensatz certificate of non-2-colorability for C_3. Label the vertices of the three-cycle from $\{1, 2, 3\}$, and consider the following polynomials: $f_1 = x_1^2 - 1$, $f_2 = x_2^2 - 1$, $f_3 = x_3^2 - 1$, $g_1 = x_1 + x_2$, $g_2 = x_2 + x_3$, $g_3 = x_1 + x_3$.

Applying Proposition 4.1, we use coefficients from $\mathbb{F}_3$ and take

$$2f_1 + 2x_1g_1 + x_1g_2 + 2x_1g_3$$

$$= 2x_1^2 - 2 + 2x_1(x_1 + x_2) + x_1(x_2 + x_3) + 2x_1(x_1 + x_3)$$

$$= (2 + 2 + 2)x_1^2 + (2 + 1)x_1x_2 + (1 + 2)x_1x_3 - 2 = -2 = 1.$$

Therefore, $2f_1 + 2x_1g_1 + x_1g_2 + 2x_1g_3 = 1$ is a degree one Nullstellensatz certificate of non-2-colorability for C_3.

5. Hamiltonian graphs

In this section, we discuss the third section of [4], which deals with a class of graphs that we will now define.

Definition 5.1. A **Hamiltonian cycle** of a graph G is a subgraph of G that is a cycle on the same number of vertices as G.

If a graph G contains a Hamiltonian cycle, we say that G is Hamiltonian.

Example 5.2. The graph in Figure 15 is Hamiltonian; $v_1, v_2, v_4, v_5, v_3, v_1$ is a Hamiltonian cycle.

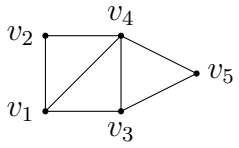

Figure 15. A Hamiltonian graph.

On the other hand, the graph in Figure 16 is not Hamiltonian.

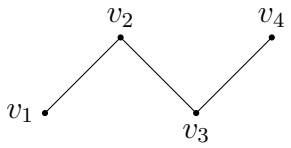

Figure 16. Paths are not Hamiltonian.

We can detect a Hamiltonian cycle in a graph using a system of polynomials, similar to our method of determining whether or not a

graph is k-colorable. The method below can be used for both directed and undirected graphs.

This section of [4] introduces an ideal whose variety is nonempty if and only if the associated graph is Hamiltonian, and uses Gröbner bases to show how this ideal can be written as an intersection of other ideals, allowing us to characterize uniquely Hamiltonian graphs.

In Section 5.1, we introduce the Hamiltonian ideal and consider subtle requirements of its definition. In Section 5.2, we explore an algebraic encoding of cycles in a graph, and discuss properties of the ideal generated by said encoding. In Section 5.3, we present the decomposition of the Hamiltonian ideal into an intersection of cycle ideals.

5.1. **The Hamiltonian Ideal.** In the following proposition, we define the Hamiltonian ideal of a graph G.

Proposition 5.3. *Let $G = (V, A)$ be a simple directed graph on vertices $V = \{1, \ldots, n\}$. Assume that the characteristic of $\mathbb{K}$ is relatively prime to n and that $\omega \in \mathbb{K}$ is a primitive n-th root of unity. Consider the following system in $\mathbb{K}[x_1, \ldots, x_n]$:*

$$H_G = \{x_i^n - 1, \prod_{j \in \delta^+(i)} (\omega x_i - x_j) \mid i \in V\}.$$

Here, $\delta^+(i)$ denotes those vertices j where there is an arc going from i to j in G. The system H_G has a solution over $\mathbb{K}$ if and only if G has a Hamiltonian cycle.

As stated, we can only prove one part of the if and only if statement. We will discuss the other part below.

Certainly, if G is Hamiltonian, H_G has a solution over $\mathbb{K}$. We will follow the argument of Lemma 3.8 in [4].

Proof. Choose a starting vertex in a Hamiltonian cycle in G and label it $\omega^0 = 1$. Then successively label vertices along the cycle with one higher power of ω than the previous vertex. All labels x_i are nth roots of unity, so all equations of the form $x_i^n - 1 = 0$ in H_G hold. Further, every vertex is connected by an arc to a vertex labeled with the next higher power of ω, so all equations of the form $\prod_{j \in \delta^+(i)} (\omega x_i - x_j) = 0$ in H_G hold. Therefore, H_G has a solution over $\mathbb{K}$. □

However, there are graphs that have a solution to H_G over $\mathbb{K}$ that are not Hamiltonian. For instance, we will show that all trees (graphs containing no cycles) have a solution to H_G over $\mathbb{K}$, and because a graph must contain at least one cycle to be Hamiltonian, these are clearly false positives. To show this, we must first prove the following two lemmas.

Lemma 5.4. *Let T be a directed tree of order n. T contains a vertex of outdegree 0.*

Proof. Suppose that all vertices in T have outdegree of at least one. Then given any vertex $v \in V(T)$ there exists a directed walk beginning at v that does not terminate, since every vertex on the walk is incident to a directed edge to a vertex on which the walk can continue. The tree has finite order, so if the walk does not terminate it must eventually repeat some number of vertices; in other words, the walk contains a cycle. However, a graph containing a cycle is not a tree, so we have

a contradiction. Therefore, there is some vertex in T with outdegree 0. □

Lemma 5.5. *Let T be a directed tree of order n. Given any vertex $v \in V(T)$, there is a directed path from v to some vertex of outdegree 0.*

Proof. Since T has finite order, the length of a directed walk in T beginning at v is bounded by the argument in Lemma 5.4. This means that there is a longest directed walk in T that begins at v. Let w be the last vertex of this walk. If w has outdegree at least one, then we can extend our longest walk to include another vertex, contradicting the fact that the walk from v to w was the longest walk beginning at v. Therefore, w has outdegree 0. Since T is a tree, there are no cycles in the walk, and since T is (simply) directed, no edges in the walk can be repeated without the walk containing a cycle, so no vertices in the walk are repeated. Thus our walk is in fact a directed path from v to a vertex of outdegree 0. □

Theorem 5.6. *Let T be a directed tree of order n. There is a solution to H_T over $\mathbb{K}$.*

Proof. We will label the vertices of T as follows: If the shortest path from a vertex $i \in V(T)$ to a vertex of outdegree 0 has length k, then let $x_i = \omega^{n-k}$. Note that the power of ω will never be negative, as T has order n. We know that all vertices in T can be labeled this way because Lemma 5.5 states that all vertices in T are at the beginning of some path to a vertex of outdegree 0, and so each has some minimal path

length k that can be used to define the vertex's associated variable in H_T.

All variables in H_T are assigned n-th roots of unity, so the equations in H_T of the form $x_i^n - 1 = 0$ hold for all $i \in V(T)$. Further, for $k \geq 1$, following the shortest path from a vertex that is k edges away from a vertex of outdegree 0 brings us to a vertex that is $k-1$ edges away from the same endpoint. Therefore, each vertex is incident to at least one arc leading to a vertex labeled with one higher power of ω, so the equations in H_T of the form $\prod_{j \in \delta^+(i)}(\omega x_i - x_j) = 0$ also hold, as only one factor in the product needs to equal zero to satisfy the equation. Thus, we have a solution to H_T. □

Trees are not the only non-Hamiltonian graphs that have a solution to H_G.

Example 5.7. The directed 3-cycle C_3 in Figure 17 with $V(C) = \{v_1, v_2, v_3\}$ and $Arcs(C) = \{\{v_1, v_2\}, \{v_1, v_3\}, \{v_2, v_3\}\}$ has

$$H_C = \{x_1^3 - 1 = 0, x_2^3 - 1 = 0, x_3^3 - 1 = 0, (\omega x_1 - x_2)(\omega x_1 - x_3) = 0, \omega x_2 - x_3 = 0\}.$$

This system has a solution: $x_1 = 1, x_2 = \omega, x_3 = \omega^2$ will satisfy H_C.

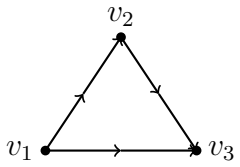

FIGURE 17. A directed cycle of length 3.

Example 5.8. The graph in Figure 18 with vertex set $V = \{v_1, v_2, v_3, v_4\}$ and directed edges $Arcs = \{\{v_1, v_2\}, \{v_2, v_3\}, \{v_3, v_1\}, \{v_1, v_4\}\}$ has

$H_G = \{x_1^4 - 1 = 0, x_2^4 - 1 = 0, x_3^4 - 1 = 0, x_4^4 - 1 = 0,$
$(\omega x_1 - x_2)(\omega x_1 - x_4) = 0, \omega x_2 - x_3 = 0, \omega x_3 - x_1 = 0\}$.
This system has a solution: $x_1 = 1, x_2 = \omega^2, x_3 = \omega^3, x_4 = \omega$ will satisfy H_G.

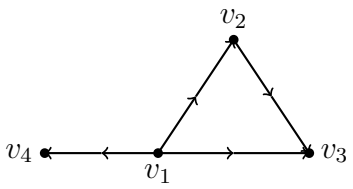

FIGURE 18. A directed graph of order four.

This trend might lead us to believe that all graphs that have a vertex of outdegree zero (none of which are Hamiltonian) have solutions to H_G, but while this is a necessary condition (we claim) for non-Hamiltonian graphs to have solutions to H_G, it is not sufficient.

Example 5.9. The graph in Figure 19 with vertex set
$V = \{v_1, v_2, v_3, v_4, v_5\}$ and directed edges
$Arcs = \{\{v_1, v_2\}, \{v_2, v_3\}, \{v_3, v_1\}, \{v_4, v_1\}, \{v_4, v_5\}\}$ has
$H_G = \{x_1^5 - 1 = 0, x_2^5 - 1 = 0, x_3^5 - 1 = 0, x_4^5 - 1 = 0, x_5^5 - 1 = 0,$
$\omega x_1 - x_2 = 0, \omega x_2 - x_3 = 0, \omega x_3 - x_1 = 0, (\omega x_4 - x_1)(\omega x_4 - x_5) = 0\}$.
The vertex 5 has outdegree 0, but there is no solution to H_G. This is because $x_1 = \omega x_3 = \omega(\omega x_2) = \omega(\omega(\omega x_1))$ because each of these vertices has outdegree one, so there are no choices for assigning their variables. Since ω is a primitive 5th root of unity, $\omega^3 \neq 1$, so there is no solution to H_G.

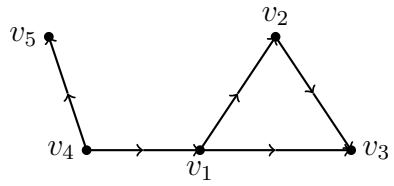

FIGURE 19. A directed graph of order five.

However, we can characterize some graphs G that we can be sure are Hamiltonian if and only if H_G has a solution.

Proposition 5.10. *Let G be a simple directed graph of order n where all vertices have outdegree at least one. If H_G has a solution over $\mathbb{K}$, then G is Hamiltonian.*

Proof. If H_G has a solution over $\mathbb{K}$, then the equations in H_G of the form $x_i^n - 1 = 0$ ensure that every vertex in G is labeled with some power of ω. Every vertex $i \in V(G)$ has outdegree at least one, which means there is an equation in H_G of the form $\prod_{j \in \delta^+(i)} (\omega x_i - x_j) = 0$ for each vertex i. Since H_G has a solution over $\mathbb{K}$, this means that every vertex is connected by an arc to a vertex labeled with the next higher power of ω. In other words, there is a directed path from a vertex labeled 1 to a vertex labeled ω to one labeled ω^2, and so on. We know that ω a primitive nth root of unity, so the nth vertex on this path, labeled ω^{n-1} must be connected by an arc to a vertex labeled $\omega^n = 1$. Because G has order n, and there have been n vertices on our path, each with different labels, there is only one vertex in G labeled 1: the vertex that began our path. Then our path must close up to become a cycle of length n, which means that G is Hamiltonian. $\square$

Notice that all points in the variety $\mathcal{V}(H_G)$ have the form

$$(\omega^i, \omega^{i+1}, \omega^{i+2}, \ldots, \omega^{i+(n-1)})$$

for some integer i.

We should note that this argument follows the line of Lemma 3.8 in [4], but that it did not hold for all connected directed graphs as claimed there. Restricting the outdegree of G is crucial to the proof.

There is an alternative way to restrict the outdegree of G, which is even simpler, but subtle. If we specify that the empty product (the product of no factors) is 1, then if $\delta^+(i)$ is empty (i.e. v_i has outdegree 0) we are left with $\prod_{j \in \delta^+(i)}(\omega x_i - x_j) = 1$. As $1 \neq 0$, this will mean that H_G has no solution, as desired.

It is likely that the authors of [4] had this in mind when formulating the system. However, the distinction is subtle enough to lead students without significant background astray, and so this fact and the consequences of its validity were worth exploring here.

5.2. **Cycle Ideals.** We proceed to a definition that will allow us to express the cycles in G algebraically.

Definition 5.11. (Cycle encodings). Let ω be a fixed primitive k-th root of unity. If C is a directed cycle of length k in a directed graph, with vertex set $V(C) = \{v_1, \ldots, v_k\}$ (we can assume without loss of generality that the vertices are indexed in order around the cycle), the cycle encoding of C is the following set of k polynomials:

$$g_i = \begin{cases} x_{v_{k-i}} - \omega^{k-i} x_{v_k} & i = 1, \ldots, k-1 \\ x_{v_k}^k - 1 & i = k \end{cases}$$

We may consider undirected cycles to be doubly-covered directed cycles; we let an edge $\{v_i, v_{i+1}\}$ in C correspond to two arcs in C: the arc from v_i to v_{i+1}, and the arc from v_{i+1} to v_i. We can choose either direction to travel around a doubly-covered undirected cycle.

If C is a doubly-covered cycle of length k in a directed graph, with vertex set $V(C) = \{v_1, \ldots, v_k\}$, the cycle encoding of C is the following set of k polynomials:

$$g_i = \begin{cases} x_{v_i} + \frac{(\omega^{2+i} - \omega^{2-i})}{(\omega^3 - \omega)} x_{v_{k-1}} + \frac{(\omega^{1-i} - \omega^{3+i})}{(\omega^3 - \omega)} x_{v_k} & i = 1, \ldots, k-2 \\ (x_{v_{k-1}} - \omega x_{v_k})(x_{v_{k-1}} - \omega^{-1} x_{v_k}) & i = k-1 \\ x_{v_k}^k - 1 & i = k \end{cases}$$

To see that this definition is a valid algebraic expression of cycles, consider the correspondence between the set of g_i and the given cycle.

In the directed case:

Certainly if we take a directed cycle in a directed graph, we can write it as a set of polynomials by choosing one of the vertices to be the first (that is, v_1).

If two cycles C_1 and C_2 have the same encoding set of polynomials, they must contain the same vertices, as the variables in the set of polynomials correspond to the vertices in the cycle. Also, these vertices must appear in the same order, since ω^j appears in the same polynomial as the variable corresponding to the j-th vertex of the cycle, and ω^j is

unique for $1 \leq j \leq k$. Therefore, C_1 and C_2 are the same cycle, and the correspondence from the set of g_i to the cycle is well-defined.

In the doubly-covered undirected case:

Again, the definition of the encoding allows us to write an undirected cycle as a set of polynomials, this time by choosing a pair of adjacent vertices to be v_k and v_{k-1}, thus fixing our choice of both the endpoint and the direction we are going around the cycle.

It is harder to see that the correspondence is well-defined in the undirected case, if only because the encoding is non-intuitive and contains many terms. We see again that two cycles C_1 and C_2 with the same encoding set of polynomials must contain the same vertices, for the reasons stated above. We also know that the last two vertices in C_1 must be the last two vertices in C_2 because only g_k has degree k, and only g_{k-1} has degree 2. Then we need to ensure that the first $k-2$ vertices of the two cycles are also in the same order. As the term x_{v_i} appears only in the polynomial g_i for $1 \leq i \leq k-2$, knowing that the i-th polynomials in the encodings of C_1 and C_2 are equal is enough to ensure that the i-th vertex of C_1 is also the i-th vertex of C_2. Then C_1 and C_2 are the same cycle in this case as well, and the correspondence from the set of g_i to the cycle is well-defined.

The correspondence is not one to one in either case, however. Taking any one of the k vertices of C to be v_1 will result in a different encoding of C, so there are k possible encodings of any directed cycle of length k. For an undirected cycle of length k, there are two possible directions

to encode the cycle for any starting vertex we choose, so there are $2k$ possible encodings of any undirected cycle of length k.

It is worth a brief explanation as to why the encoding of an undirected cycle is so complicated compared to that of a directed one. The encoding is designed so that the terms cancel at certain points, so its form is a result of reverse-engineering. We will discuss this cancellation more below, after the proof of Lemma 5.19.

Definition 5.12. (Cycle Ideals). The cycle ideal associated to a cycle C is $H_{G,C} = \langle g_1, \ldots, g_k \rangle \subseteq \mathbb{K}[x_{v_1}, \ldots, x_{v_k}]$, where the g_i are the cycle encoding of C. [4]

Lemma 5.13. *If the leading monomials of a generating set $F = \{f_1, \ldots, f_s\}$ of an ideal $I \subset \mathbb{K}[x_1, \ldots, x_n]$ are relatively prime, then F is a Gröbner basis of I. (Exercise 2.6.11 in* [2]*)*

Proof. We will prove the contrapositive of this statement.

Fix a monomial order. Suppose that a generating set $F = \{f_1, \ldots, f_s\}$ of an ideal $I \subset \mathbb{K}[x_1, \ldots, x_n]$ is not a Gröbner basis of I.

Then there is some $g \in I$ such that $LT(g)$ is not in $\langle LT(F) \rangle$. This means that $LT(g) = cr$ for some monomial $r \in \mathbb{K}[x_1, \ldots, x_n]$ that is less than any $LM(f_i)$ in the given monomial order, and some $c \in \mathbb{K}$. Then $LM(g) = r$.

We also know that since F generates I, $g = h_1 f_1 + \cdots + h_s f_s$ for some $h_i \in \mathbb{K}[x_1, \ldots, x_n]$. Then certainly $LM(g) = LM(h_1 f_1 + \cdots + h_s f_s)$, so

$$LM(h_1 f_1 + \cdots + h_s f_s) = r.$$

The only way that the leading monomial of the sum $h_1 f_1 + \cdots + h_s f_s$ can be less than the leading monomials of all generators f_i is if $LM(h_j f_j) = LM(h_t f_t)$ for some pair $j \neq t$, which makes the leading terms cancel. This implies that $LM(h_j)LM(f_j) = LM(h_t)LM(f_t)$. If $LM(h_j) = LM(h_t)$, then $LM(f_j) = LM(f_t)$, which means the leading monomials of F are not relatively prime, and we are done. Therefore, assume without loss of generality that $LM(h_j) > LM(h_t)$. Then we can divide both side of the equation by the smaller monomial and still be left with a monomial equation:

$$\frac{LM(h_j)}{LM(h_t)} LM(f_j) = LM(f_t).$$

This means that the monomial $LM(f_t)$ is a multiple of $LM(f_j)$, and therefore the leading monomials of F are not relatively prime.

Hence, if F is not a Gröbner basis of I, the leading monomials of F are not relatively prime, and the lemma is true. □

Lemma 5.14. *The cycle encoding polynomials $F = \{g_1, \ldots, g_k\}$ are a reduced Gröbner basis for the cycle ideal $H_{G,C}$ with respect to any term order in which $x_{v_1} > x_{v_2} > \cdots > x_{v_k}$.*

Proof. First note that the leading monomials in a cycle encoding of a doubly covered cycle or a directed cycle of length k are

$$\{x_{v_1}, \ldots, x_{v_{k-2}}, x_{v_{k-1}}^2, x_{v_k}^k\}$$

or

$$\{x_{v_1}, \ldots, x_{v_{k-1}}, x_{v_k}^k\},$$

respectively. We observe that the monomials in each set are relatively prime. Therefore by Lemma 5.13, F is a Gröbner basis for $H_{G,C}$.

In the directed case, the monomial $x_{v_{k-i}}$ appears only in g_i, and the monomials 1 and $x_{v_k}^k$ only appear in g_k. Further, the monomial x_{v_k} that appears in g_i for $1 \leq i \leq k-1$ is not in $\langle LT(H_{G,C})\rangle$, and so cannot be in $\langle LT(H_{G,C}\backslash\{g_i\})\rangle$. Thus, no monomial in g_i is in $\langle LT(H_{G,C}\backslash\{g_i\})\rangle$, and F is a reduced Gröbner basis of $H_{G,C}$.

In the undirected case, the monomial x_{v_i} appears only in g_i for $1 \leq i \leq k-2$, the monomials $x_{v_{k-1}}^2$, $x_{v_k}^2$, and $x_{v_{k-1}}x_{v_k}$ appear only in g_{k-1}, and the monomials 1 and $x_{v_k}^k$ only appear in g_k. Further, the monomials $x_{v_{k-1}}$ and x_{v_k} that appear in g_i for $1 \leq i \leq k-2$ are not in $\langle LT(H_{G,C})\rangle$, and hence are not in $\langle LT(H_{G,C}\backslash\{g_i\})\rangle$. Thus, no monomial in g_i is in $\langle LT(H_{G,C}\backslash\{g_i\})\rangle$, and F is again a reduced Gröbner basis of $H_{G,C}$. □

Lemma 5.15. *The ideal $H_{G,C}$ has $|\mathcal{V}(H_{G,C})| = k$ if C is directed.*

Proof. Fix any monomial order in which $x_{v_1} > x_{v_2} > \cdots > x_{v_k}$. We saw above that $\{g_i\}$ form a Gröbner basis for $H_{G,C}$.

This means that every point in $\mathcal{V}(H_{G,C})$ must be a solution to g_i for all i.

To find a solution to $g_k = x_{v_k}^k - 1$, we must set $x_{v_k} = \omega^j$ to be some k-th root of unity ($1 \leq j \leq k$). Fixing this value leaves us with a system of k linear equations coming from the g_i, which we express in

an augmented $k \times k$ matrix as follows:

$$\left(\begin{array}{ccccc|c} 1 & 0 & \dots & 0 & -\omega & 0 \\ 0 & 1 & \dots & 0 & -\omega^2 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -\omega^{k-1} & 0 \\ 0 & 0 & \dots & 0 & 1 & \omega^j \end{array}\right)$$

Conveniently, this matrix is already in Reduced Row-Echelon form. We note that all entries below the diagonal are zero, and that the diagonal itself is entirely comprised of ones. This means we have a pivot in every column of the matrix, so it has a unique solution. In other words, this matrix represents exactly one point in $\mathcal{V}(H_{G,C})$.

Now, there were k choices for our value of $x_{v_k} = \omega^j$, meaning that the degree of $H_{G,C}$ is k. Each choice gives us a distinct augmented matrix: the left-hand side will be the same for all k values of j, but the right-hand side will be different for each value of j (since the k k-th roots of unity are unique). We have found k points in $\mathcal{V}(H_{G,C})$, and we know that we have considered all possible points, as any other value we assign to x_{v_k} will not be a solution to g_k. Therefore, $|\mathcal{V}(H_{G,C})| = k$. $\square$

Lemma 5.16. *The ideal $H_{G,C}$ has $|\mathcal{V}(H_{G,C})| = 2k$ if C is undirected.*

Proof. Again we must find a solution for all g_i, so to solve $g_k = x_{v_k}^k - 1 = 0$, we set $x_{v_k} = \omega^j$ $(1 \leq j \leq k)$. We must also solve $g_{k-1} = (x_{v_{k-1}} - \omega x_{v_k})(x_{v_{k-1}} - \omega^{-1} x_{v_k}) = 0$. To do this, we will choose one of the quadratic's two factors to be zero: either $x_{v_{k-1}} = \omega x_{v_k}$, or $x_{v_{k-1}} = \omega^{-1} x_{v_k}$. (Note that since C is a cycle, it has length $k \geq 3$,

so $\omega \neq \omega^{-1}$ and the two factors are distinct.) Fixing these two values gives us a system of k linear equations, which we again express in an augmented $k \times k$ matrix:

$$\left(\begin{array}{ccccccc|c} 1 & 0 & \dots & 0 & 1 & \frac{(1-\omega^4)}{(\omega^3-\omega)} & 0 \\ 0 & 1 & \dots & 0 & \frac{(\omega^4-1)}{(\omega^3-\omega)} & \frac{(\omega^{-1}-\omega^5)}{(\omega^3-\omega)} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & \frac{(\omega^{2+(k-2)}-\omega^{2-(k-2)})}{(\omega^3-\omega)} & \frac{(\omega^{1-(k-2)}-\omega^{3+(k-2)})}{(\omega^3-\omega)} & 0 \\ 0 & 0 & \dots & 0 & 1 & \alpha & 0 \\ 0 & 0 & \dots & 0 & 0 & 1 & \omega^j \end{array}\right)$$

where α is either ω or ω^{-1}. This matrix is also in Reduced Row-Echelon form, and again all entries below the diagonal are zero, and the diagonal is entirely comprised of ones. Thus we have a pivot in every column of the matrix, so it has a unique solution, and represents exactly one point in $\mathcal{V}(H_{G,C})$.

As before, there were k choices for our value of $x_{v_k} = \omega^j$, but now there are also two choices for α. This gives us a total of $2k$ choices, meaning that $H_{G,C}$ has degree $2k$. We will show that each choice results in a distinct augmented matrix. The key fact to note is that there is exactly one way to write a matrix in reduced row-echelon form.

If we look at two matrices with the same j value and different α values, they will have two distinct reduced row-echelon forms, and so be different matrices. The same is true of two matrices with the same α value and different j values. This means we have found $2k$ points in $\mathcal{V}(H_{G,C})$. We know that any j value other than the k we considered that we assign to x_{v_k} will not be a solution to g_k, and any α value other

than the two we considered will not be a solution to g_{k-1}. Therefore, $|\mathcal{V}(H_{G,C})| = 2k$. □

Now that we know the size of its variety, we can explore some other properties of the ideal $H_{G,C}$.

Definition 5.17. An ideal I is **radical** if $f^m \in I$ for some integer $m \geq 1$ implies $f \in I$.

The radical of an ideal I is the ideal

$$\sqrt{I} = \langle f | f^m \in I \text{for some integer } m \geq 1 \rangle.$$

Then I is radical if and only if $I = \sqrt{I}$. [2]

Example 5.18. The ideal $\langle x_1 x_2 \rangle \subset \mathbb{K}[x_1, \ldots, x_n]$ is radical. If $f^m \in \langle x_1 x_2 \rangle$, then $f^m = h x_1 x_2$ for some $h \in \mathbb{K}[x_1, \ldots, x_n]$. Suppose $f \notin \langle x_1 x_2 \rangle$. Then $x_1 x_2$ does not divide f. However, if $x_1 x_2$ does not divide f, there is no way the $x_1 x_2$ can divide f^m, since the only factors of f^m are f. By contradiction, f is in $\langle x_1 x_2 \rangle$, and hence the ideal is radical.

Lemma 5.19. *The ideal $H_{G,C}$ is radical.*

To prove this, we must first give a definition.

Definition 5.20. An ideal I is **zero-dimensional** if its variety $\mathcal{V}(I)$ contains finitely many points.

The **degree** of a zero-dimensional ideal is the number of points in its variety, counted according to multiplicity.

Proof. By Theorem 2.10 of [3], since the degree of $H_{G,C}$ is equal to $|\mathcal{V}(H_{G,C})|$, $H_{G,C}$ is radical. □

In Lemma 5.15 and Lemma 5.16, we found the size of $\mathcal{V}(H_{G,C})$, but we can go further and explicitly describe the points in the variety.

In the directed case, we find that the point

$$(\omega, \omega^2, \ldots, \omega^k)$$

is in the variety, and therefore all cyclic permutations of this point are in the variety as well (these look like the point given, except that every term is multiplied by ω^i for some $1 \leq i \leq k$).

To check this, for each $1 \leq i \leq k-1$ evaluate

$$g_i(\omega, \omega^2, \ldots, \omega^k) = \omega^{k-i} - \omega^{k-i}\omega^k = \omega^{k-i} - \omega^{k-i} = 0$$

as $\omega^k = 1$. Further, any k-th root of unity will be a solution to $g_k = x_{v_k}^k - 1$, so the point is in $\mathcal{V}(H_{G,C})$ as desired.

In the undirected case, we find that both the points

$$(\omega, \omega^2, \ldots, \omega^k) \text{ and } (\omega^k, \omega^{k-1}, \ldots, \omega)$$

are in the variety, along with all the cyclic permutations of each.

To check this, for each $1 \leq i \leq k-2$ evaluate

$$\begin{aligned}
&(\omega^3 - \omega) g_i(\omega, \omega^2, \ldots, \omega^k) \\
&= (\omega^3 - \omega)\omega^i + (\omega^{2+i} - \omega^{2-i})\omega^{k-1} + (\omega^{1-i} - \omega^{3+i})\omega^k \\
&= \omega^{3+i} - \omega^{1+i} + \omega^{1+i+k} - \omega^{1-i+k} + \omega^{1-i+k} - \omega^{3+i+k} \\
&= 0
\end{aligned}$$

as $\omega^k = 1$, and

$$\begin{aligned}
&(\omega^3 - \omega) g_i(\omega^k, \omega^{k-1}, \ldots, \omega) \\
&= (\omega^3 - \omega)\omega^{1-i} + (\omega^{2+i} - \omega^{2-i})\omega^2 + (\omega^{1-i} - \omega^{3+i})\omega \\
&= \omega^{4-i} - \omega^{2-i} + \omega^{4+i} - \omega^{4-i} + \omega^{2-i} - \omega^{4+i} \\
&= 0.
\end{aligned}$$

Further,

$$\begin{aligned}
g_{k-1}(\omega, \omega^2, \ldots, \omega^k) &= (\omega^{k-1} - \omega^{k+1})(\omega^{k-1} - \omega^{k-1}) \\
&= (\omega^{k-1} - \omega^{k+1})(0) \\
&= 0
\end{aligned}$$

and

$$\begin{aligned}
g_{k-1}(\omega^k, \omega^{k-1}, \ldots, \omega) &= (\omega^{1-(k-1)} - \omega^{1-k+1})(\omega^{1-(k-1)} - \omega^{(1-k)-1}) \\
&= (0)(\omega^{1-(k-1)} - \omega^{(1-k)-1}) \\
&= 0
\end{aligned}$$

so both points are solutions to g_{k-1}, and any k-th root of unity will be a solution to $g_k = x_{v_k}^k - 1$, so the point is in $\mathcal{V}(H_{G,C})$ as desired.

5.3. **Decomposition of the Hamiltonian Ideal.** We now understand important properties of the cycle ideals $H_{G,C}$ of G, but we must consider how these relate to the Hamiltonian ideal H_G.

Lemma 5.21. *Let G be a connected directed or doubly-covered undirected graph on n vertices. Then*

$$\mathcal{V}(H_G) = \bigcup_C \mathcal{V}(H_{G,C}),$$

where the union is over all Hamiltonian cycles C in G.

Proof. By Proposition 5.10, every point in $\mathcal{V}(H_G)$ has the form $p_1 = (\omega^{i+1}, \omega^{i+2}, \ldots, \omega^{i+k})$ (or possibly $p_2 = (\omega^{i+k}, \omega^{i+k-1}, \ldots, \omega^{i+1})$ if G is undirected) if we index the vertices so that some Hamiltonian cycle of G is written $C_i = \{v_1, v_2, \ldots, v_n\}$. Then the point p_1 or p_2 is in $\mathcal{V}(H_{G,C_i})$, as we found in our discussion of variety membership above, and

$$\mathcal{V}(H_G) \subseteq \bigcup_C \mathcal{V}(H_{G,C}).$$

From said discussion we also know that every point in $\mathcal{V}(H_{G,C_i})$ has the form $p_1 = (\omega^{i+1}, \omega^{i+2}, \ldots, \omega^{i+k})$ (or possibly $p_2 = (\omega^{i+k}, \omega^{i+k-1}, \ldots, \omega^{i+1})$ if G is undirected), again indexing the vertices of G so that $C_i = \{v_1, v_2, \ldots, v_n\}$. Further, as every point in

$\bigcup_C \mathcal{V}(H_{G,C})$ is in $\mathcal{V}(H_{G,C_i})$ for some i, we can use the indexing appropriate to C_i to write any point in $\bigcup_C \mathcal{V}(H_{G,C})$ in the form of p_1 (p_1 or p_2 if G is undirected).

From our proof of the first part of Proposition 5.3, we see that points of the form of p_1 (p_1 or p_2 if G is undirected) are in $\mathcal{V}(H_G)$, so

$$\bigcup_C \mathcal{V}(H_{G,C}) \subseteq \mathcal{V}(H_G).$$

Thus, the equality holds. □

We will now make a brief detour to present a few properties about intersecting ideals.

Theorem 5.22. *(4.3 Theorem 15 in* [2]*) If I and J are ideals in $\mathbb{K}[x_1, \ldots, x_n]$, then $\mathcal{V}(I \cap J) = \mathcal{V}(I) \cup \mathcal{V}(J)$.*

Theorem 5.23. *(4.3 Proposition 16 in* [2]*) If I and J are ideals in $\mathbb{K}[x_1, \ldots, x_n]$, then $\sqrt{I \cap J} = \sqrt{I} \cap \sqrt{J}$.*

Theorem 5.24. *(The Strong Nullstellensatz, 4.2 Theorem 6 in* [2]*) Let $\mathbb{K}$ be an algebraically closed field. If I is an ideal in $\mathbb{K}[x_1, \ldots, x_n]$, then*

$$I(\mathcal{V}(I)) = \sqrt{I}.$$

From this, we have enough information to prove the decomposition theorem that is the culmination of Section 3 of [4].

Theorem 5.25. *Let G be a connected directed graph on n vertices. Then*

$$H_G = \bigcap_C H_{G,C}$$

where the intersection is over all Hamiltonian cycles C in G.

Proof. Recall that

$$H_G = \{x_i^n - 1 = 0, \prod_{j \in \delta^+(i)} (\omega x_i - x_j) = 0 \mid 1 \leq i \leq n\}$$

where $\delta^+(i)$ denotes those vertices j where there is an arc going from i to j in G.

Since H_G contains a square-free single variable polynomial in each variable (the generator of the form $x_i^n - 1$), by Proposition 2.7 of [3], H_G is radical. Then by Theorem 5.24 we see that

$$H_G = I(\mathcal{V}(H_G))$$

and by Lemma 5.21 this implies

$$H_G = I(\bigcup_C \mathcal{V}(H_{G,C}))$$

for Hamiltonian cycles C in G. By Theorem 5.22, we can then write the union of varieties as a variety of intersections:

$$H_G = I(\mathcal{V}(\bigcap_C H_{G,C})).$$

From Lemma 5.19 we know that $H_{G,C}$ is radical, so by Theorem 5.23 $\bigcap_C H_{G,C}$ is radical as well. Therefore by Theorem 5.24

$$I(\mathcal{V}(\bigcap_C H_{G,C})) = \bigcap_C H_{G,C}.$$

Then

$$H_G = \bigcap_C H_{G,C}$$

as desired. □

Definition 5.26. If a graph has only one Hamiltonian cycle, (assuming that we regard cycles as the same if we can write them in the same way by choosing a different starting vertex or, if the graph is undirected, a direction) we say the graph is **uniquely Hamiltonian**.

This means that for a uniquely Hamiltonian directed graph there are n possible ways to write down a cycle of length n, and for a uniquely Hamiltonian undirected graph there are $2n$ possible ways to write down a cycle of length n.

Corollary 5.27. *A graph G is uniquely Hamiltonian if and only if the Hamiltonian ideal H_G is of the form $H_{G,C}$ for some length n cycle C.*

Proof. If G is uniquely Hamiltonian, then the n or $2n$ ways of writing down the Hamiltonian cycle all correspond to the same ideal $H_{G,C}$. Then there is only one element in the intersection of Hamiltonian Cycle Ideals, so

$$H_G = \bigcap_C H_{G,C} = H_{G,C}.$$

If $H_G = H_{G,C}$ for some length n cycle C, then $\bigcap_C H_{G,C} = H_{G,C}$, so there can only be one distinct ideal in the intersection. Since we took the intersection over all Hamiltonian cycles, it follows that there is only one Hamiltonian Cycle Ideal. Then G must be uniquely Hamiltonian.

□

From this corollary, we can check if a given graph is uniquely Hamiltonian using the following algorithm:

(1) Compute a reduced Gröbner basis of H_G.

(2) Check that this basis has the form of an ideal $H_{G,C}$.

While this is a short algorithm, finding Gröbner bases is computationally intensive, and takes longer than polynomial time (a measure of complexity) to complete.

Finally, we will look at some examples of Theorem 5.25.

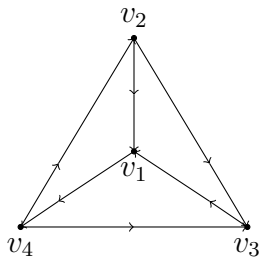

FIGURE 20. A directed K_4 graph.

Example 5.28. Let G be the directed K_4 graph in Figure 20, and let ω be a primitive 4th root of unity. The ideal $H_G \subseteq \mathbb{K}[x_1, x_2, x_3, x_4]$ is generated by the polynomials

$$\{x_i^4 - 1 | 1 \leq i \leq 4\} \cup \{(\omega x_1 - x_2), (\omega x_2 - x_3)(\omega x_2 - x_4), \\ (\omega x_3 - x_1)(\omega x_3 - x_4), (\omega x_4 - x_1)\}.$$

Given the ordering $x_1 > x_2 > x_3 > x_4$,

$$\{x_4^4 - 1, x_1 - \omega x_4, x_2 - \omega^2 x_4, x_3 - \omega^3 x_4\}$$

is a reduced Gröbner basis of H_G, as well as a generating set for the Hamiltonian cycle ideal $H_{G,C}$ with $C = \{(v_1, v_2), (v_2, v_3), (v_3, v_4), (v_4, v_1)\}$.

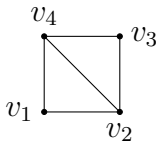

FIGURE 21. An undirected graph of order four.

Example 5.29. Let G be the undirected graph in Figure 21, and again let ω be a primitive 4th root of unity. The ideal $H_G \subseteq \mathbb{K}[x_1, x_2, x_3, x_4]$ is generated by the polynomials

$$\{x_i^4 - 1 | 1 \leq i \leq 4\} \cup \{(\omega x_1 - x_2)(\omega x_1 - x_4),$$
$$(\omega x_2 - x_1)(\omega x_2 - x_3)(\omega x_2 - x_4), (\omega x_3 - x_2)(\omega x_3 - x_4),$$
$$(\omega x_4 - x_1)(\omega x_4 - x_2)(\omega x_4 - x_3)\}.$$

We see that G has eight directed cycles:

$$C_1 = \{(v_1, v_2), (v_2, v_3), (v_3, v_4), (v_4, v_1)\},$$
$$C_2 = \{(v_2, v_3), (v_3, v_4), (v_4, v_1), (v_1, v_2)\},$$
$$C_3 = \{(v_3, v_4), (v_4, v_1), (v_1, v_2), (v_2, v_3)\},$$
$$C_4 = \{(v_4, v_1), (v_1, v_2), (v_2, v_3), (v_3, v_4)\},$$

$$C_5 = \{(v_1, v_4), (v_4, v_3), (v_3, v_2), (v_2, v_1)\},$$

$$C_6 = \{(v_4, v_3), (v_3, v_2), (v_2, v_1), (v_1, v_4)\},$$

$$C_7 = \{(v_3, v_2), (v_2, v_1), (v_1, v_4), (v_4, v_3)\},$$

$$C_8 = \{(v_2, v_1), (v_1, v_4), (v_4, v_3), (v_3, v_2)\}.$$

This means that G is uniquely Hamiltonian.

Given the ordering $x_1 > x_2 > x_3 > x_4$,

$$\{x_4^4 - 1, x_1 + x_3 + \frac{\omega^2 - 1}{\omega^3 - \omega} x_4, x_2 + \frac{\omega^{-1} - \omega}{\omega^3 - \omega} x_4, (x_3 - \omega x_4)(x_3 - \omega^{-1} x_4)\}$$

is a reduced Gröbner basis of H_G, as well as a generating set for the Hamiltonian cycle ideal H_{G,C_1}.

6. Conclusion and Further Questions

From the examples we have explored, it is clear that the polynomial method is a powerful tool in analysis of graphs. The proofs that are entirely based in graph theory, such as that of Proposition 3.22, and those that are more algebraic, such as that of Lemma 3.16, complement one another, and provide a range of options for mathematicians exploring a graph's properties.

Further research could consider what graphs have linear Nullstellensatz certificates of non-4-colorability, or examine the degree of Nullstellensatz certificates of non-3-colorability for planar graphs, or perhaps polytopes.

It would be satisfying to find a graph-theoretic proof of Corollary 3.18, though initial attempts led to no significant insight on how this might be done.

We have verified that the graph in Figure 10 has the least number of vertices of any graph that has a linear Nullstellensatz certificate of non-3-colorability but does not contain an odd wheel. However, we have not rigorously checked that it has the least number of edges of any graph of this type, so the question of whether it is truly the smallest remains open.

References

[1] David Allen Bayer, *The Division Algorithm and the Hilbert Scheme*, ProQuest LLC, Ann Arbor, MI, 1982. Thesis (Ph.D.)–Harvard University. MR2632095

[2] David Cox, John Little, and Donal O'Shea, *Ideals, varieties, and algorithms*, 3rd ed., Undergraduate Texts in Mathematics, Springer, New York, 2007. An introduction to computational algebraic geometry and commutative algebra. MR2290010 (2007h:13036)

[3] David A. Cox, John Little, and Donal O'Shea, *Using algebraic geometry*, 2nd ed., Graduate Texts in Mathematics, vol. 185, Springer, New York, 2005. MR2122859

[4] Jesús A. De Loera, Christopher J. Hillar, Peter N. Malkin, and Mohamed Omar, *Recognizing graph theoretic properties with polynomial ideals*, Electron. J. Combin. **17** (2010), no. 1, Research Paper 114, 26. MR2679568

[5] Jesús A. De Loera, Jon Lee, Peter N. Malkin, and Susan Margulies, *Hilbert's Nullstellensatz and an algorithm for proving combinatorial infeasibility*, ISSAC 2008, ACM, New York, 2008, pp. 197–206, DOI 10.1145/1390768.1390797. MR2513506

[6] Joseph Gallian, *Contemporary Abstract Algebra*, 7th ed., Houghton Mifflin, 2010.

[7] John M. Harris, Jeffry L. Hirst, and Michael J. Mossinghoff, *Combinatorics and graph theory*, 2nd ed., Undergraduate Texts in Mathematics, Springer, New York, 2008. MR2440898

[8] Serge Lang, *Algebra*, 3rd ed., Graduate Texts in Mathematics, vol. 211, Springer-Verlag, New York, 2002. MR1878556

[9] Bo Li, Benjamin Lowenstein, and Mohamed Omar, *Low Degree Nullstellensatz Certificates for 3-Colorability*, ArXiv e-prints (2015), available at `arXiv.1503.04680`.