SJÄLVSTÄNDIGA ARBETEN I MATEMATIK

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET

An algebraic approach to the problem of graph isomorpism

av

Alex Loiko

2014 - No 6

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET, 106 91 STOCKHOLM

An algebraic approach to the problem of graph isomorpism

Alex Loiko

Självständigt arbete i matematik 15 högskolepoäng, Grundnivå

Handledare: Boris Shapiro

2014

An algebraic approach to the problem of graph isomorpism

Alex Loiko

February 18, 2014

Acknowledgements

I would like to thank my advisor Boris Shapiro for introducing me to a topic that suited my interests and background so well and for his help with this project. I would also like to thank Torbjörn Tambour for his insightful comments.

Abstract

This work discusses invariant theory and its application on a particular question, *the graph isomorphism problem.* We develop the commutative algebra theory required to prove that the methods of invariant theory apply to graph isomorphism, implement several algorithms in `Mathematica` for solving graph isomorphism between complex-weighted graphs and analyze its complexity. Along the way, we discuss representation theory, group theory, and various algebraic methods.

1 Introduction

1.1 Background and overview

The problem of whether two given graphs are isomorphic is a well-known difficult problem. From the point of view of complexity theory, it belongs to the class NP. It is believed not to be NP-complete [Sch87] and has led to defining and studying the class **GI** of all languages with polynomial reduction to the graph isomorphism problem.

The graphs we will study are finite, the edges are undirected and weighted by complex numbers. The problem of undirected weighted graph isomorphism seems to be much more general than the ordinary undirected version, but it turns out that undirected complex-weighted graph isomorphism is **GI-complete**, meaning that there is polynomial reductions going in both directions between the more general and the undirected unlabeled versions.

Definition 1. *Let $V = \{1, 2, \ldots, n\}$ and let $\binom{V}{2}$ be the set of subsets containing 2 elements of V. Let G_1, G_2 be two complete undirected graphs with vertex set V and edge set $\binom{V}{2}$ and weight functions $w_i : \binom{V}{2} \longrightarrow \mathbb{C}$, $i = 1, 2$.*

The graphs G_1, G_2 are said to be ***isomorphic****, if there exists a permutation $\sigma \in S_n$ such that for each edge pair $\{a, b\}$,*

$$w_1(\{a, b\}) = w_2(\{\sigma(a), \sigma(b)\})$$

1.1.1 Reduction

Theorem 1. *There is a reduction polynomial in the size of input from the problem of edge-weighted undirected graph isomorphism to unweighted undirected graph isomorphism.*

Proof. Let G_1, G_2 be two edge-weighted graphs and let k be the total number of different edge-weights in both graphs. Let $|E_i|$ be the number of edges in graph G_i. We will abbreviate $|E|$ for both $|E_1|$ and $|E_2|$ because a necessary condition for isomorphy is an equal number of edges. Instead of weights we will color the edges in k colors. For edge (x, y) of color i, replace it with the structure below, x adjacent to a new vertex x' which is adjacent to $i + |E|$ (adding $|E|$ to ensure uniqueness) new vertices that are adjacent to a new vertex y' that is adjacent to y. This makes it only possible for equally marked edges to be identified. This reduction is obviously polynomial and the new graphs have $O(V+E^2)$ vertices.

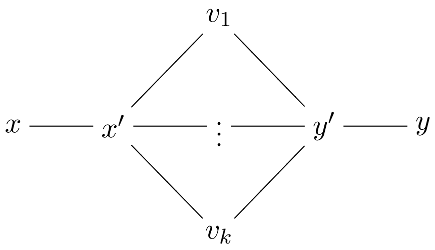

□

Invariant theory for the weighted GI problem

It turns out that studying the more general problem of weighted graph isomorphism makes it possible to use techniques from *invariant theory* - a branch of algebra using results from group theory, commutative algebra, representation theory and algebraic geometry.

Complexity

One asymptotically fast algorithm solves graph isomorphism in $2^{O(\sqrt{n \log n})}$ [BL83], and there are heuristics that are efficient for almost all pairs of graphs.

The time complexity of the first algorithm that we will develop using invariant theory is $\Omega\left(\left(\binom{\binom{n}{2}+n!}{\binom{n}{2}}\right)\right)$. The time complexity of the improved version is difficult to give an upper bound for, but it does $O\left(\left(\binom{\binom{n}{2}+\binom{n}{2}^2}{\binom{n}{2}}\right)\right)$ many Gröbner basis computations in the worst case to generate a separating set. In practice, the number of Gröbner basis calculations is much lower. We refer to **section 3.2** for the implementation of the first version, **section 4** for for discussion and analysis, **theorem 8** for the second version and **section 6.4** for its analysis.

1.2 Introducing the algebraic approach

1.2.1 Graph counting and comparison

We consider the main problem, namely, the one of counting and comparing graphs on n vertices.

Let S' be the set of all labeled graphs over n nodes with edges without weights. We consider the action of the symmetric group S_n on S' by permuting the vertices. If a graph $G \in S'$ has an edge (i,j) and $\sigma \in S_n$, the graph $\sigma \cdot G$ has an edge $(\sigma(i), \sigma(j))$.

It turns out that this approach is not well suited for weighted or directed graphs or any kind of numerical computation. The problem lies in the set S' that contains all possible $2^{\binom{n}{2}}$ labeled graphs.

1.2.2 Reducing the size of S'

For our purposes, it turns out to be better to see an unweighted undirected graph as a sum of its edges, e.g. $G = \{1,2\} + \{2,3\} + \{3,1\}$ is a triangle with edges $1 \longleftrightarrow 2, 2 \longleftrightarrow 3, 3 \longleftrightarrow 1$. We let the symmetric group S_n act on this set of edges (represented by unordered pairs $\{i,j\}$) by

$$\sigma\{i,j\} = \{\sigma(i), \sigma(j)\}.$$

Then the action of σ on G is essentially the same, we get a new graph with permuted labels. The advantage of acting on edges instead of on whole graphs is that the size of the set S is reduced from $2^{\binom{n}{2}}$ (number of possible undirected labeled graphs) to only $\binom{n}{2}$ (the number of unordered pairs of numbers $1, 2, \ldots, n$)

This approach has another advantage: weighted graphs can be represented as sums $G = \sum_{\{i,j\}} c_{\{i,j\}}\{i,j\}$ over the edges and directed graphs can be thought of as weighted graphs where the weight contains information about the direction of the edge.

1.2.3 Defining the S_n-representation

Given an action of a finite group G on a finite set S, there is a natural way to construct a G-module out of S. We introduce formal variables for each element $\mathbf{s} \in S$, and define M to be the $|S|$-dimensional vector space

$$M = \bigoplus_{\mathbf{s}\in S} \mathbb{C}\mathbf{s}$$

Then we can define a $G-$representation on M from the action on S by

$$g \cdot \sum_{\mathbf{s}\in S} a_{\mathbf{s}}\mathbf{s} = \sum_{\mathbf{s}\in S} a_{\mathbf{s}} g \cdot \mathbf{s}$$

(This is just linear extension of the inverse G-action on the basis vectors).

This turns out to be a G-module. We use this construction on S_n and S from **section 1.2.2** to turn $\mathbb{C}S$ into a S_n-module.

1.2.4 Applications

This algebraic approach leads to a general way of solving symmetry problems like GI and is believed [Thi00] to lead to progress in deciding **Ulam's reconstruction conjecture**, defined as follows:

Conjecture 1. *For a given graph G, we define $D(G)$ to be the multiset of all graphs obtained from G by deleting exactly one vertex. Then two graphs G, H are isomorphic if and only if $D(G) = D(H)$.*

1.3 Summary

In this work, we develop a solid and mostly self-contained foundation for invariant theory. The literature referenced for this work often omits important details, uses less elementary techniques or refers the proofs to external sources. Therefore, I believe that the subject has been made accessible for a wider audience.

We state the basic definitions of invariant theory and multilinear algebra in **sections 2, A**. We show using commutative algebra and algebraic geometry that invariant theory can be applied to the problem of graph isomorphism in **section 2.5**. We describe how a theorem of Emmy Noether can be used to apply the theoretical results to comparing graphs in practice in **section 3** and implement a complete algorithm for solving graph isomorphism in **section 3.2**. We indicate several ways the algorithm can be improved in **section 4**. In **section 5**, we define and show properties of Molien series, which is another theoretical tool that combined with a structure theorem of invariant rings described in **section 6** leads to a more efficient algorithm discussed in **section 6.3**. The algorithm is demonstrated for special cases in **section 6.3.1**.

2 Invariant ring

In **section 1.2** we reinterpreted combinatorial problems of counting and equivalence as algebra. This section develops the crucial theory behind the algorithm for deciding whether two elements belong to the same orbit.

2.1 Separating sets

Given a set S with an equivalence relation $\sim$ (the set of objects and symmetries), an arbitrary set D and a collection of functions $\mathcal{F} = \{f_i : S \longrightarrow D\}_i$ invariant on equivalence classes, we say that $\mathcal{F}$ is a *separating set* if $x \sim y$ if and only if $f(x) = f(y)$ for each $f \in \mathcal{F}$. Finding a finite separating set of computable functions solves the symmetry/isomorphism problem. Representation theory gives an algorithm for constructing a finite set of polynomial separating functions for any group acting on any finite set [FH91]. In what follows, we will develop the theory, give examples, implement this method in **Mathematica** and analyze complexity.

2.2 Definitions

Let V be a $\mathbb{C}[G]$-module. Then V^* is also a $\mathbb{C}[G]$-module under [1]

$$(g \bullet f)(v) = f(g^{-1}(v)) \iff g \bullet f = f \circ g^{-1}$$

The symmetric tensor algebra $\bigodot V^*$ is a polynomial ring by **Theorem 10** below. We let each generator $\bigotimes_{i=1}^{n} \widehat{v_i}$ act on $v \in V$ by

$$\left(\bigotimes_{i=1}^{n} \widehat{v_i} \right) \bullet v = \prod_{i=1}^{n} \widehat{v_i}(v)$$

and extend this action linearly over the whole $\bigodot V^*$.

2.2.1 Interpretation of $\bigodot V^*$

$\bigodot V^*$ can be interpreted as the ring of all polynomial functions[2] on V. As an example, let $V = \mathbb{C}\mathbf{x} \oplus \mathbb{C}\mathbf{y}$ be a 2-dimensional vector space over $\mathbb{C}$ with basis vectors $\mathbf{x}, \mathbf{y}$. Then the dual space V^* has a dual basis $\hat{x}, \hat{y}$ and by **Theorem 10**, V^* is isomorphic to $\mathbb{C}[x, y]$ as an algebra under concatenation of tensor products. Let $p(x, y) \in \mathbb{C}[x, y]$, $p(x, y) = x^2 y + 2y$. Then p acts on $v = a\mathbf{x} + b\mathbf{y} \in V$ by $p(x, y) \bullet v = a^2 b + 2b$

2.2.2 G-action on $\bigodot V^*$

Let V a vector space and G be a group acting linearly on V. Following the discussion in **Section 2.2.1**, every $f \in \bigodot V^*$ is a polynomial function $V \longrightarrow \mathbb{C}$. We define $g \bullet f = f \circ g^{-1}$ for $f \in \bigodot V$ and $g \in G$. This defines a group action because a polynomial function composed with a linear operator on V remains polynomial.

2.3 Relation to separating sets and equality on orbits

Recall the discussion of *separating sets* in **Section 2.1** We have constructed a large ring of all polynomial functions on V. We also have a group G acting on V partitioning the vector space into orbits. To construct a separating set, we need a set $\mathcal{F}$ of functions constant on the orbits. It turns out that for finite groups G and finite-dimensional $\mathbb{C}$-vector spaces V, it is enough to let $\mathcal{F}$ be the set of *all* polynomial functions constant on the orbits, which will be proven in **Theorem 2**.

[1] The reason for g^{-1} instead of g is that acting by g may lead to non-representations that do not satisfy $gh = g \circ h$ fon non-abelian groups G. Indeed, if the action was defined by $g \bullet f = f \circ g$, we would have $(gh) \bullet f = f \circ (gh) = f \circ g \circ h = h \bullet (f \circ g) = h \bullet (g \bullet f) = (hg) \bullet f$

[2] By which we mean all functions that map $v = a_1 \mathbf{e_1} + \ldots + a_n \mathbf{e_n}$ on an element in $\mathbb{C}[x_1, \ldots, x_n]$ and then evaluate in $(a_1, \ldots, a_n)$.

We say that a polynomial $p \in \bigodot V^*$ is **invariant** if $p = g \bullet p$ for all $g \in G$. This condition guarantees that p is invariant on orbits, since $p \bullet v = (g^{-1} \bullet p)(v) = p \circ g(v) = p(g \bullet v)$.

The set of all invariant polynomials is closed under $+$ and $\cdot$ and therefore is a subring of $\bigodot V^*$. We call this ring the **invariant ring** of G and denote it by $\bigodot V^{*G}$.

For finite groups G, the invariant ring is non-empty and does not only contain constants because for each $p \in \bigodot V^*$,

$$h \bullet \left(\sum_{g \in G} g \bullet p \right) = \sum_{g \in G} (hg) \bullet p = \sum_{g \in G} g \bullet p$$

and thus $\sum_{g \in G} g \bullet p \in \bigodot V^{*G}$.

In later sections, we will show that $\bigodot V^{*G}$ is a separating set and how to compute a finite generating set of $\bigodot V^{*G}$

2.4 Reynolds operator

The Reynolds operator $R^G : \bigodot V^* \longrightarrow \bigodot V^{*G}$ is defined by

$$R^G(p) = \frac{1}{|G|} \sum_{g \in G} g \bullet p.$$

We alraedy showed that $R^G(p) \in \bigodot V^{*G}$ for all $p \in \bigodot V^*$ in **Section 2.3**

Varieties - algebraic geometry

Recall that a *variety* of an ideal $I \subset \mathbb{C}[x_1, \dots, x_n]$ is the set of common zeros of I in V,

$$\mathcal{V}(I) = \{v \in V | \forall f \in V : f(v) = 0\}$$

Hilberts Nullstellensatz states that $\mathcal{V}(I)$ is never empty,

$$\mathcal{V}(I) = \emptyset \Longleftrightarrow I = \mathbb{C}[x_1, \dots, x_n]$$

The following properties follows from the definitions:

- Finite union:
$$\mathcal{V}(I) \cup \mathcal{V}(J) = \mathcal{V}(I \cap J)$$

- $\emptyset$ and V:
$$\emptyset = \mathcal{V}(\mathbb{C}[x_1, \dots, x_n]$$

-
$$V = \mathcal{V}(\{0\})$$

- Arbitrary intersection:

$$\bigcap_\alpha \mathcal{V}(I_\alpha) = \mathcal{V}(\sum_\alpha I_\alpha)$$

- Single points are varieties:

$$\{(a_1, \ldots, a_n)\} = \mathcal{V}(x_1 - a_1, \ldots, x_n - a_n)$$

  This implies that finite sets are varieties.

This shows that the sets $\mathcal{V}(I)$ define an topology on V called the *Zariski* topology where the $\mathcal{V}(I)$ are assumed as closed. The last property implies that it is at least as fine as the finite complement topology. We define $r(\Omega)$ for $\Omega \subset V$ as

$$r(\Omega) = \{f \in \mathbb{C}[x_1, \ldots, x_n] | \forall \boldsymbol{x} \in \Omega : \; f(\boldsymbol{x}) = 0\}$$

the ideal of all functions that are 0 on Ω.

The coordinate ring

If W is a variety in V, the **coordinate ring** of W is defined as

$$\mathbb{C}[x_1, \ldots, x_n] \Big/ r(W) .$$

It can be identified with the algebra of **all polynomial functions** on W: Let A be all polynomial functions on W and let ϕ be the map that restricts a polynomial $p \in \mathbb{C}[x_1, \ldots, x_n]$ to W. The kernel of ϕ is all polynomials that are 0 on W, which is $r(W)$.

2.5 The separation property

The following theorem is the main theoretical result behind our use or invariant theory to distinguish combinatorial objects.

Theorem 2. *Let V be a finite-dimensional vector space over $\mathbb{C}$ and G a finite group acting on V. Then $\bigodot V^{*G}$ separates the orbits of V, that is, if $v_1, v_2 \in V$ are in different orbits, there is $p \in \bigodot V^{*G}$ such that $p(v_1) \neq p(v_2)$.*

Proof. The orbit $orb_G(v)$ of v under G-action contains $\leq |G|$ elements and therefore is finite. Under the Zariski topology finite sets in V are closed. Closed sets are by definition varieties. Let $I_{v_1} = r(orb(v_1)), I_{v_2} = r(orb(v_2))$. Then by the variety properties,

$$\mathcal{V}(I_{v_1} + I_{v_2}) = orb(v_1) \cap orb(v_2) = \emptyset$$

because v_1 and v_2 were assumed to lie in different orbits.

By the Nullstellensatz, we have $I_{v_1} + I_{v_2} = \bigodot V^*$. It follows that $1 \in I_{v_1} + I_{v_2} \implies 1 = a + b$ with $a \in I_{v_1}$ and $b \in I_{v_2}$. Let $A = R^G(a), B = R^G(b)$ where R^G is the Reynolds operator. Then $A + B = R^G(a+b) = R^G(1) = 1$. We have $(g \bullet a)(v_1) = a(g^{-1}(v_1)) = 0$ because $g^{-1}(v_1)$ is a common zero of I_{v_1} and it follows that $A(v_1) = \frac{1}{|G|} \sum_{g \in G} 0 = 0$. Likewise, $B(v_2) = 0$. But the sum $(A + B)(v_2) = 1 = A(v_2) + B(v_2) = A(v_2)$. Now we get $A(v_2) = 1$, so A separates v_1 and v_2. □

3 Computations

In **Section 2.3** we stated that the invariant ring is a separating set which was proven in **Theorem 2**. But it is still not clear how knowledge that a certain infinite set of invariant functions is separating can be used to construct an algorithm for determining graph isomorphism.

We will compute a finite set of generators of $\bigodot V^{*G}$ for arbitrary finite groups G and finite-dimentional complex vector spaces V.

This set will be shown to constitute a finite separating set for graphs under isomorphism.

3.1 Finite set of generators

The following theorem is due to Emmy Noether and is one of the two main computational results. It directly leads to an algorithm implemented in **section 3.2**. The second important computational result is **theorem 6** discussed in **section 6**. It leads to a major improvement of the algorithm in **section 3.2**, discussed in **section 6.3**.

Theorem 3. *Let G be a finite group and V a n-dimensional complex vector space.*

Then the invariant ring $\bigodot V^{*G}$ is generated by the set

$$\{R_G(x_1^{a_1} \cdot \ldots \cdot x_n^{a_n}) | a_1 + \ldots + a_n \leq |G|, a_i \geq 0\}$$

Proof. The strategy of the proof will be to use the fundamental theorem of symmetric polynomials to rewrite a symmetric polynomial with degree larger than $|G|$ in $|G|$ variables to a polynomial in the $|G|$ symmetric power sum polynomials.

We introduce some notation:

$\mathbf{x}^{\boldsymbol{\alpha}}$ with $\mathbf{x} = (x_1, \ldots, x_n)$ and $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ stands for $x_1^{\alpha_1} \ldots x_n^{\alpha_n}$

As an example, let $V = \mathbb{C}\mathbf{e_x} \oplus \mathbb{C}\mathbf{e_y}$ be a 2-dimensional vector space, on which a group element g^{-1} acts by $g^{-1} \bullet (a\mathbf{e_x} + b\mathbf{e_y}) = (ag_{1,1} + bg_{1,2})\mathbf{e_x} + (ag_{2,1} + bg_{2,2})\mathbf{e_y}$. Then $p = x^2y$ is a polynomial function on V with $p\,(a\mathbf{e_x} + b\mathbf{e_y}) = a^2b$. After action of g, we have

$$\begin{aligned}&(g \bullet p)\,(a\mathbf{e_x} + b\mathbf{e_y})\\ &= p\left(g^{-1} \bullet (a\mathbf{e_x} + b\mathbf{e_y})\right)\\ &= p\,((ag_{1,1} + bg_{1,2})\mathbf{e_x} + (ag_{2,1} + bg_{2,2})\mathbf{e_y})\\ &= (ag_{1,1} + bg_{1,2})^2(ag_{2,1} + bg_{2,2})\end{aligned}$$

With the exponent notation, we have $x^2y = (x, y)^{(2,1)}$ and $g \bullet x^2y = ((xg_{1,1} + yg_{1,2}), (xg_{2,1} + yg_{2,2}))^{(2,1)}$

The Reynolds operator on $\mathbf{x}^{\boldsymbol{\alpha}}$ can with this notation be expressed as

$$R_G(\mathbf{x}^{\boldsymbol{\alpha}}) = \frac{1}{|G|}\sum_{g\in G}(g \bullet \mathbf{x})^{\boldsymbol{\alpha}}$$

U_g for $g \in G$ is defined by letting g act on $\mathbf{x}$ resulting in a vector $(y_1, y_2, \ldots, y_n)$ as in the example above by

$$U_g = u_1y_1 + \ldots + u_ny_n$$

where $u_1, \ldots, u_n$ are formal independent variables (making this is an expression in $\mathbb{C}[x_1, \ldots, x_n, u_1, \ldots, u_n]$).

S_k we define by

$$S_k = \sum_{g\in G} U_g^k$$

$|\alpha|$ is defined as $\alpha_1 + \ldots + \alpha_n$

Let $a_1, \ldots, a_{|G|}$ be formal variables in $\mathbb{C}[a_1, \ldots, a_{|G|}]$. Then each symmetric polynomial $p(a_1, \ldots, a_{|G|})$ is a unique polynomial q in the elementary symmetric polynomials. By an alternative version of the theorem p is also a unique polynomial r in the first $|G|$ power sum symmetric polynomials. Thus each S_k can be expressed as a polynomial $S_k = T_k(S_1, \ldots, S_{|G|})$.

We expand each U_g^k as

$$U_g^k = \sum_{|\boldsymbol{\alpha}|=k}\binom{k}{\boldsymbol{\alpha}}\mathbf{u}^{\boldsymbol{\alpha}}(g \bullet \mathbf{x})^{\boldsymbol{\alpha}}$$

where $\mathbf{u} = (u_1, \ldots, u_n)$. Now we sum over $g \in G$,

$$\sum_{g \in G} U_g^k = \sum_{g \in G} \left(\sum_{|\boldsymbol{\alpha}|=k} \binom{k}{\boldsymbol{\alpha}} \mathbf{u}^{\boldsymbol{\alpha}} (g \bullet \mathbf{x})^{\boldsymbol{\alpha}} \right)$$
$$= \sum_{|\boldsymbol{\alpha}|=k} \binom{k}{\boldsymbol{\alpha}} |G| \mathbf{u}^{\boldsymbol{\alpha}} R_G(\mathbf{x}^{\boldsymbol{\alpha}})$$

If we consider the expressions to be polynomials in $u_1, \ldots, u_n$ with coefficients in $\mathbf{K}[x_1, \ldots, x_n]$, we get

$$\sum_{|\boldsymbol{\alpha}|=k} \left(\binom{k}{\boldsymbol{\alpha}} |G| R_G(\mathbf{x}^{\boldsymbol{\alpha}}) \right) \mathbf{u}^{\boldsymbol{\alpha}} = T_k(S_1, \ldots, S_{|G|})$$

and equating coefficients of $\mathbf{u}^{\boldsymbol{\alpha}}$ shows that the coefficient $\left(\binom{k}{\boldsymbol{\alpha}}|G|R_G(\mathbf{x}^{\boldsymbol{\alpha}})\right)$ of $\mathbf{u}^{\boldsymbol{\alpha}}$ on the left hand side is a polynomial in the coefficient $\mathbf{u}^{\boldsymbol{\alpha}}$ of the left hand side, which itself is a polynomial in $R_G(\mathbf{x}^{\boldsymbol{\beta}})$ for all $\beta \leq |G|$.

This shows that the quantities $R_G(\mathbf{x}^{\boldsymbol{\alpha}})$ are elements of the ring generated by all $R_G(\mathbf{x}^{\boldsymbol{\beta}})$ with $|\boldsymbol{\beta}| \leq |G|$.

It remains to show that R_G is surjective. Let $p \in \bigodot V^{*G}$. We compute $R_G(p) = \frac{1}{|G|} \sum_{g \in G} p = p$ and the surjection property follows.

As there are $\binom{n+|G|}{n}$ values of $\boldsymbol{\alpha}$ with $|\boldsymbol{\alpha}| \leq |G|$, this theorem proves the bound $\binom{n+|G|}{n}$ for the number of generators of the invariant ring. For the graph case, S_n has $n!$ elements and V has dimension $\binom{n}{2}$. Emmy Noether's theorem applied to this particular case gives the number of basis elements

$$\binom{\binom{n}{2} + n!}{\binom{n}{2}}$$

For $n = 4$, the expression is 593775, for $n = 5$ it is 266401260897200. □

Based on Noether's theorem and **Theorem 2**, we can now construct a simple but rather inefficient algorithm that determines graph isomorphy on graphs weighted with complex numbers. For each of the $\binom{\binom{n}{2}+n!}{\binom{n}{2}}$ polynomials in the generating set of $\bigodot V^{*S_n}$, we evaluate the polynomials on the two graphs and compare the results. If all values coincide, the graphs are isomorphic, otherwise they are not.

3.2 Mathematica implementation

We implement this algorithm in `Mathematica`:

`SnAction[perm, poly, var, n]` computes the action of the permutation `perm`$^{-1}$ on the polynomial `poly` in the $\binom{n}{2}$ variables $var_{1,2}, \dots, var_{n-1,n}$

```
SnAction[perm_, poly_, var_, n_] :=
 poly /.
  Table[var[t[[1]], t[[2]]] ->
    If[perm[[t[[1]]]] < perm[[t[[2]]]],
     var[perm[[t[[1]]]], perm[[t[[2]]]]],
     var[perm[[t[[2]]]], perm[[t[[1]]]]]], {t, Subsets[Range[n], {2}]}]
```

`variables[var, n]` is the $\binom{n}{2}$ variables $var_{1,2}, \dots, var_{n-1,n}$:

```
variables[var_, n_] :=
Table[var[t[[1]], t[[2]]],
  {t, Subsets[Range[n], {2}]}]
```

`reynolds[poly, var,n]` computes the Reynolds operator R_{S_n}(`poly`)

```
reynolds[poly_, var_, n_] :=
 1/Factorial[n] Total[
    Table[
     SnAction[perm, poly, var, n], {perm,
      Permutations[Table[i, {i, n}]]}
     ]
    ] // Expand
```

`expVectors[n]` computes the list of all $\binom{n}{2}$-dimensional tuples $\boldsymbol{\alpha}$ with $|\boldsymbol{\alpha}| \leq n!$

```
(* expVectors - computes exponent vectors for Noether's generator-bound theorem on the invariant ring. Uses FrobeniusSolve that returns all solutions (a1,a2,...,an) to m1*a1 + m2*a2+..+mn*an=d for input (m1,m2,...,mn) and d. Much more efficient that the previous version that generated all tuples (a1,...,an) with ai<Factorial[n] and selected the ones with sum<Factorial[n]*)

expVectors[n_] :=
 Flatten[
  Table[
   FrobeniusSolve[
    ConstantArray[1, Binomial[n, 2]],
     k],
   {k, Range[0, Factorial[n]]}],
  1]
(* Example *)
expVectors[3];
```

generators[var, n] computes the list of generators from **Theorem 3**

```
generators[var_, n_] := Table[
  reynolds[
   Times @@ MapThread[#1^#2 &,
     {variables[var, n],
      expVector}],
   var, n]
  ,
  {expVector, expVectors[n]}]
```

Finally, isomorphic[graph1, graph2] tests isomorphy

```
(* Evaluates poly on graph. 'graph' is expected to be a list of weights for each of the Binomial[n, 2]edges. *)

evaluateGraph[poly_, var_, n_, graph_] :=

 poly /. MapThread[#1 -> #2 &, {variables[var, n], graph}] // Expand

(* Example: *)
evaluateGraph[
 x[1, 2]^2 + x[2, 3] x[1, 3], x, 3, {0, a, b}]

(* graph1, graph2 is a list of weights, have to be same length. Solves the graph isomorphy problem with Noether's theorem and invariant theory! *)

isomorphic[graph1_, graph2_, n_] :=
 Module[{n1 = Length[graph1],
   n2 = Length[graph2]},
  If[n1 != n2, False,
   Fold[And, True,
    Table[
     evaluateGraph[poly, x, n, graph1] ===
      evaluateGraph[poly, x, n, graph2],
     {poly, generators[x, n]}
     ]
    ]
   ]
  ]
```

The implementation was tested to work correctly and near instantaneously for $n = 3$ but runs out of memory for $n = 4$.

4 Improving the algorithm

We will now focus on improving the algorithm from **Section 3**. The version implemented in `Mathematica` worked for graphs with up to 3 vertices and had a time complexity of $\Omega\left(\left(\!\!\binom{\binom{n}{2}+n!}{\binom{n}{2}}\!\!\right)\right)$.

We will demonstrate a simple improvement using only algebra: Let $\mathcal{G}$ be the set of generators from **Theorem 3**. We know $\mathcal{G}$ to be a generating set for $\bigodot V^{*S_n}$, that is $\mathbb{C}[\mathcal{G}] = \bigodot V^{*S_n}$. If we find a set $\mathcal{T}$ with $\mathbb{C}[\mathcal{T}] = \mathbb{C}[\mathcal{G}]$, then $\mathcal{T}$ also has to be a generating set and therefore a separating set for graphs.

As an example, consider the list $\mathcal{G}$ of the $\binom{\binom{3}{2}+3!}{\binom{3}{2}} = 84$ generators of $\bigodot V^{*S_3}$ generated by `generators[x, 3]` from **Section 3.2**. It contains, among 81 others, polynomials

$$
\begin{aligned}
x[1,2]x[2,3]x[1,3] &= R_{S_3}(x[1,2]x[2,3]x[1,3]) \\
1/3x[1,2] + 1/3x[1,3] + 1/3x[2,3] &= R_{S_3}(x[1,2]) \\
1/3x[1,2]x[2,3] + 1/3x[1,3]x[2,3] + 1/3x[1,3]x[2,3] &= R_{S_3}(x[1,3]x[2,3])
\end{aligned}
$$

and therefore all the elementary symmetric polynomials of 3 variables. This implies by the Fundamental theorem of symmetric polynomials that the ring of all symmetric polynomials on $x[1,2], x[2,3], x[1,3]$ is a subring of $\bigodot V^{*S_n}$. On the other hand, every polynomial in $\bigodot V^{*S_n}$ is symmetric, therefore

$$
\begin{aligned}
\bigodot V^{*S_3} = \mathbb{C}[\mathcal{G}] = \mathbb{C}\Big[&x[1,2]x[2,3]x[1,3], \\
&1/3x[1,2] + 1/3x[1,3] + 1/3x[2,3], \\
&1/3x[1,2] + 1/3x[1,3] + 1/3x[2,3]\Big]
\end{aligned}
$$

This argument shows that it is enough to evaluate two graphs of 3 vertices on the elementary symmetric polynomials and none others - not surprising as two weighted triangle graphs are isomorphic if and only if the multiset of edge weights are equal.

In **Theorem 13** of **section B**, we describe an algorithm based on the theory of Gröbner bases that given a list of polynomials $\mathcal{G}$ and another polynomial p, tests whether $p \in \mathbb{C}[\mathcal{G}]$. This algorithm can then be applied to reduce the size of the generating set. As is described below in **section B.3.1** the algorithm applied to the 84 generators from **section 3.2** results in only 3 generators

$$\left\{\frac{1}{3}x_{1,2} + \frac{1}{3}x_{1,3} + \frac{1}{3}x_{2,3}, \frac{1}{3}x_{1,2}^2 + \frac{1}{3}x_{1,3}^2 + \frac{1}{3}x_{2,3}^2, \frac{1}{3}x_{1,2}^3 + \frac{1}{3}x_{1,3}^3 + \frac{1}{3}x_{2,3}^3\right\}$$

In **section 5** we gain information about the invariant ring by decomposing it into a direct sum of degree-homogeneous components and defining the generating function for the sequence of dimensions of these homogeneous components. The technique is useful for faster computation of smaller generating sets and therefore for smaller sets of separating functions, which leads to the more efficient graph isomorphy algorithm described in **theorem 8**.

5 Molien series

A very useful tool for computing the invariant ring of a group G acting on a vector space V and its generators is the *Molien series* of the invariant ring. The results from this section are used together with commutative algebra theorems from **section 6** to develop and prove correctness of the algorithm in **theorem 8**.

Assume $\dim_{\mathbb{C}} V = \dim_{\mathbb{C}} V^* = n$. Then by **theorem 10**, $\bigodot V^* = \bigoplus_{d\geq 0}\left(\bigodot^d V^*\right) \cong \mathbb{C}[x_1,\ldots,x_n]$ is a free graded algebra of dimension n.

The subalgebra $\bigodot V^{*G}$ inherits the grading with

$$\bigodot V^{*^G} = \bigoplus_{d\geq 0}\left(\bigodot^d V^*\right)^G$$

The *Molien series* of $\bigodot V^{*G}$ are defined as the generating function of the sequence

$$\left\{\dim_{\mathbb{C}}\left(\bigodot^d V^*\right)^G\right\}_d^\infty$$

of dimensions of the components $\left(\bigodot^d V^*\right)^G$,

$$H_{\bigodot V^{*G}}(T) =_{\text{def}} \sum_{d\geq 0} \dim_{\mathbb{C}}\left(\bigodot^d V^*\right)^G T^d \in \mathbb{C}\,[\![T]\!]$$

5.1 Computing Molien series

As will follow from theorems 1, 4, the Molien series of an invariant ring can be computed without finding a generating set, which allows for improved methods of computing generating and consequently constructing separating sets of functions.

5.1.1 Diagonalizability

Let G be a finite group acting on a finite-dimensional vector space V over the field $\mathbb{C}$ of complex numbers. Let $g \in G$ be an element of G that acts on V with invertible linear map $M_g \in GL(V)$. To be able to compute Molien series, we show the following lemma:

Lemma 1.

1. *The corresponding linear map $M_g \in GL(V)$ for the group action of $g \in G$ on a finite-dimensional complex vector space V is diagonalizable for finite $|G|$.*

2. *The eigenvalues of M_g lie on the unit circle and can be written as $\lambda = e^{\pi i r}$ for rational r e.g. the eigenvalues are roots of unity.*

Proof.

1. Let $\langle \bullet | \bullet \rangle$ be any scalar product on V. Then define $[\bullet | \bullet] : V \times V \longrightarrow \mathbb{C}$ by

$$[v_1 | v_2] = \frac{1}{|G|} \sum_{h \in G} \langle h \bullet v_1 | h \bullet v_2 \rangle$$

$[\bullet | \bullet]$ defines another scalar product that has the property of being *G-invariant*, meaning that $[g \bullet v_1 | g \bullet v_2] = [v_1 | v_2]$ for all $v_1, v_2 \in V$.

From elementary linear algebra, it follows that there is a basis $\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_n$ orthonormal with respect to $[\bullet | \bullet]$.

Consider the vectors $\boldsymbol{w_1} = M_g \boldsymbol{u}_1, \boldsymbol{w_2} = M_g \boldsymbol{u}_2, \ldots, M_g \boldsymbol{u}_n$. These vectors form a basis for V because $M_g \in GL(V)$ and $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$ were assumed to be a basis. We compute

$$\begin{aligned} [\boldsymbol{w_i} | \boldsymbol{w_j}] = [g \bullet \boldsymbol{u}_i | g \bullet \boldsymbol{u}_j] = \\ [\boldsymbol{u}_i | \boldsymbol{u}_j] = \delta_{i,j} \end{aligned}$$

where $\delta_{i,j}$ is the Kronecker δ. As seen, the vectors $\boldsymbol{w_1}, \ldots, \boldsymbol{w_n}$ form an orthonormal basis under $[\bullet | \bullet]$. In the basis $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$ the coordinates of the vectors $\boldsymbol{w_1}, \ldots, \boldsymbol{w_n}$ are the rows of M_g. We have showed that the rows of M_g form an orthonormal basis. This is equivalent to M_g unitary. The Spectral theorem implies that M_g is diagonalizable.

2. Let $d = \operatorname{ord}_G g$. Then $g^d = e$ and it follows from group action that $M_g^d = I_{\dim V}$. M_g is diagonalizable by **1.** It follows that $\lambda^d = 1$ for every eigenvalue λ of M_g, which implies that $\lambda = e^{2\pi i \frac{k}{d}}$ for some integer k.

□

5.2 Trace of the Reynolds operator

Recall the *Reynolds operator* from 2.4 defined by

$$R^G(p) = \frac{1}{|G|} \sum_{g \in G} g \bullet p$$

for finite groups G. We will be interested in its projection and grading preserving property:

Lemma 2.

1. *As a linear map $\bigodot V^* \longrightarrow \bigodot V^{*G}$, the Reynolds operator is a projection.*

2. *The Reynolds operator preserves degree of monomials. Let p me a homogeneous polynomial $p \in \bigodot V^*$. Then $R_G(p)$ homogeneous with the same total degree.*

Proof.

1. Let $p \in \bigodot V^*$. Then $R^G(p) \in \bigodot V^{*G}$. From the proof of **theorem 3**, we know that R^G is identity on $\bigodot V^{*G}$. We compute $R^G \circ R^G(p) = R^G(R^G(p)) = R^G(p)$

2. Let $\boldsymbol{x}^{\boldsymbol{\alpha}} \in \bigodot V^*$ be a monomial. Then as in the proof of **theorem 3**, $g \bullet \boldsymbol{x}^{\boldsymbol{\alpha}} = (g \bullet \boldsymbol{x})^{\boldsymbol{\alpha}}$ which is homogeneous. It follows that $R^G(\boldsymbol{x}^{\boldsymbol{\alpha}}) = \frac{1}{|G|}\sum_{g\in G}(g \bullet \boldsymbol{x})^{\boldsymbol{\alpha}}$ is a sum of homogeneous components with total degree $|\boldsymbol{\alpha}|$ and is therefore homogeneous with total degree $|\boldsymbol{\alpha}|$.

□

Lemma 2 implies that R^G can be restricted to homogeneous components,

$$R^G\left(\bigodot^d V^*\right) = \bigodot^d V^{*^G}$$

The following lemma relates the trace of R^G to the dimension of a graded component:

Lemma 3. *Let $R^G_d : \bigodot^d V^* \longrightarrow \bigodot^d V^*$ be the restriction of R^G to $\bigodot^d V^*$. Then*

$$\operatorname{Tr}(R^G_d) = \dim_{\mathbb{C}} \bigodot^d V^{*^G}$$

Proof. By **lemma 2**, R^G is a projection. Therefore, the restriction $R^G_d : \bigodot^d V^* \longrightarrow \bigodot^d V^*$ is also a projection on $\bigodot^d V^{*^G}$. By **theorem 9**, $\bigodot^d V^*$ is finite-dimensional with dimension $\binom{n+d-1}{d}$. Thus the trace is defined. Since R^G_d is a projection, there is a basis where the corresponding matrix of R^G_d has block structure

$$R^G_d \sim \begin{pmatrix} I_k & B \\ 0 & 0 \end{pmatrix}$$

where $k = \dim \operatorname{Im} R^G_d$ and B is some $k \times (\binom{n+d-1}{d} - k)$ matrix. Since R^G_d is a projection onto $\bigodot^d V^{*^G}$, $k = \operatorname{Im} R^G_d = \dim \bigodot^d V^{*^G}$.

We compute $\operatorname{Tr} R_d^G = \operatorname{Tr} \begin{pmatrix} I_k & B \\ 0 & 0 \end{pmatrix} = k$. $k = \dim {\bigodot^d V^*}^G$ and the proof is finished. □

The lemma leads to the following corollary concerning Molien series:

Corollary 1. *The Molien series of* ${\bigodot V^*}^G$ *can be computed by*

$$H_{\bigodot V^{*G}}(T) = \sum_{d \geq 0} (\operatorname{Tr} R_d^G) T^d \tag{1}$$

5.3 Traces of group elements

The aim of this section is to state and prove **theorem 4**, which will be done by expanding the definition of R^G in **corollary 1**.

To compute $\operatorname{Tr} R_d^G$, we use the definition of the Reynolds operator to rewrite R^G as $R^G(p) = \frac{1}{|G|} \sum_{g \in G} g \bullet p$. We saw in the proof of **theorem 3** that g-action preserves total degree. Therefore, we can define a restriction of g onto $\bigodot^d V^*$. Denote the linear map corresponding to this restriction as g_d. Then R_d^G can be expressed as

$$R_d^G = \frac{1}{|G|} \sum_{g \in G} g_d$$

By linearity of trace,

$$\operatorname{Tr} R_d^G = \frac{1}{|G|} \sum_{g \in G} \operatorname{Tr} g_d$$

We substitute into 1:

$$\begin{gathered} H_{\bigodot V^{*G}}(T) = \sum_{d \geq 0} (\operatorname{Tr} R_d^G) T^d = \\ \sum_{d \geq 0} \left(\frac{1}{|G|} \sum_{g \in G} \operatorname{Tr} g_d \right) T^d = \\ \frac{1}{|G|} \sum_{g \in G} \left(\sum_{d \geq 0} \operatorname{Tr} g_d T^d \right) \end{gathered} \tag{2}$$

The expression $\left(\sum_{d \geq 0} \operatorname{Tr} g_d T^d\right)$ will be simplified below using **Lemma 1**. The result is the following theorem:

Theorem 4. *A finite group G acts on a finite-dimensional vector space V by $M_g \in GL(V)$ for every $g \in G$. Then the Molien series $H_{\bigodot V^{*G}}(T)$ are given by the formula*

$$H_{\bigodot V^{*G}}(T) = \frac{1}{|G|} \sum_{g \in G} \frac{1}{\det(I_n - T M_g)}$$

Proof. By **Lemma 1**, $M_g : V \longrightarrow V$ has eigenvectors $\lambda_1, \ldots, \lambda_n$, for which there is a corresponding eigenbasis $\boldsymbol{v_1}, \boldsymbol{v_2}, \ldots, \boldsymbol{v_n}$. Consider the dual basis $\boldsymbol{f_1}, \boldsymbol{f_2}, \ldots, \boldsymbol{f_n}$ of V^*. Then by **theorem 9**, $\bigodot^d V^*$ is has a basis consisting of every monomial $\otimes_{i=1}^n f_i^{\alpha_i}$ for any choice of $|\boldsymbol{\alpha}| = \alpha_1 + \alpha_2 + \ldots + \alpha_n = d$ that by **theorem 10** corresponds to monomials $\boldsymbol{x}^{\boldsymbol{\alpha}}$. By the discussion in **theorem 3**, g acts on $\boldsymbol{x}^{\boldsymbol{\alpha}}$ by $g \bullet \boldsymbol{x}^{\boldsymbol{\alpha}} = \boldsymbol{\lambda}^{\boldsymbol{\alpha}} \boldsymbol{x}^{\boldsymbol{\alpha}}$ where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ and $\boldsymbol{\lambda}^{\boldsymbol{\alpha}} = \prod_{i=1}^n \lambda_i^{\alpha_i}$.

It follows that in the chosen monomial basis, g_d has an associated diagonal basis where every monomial $\boldsymbol{x}^{\boldsymbol{\alpha}}$ is an eigenvector and $\boldsymbol{\lambda}^{\boldsymbol{\alpha}}$ corresponding eigenvalue. Then $\operatorname{Tr} g_d$ is the sum of all eigenvalues which we compute as

$$\operatorname{Tr} g_d = \sum_{|\boldsymbol{\alpha}|=d} \boldsymbol{\lambda}^{\boldsymbol{\alpha}} =$$

$$\left[T^d\right] \left(\prod_{i=1}^n \frac{1}{1 - T\lambda_i} \right)$$

is the coefficient of T^d in the expansion of $\prod_{i=1}^n \frac{1}{1-T\lambda_i}$.

We conclude that

$$\left(\sum_{d \geq 0} \operatorname{Tr} g_d \right) = \prod_{i=1}^n \frac{1}{1 - T\lambda_i}$$

Let M_g be the matrix associated with g_d in the eigenbasis. The eigenvectors are $\lambda_1, \ldots, \lambda_n$, so $M_g = \operatorname{Diag}(\lambda_1, \ldots, \lambda_n)$. It follows that $I_n - TM_g = \operatorname{Diag}(1 - T\lambda_1, \ldots, 1 - T\lambda_n)$. We compute the determinant $\det(I_n - TM_g) = \prod_{i=1}^n (1 - T\lambda_i)$ and therefore simplify

$$\left(\sum_{d \geq 0} \operatorname{Tr} g_d \right) = \prod_{i=1}^n \frac{1}{1 - T\lambda_i} = \frac{1}{\det(I_n - TM_g)} \tag{3}$$

Simple substitution into 3 gives

$$H_{\bigodot V^{*G}}(T) = \sum_{d \geq 0} (\operatorname{Tr} R_d^G) T^d = \frac{1}{|G|} \sum_{g \in G} \left(\frac{1}{\det(I_n - TM_g)} \right)$$

□

5.4 Molien series of special rings

We will calculate the Molien series for invariant rings of permutation actions and for invariant rings generated by algebraically independent homogeneous polynomials.

5.4.1 Algebraically independent homogeneous polynomials

Lemma 4. *Let V be an n-dimensional vector space, G - a group acting on it as a subgroup of $GL(V)$. Assume $\bigodot V^{*G} \subseteq \bigodot V^*$ is generated by algebraically independent homogeneous polynomials $\theta_1, \dots, \theta_n$ with degrees $d_1, d_2, \dots, d_k$. Then the molien series of the invariant ring*

$$\bigodot V^{*^G} \cong \mathbb{C}[\theta_1, \dots, \theta_n] \subseteq \mathbb{C}[x_1, \dots, x_n]$$

can be computed by

$$H_{\bigodot V^{*G}}(T) = \frac{1}{(1-T^{d_1})(1-T^{d_2})\dots(1-T^{d_k})}$$

Proof. $\bigodot^d V^{*^G}$ has one spanning component for each product $\boldsymbol{\theta}^{\boldsymbol{\alpha}}$ with $|\boldsymbol{\alpha}| = d$. It follows that $\dim \bigodot^d V^{*^G}$ is the coefficient of T^d in

$$(1+T^{d_1}+T^{2d_1}+\dots)(1+T^{d_2}+T^{2d_2}+\dots)\dots(1+T^{d_k}+T^{2d_k}+\dots) = \frac{1}{(1-T^{d_1})(1-T^{d_2})\dots(1-T^{d_k})}$$

□

5.4.2 Example, graphs with $n = 3$

The computation in **section** 4 shows that $\bigodot V^{*S_3}$ is generated by the 3 symmetric power-sum polynomials. These polynomials are algebraically independent. We can therefore apply **lemma** 4:

$$H_{\bigodot V^{*S_3}}(T) = \frac{1}{(1-T)(1-T^2)(1-T^3)} = \\ 1 + T + 2T^2 + 3T^3 + 4T^4 + 5T^5 + 7T^6 + \dots$$

We will obtain another way to compute Molien series for this invariant ring from **lemma** 5

5.4.3 Molien series for permutation actions

Assume a finite group G acts on a finite set S. Then G acts linearly on $V = \mathbb{C}S$. The Molien series of the invariant ring $\bigodot V^{*G}$ of this action can be found with the following result:

Lemma 5. *A finite group G acts on a finite set X by permutation. Each $g \in G$ acts by a corresponding permutation $\pi_g \in S_X$. Let $cyc(g) = (a_1^g, a_2^g, \dots, a_n^g)$ be the exponent vector of the cycle type $1^{a_1^g} 2^{a_2^g} \dots n^{a_n^g}$ of π_g for each $g \in G$. Then the Molien series of the invariant ring $\bigodot \mathbb{C}X^{*G}$ can be found by*

$$H_{\bigodot \mathbb{C}X^{*G}}(T) = \frac{1}{|G|} \sum_{g \in G} \frac{1}{\prod_{i=1}^n (1-T^i)^{a_i^g}}$$

Proof. $g \in G$ acts on X with cycle type $1^{a_1^g} 2^{a_2^g} \dots n^{a_n^g}$. This means that there are a_i^g cycles of length i for each $i = 1, 2, \dots n$ and $\sum_{i=1}^n i \cdot a_i^g = n$.

We choose the standard basis $B = \{\boldsymbol{s} | s \in X\}$, and reorder the basis vectors in a way so that elements are grouped by cycles. Then the matrix M_g of g has the form

$$M_g = \oplus_{i=1}^n C_i^{\oplus a_i^g}$$

where each C_k is a matrix of a cyclic action on k elements, one way such matrices could look like is

$$C_2 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad C_3 = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

We recognize C_k as a matrix for the shift operator of the module of all polynomial remainders under division by $z^k - 1$, denoted X_{z^k-1} in [Fuh11]. According to **proposition 5.4 of [Fuh11]** the eigenvalues of the shift operator are the roots $\{e^{\frac{l}{k}2\pi i} | l = 0, 1, \dots (k-1)\}$ of $z^k - 1$.

If we instead choose an eigenbasis, C_k is similar to

$$C_k \sim \mathrm{Diag}(\omega, \omega^2, \dots, \omega^k)$$

where $\omega = e^{\frac{2\pi i}{k}}$. Then the determinant $\det(I_k - TC_k)$ is

$$\prod_{i=1}^k (1 - T\omega^k) = \prod_{i=1}^k T\left(\frac{1}{T} - \omega^k\right) =$$

$$T^k \left(\left(\frac{1}{T}\right)^k - 1\right) = 1 - T^k$$

and we can factor $\det(I_n - TM_g)$ as

$$\det(I_n - TM_g) = \prod_{i=1}^n (1 - T^k)^{a_i^g}$$

which substituted in **equation (4)** gives

$$\sum_{d \geq 0} \operatorname{Tr} g_d T^d = \frac{1}{\det(I_n - TM_g)} =$$

$$\prod_{i=1}^k \frac{1}{(1 - T^k)^{a_i^g}}$$

We substitute this into **equation (3)** and compute

$$H_{\bigodot \mathbb{C}X^{*G}}(T) = \sum_{d\geq 0} (\operatorname{Tr} R_d^G) T^d =$$

$$\frac{1}{|G|} \sum_{g\in G} \left(\prod_{i=1}^{k} \frac{1}{(1-T^k)^{a_i^g}} \right) =$$

$$\frac{1}{|G|} \sum_{g\in G} \frac{1}{\prod_{i=1}^{n} (1-T^i)^{a_i^g}}$$

□

5.4.4 Example, graphs with $n = 3, 4$

Again, we consider graphs with $n = 3$. The set X of edges upon which S_3 acts has 3 elements. There is 1 permutation with cycle type 1^3 that is the identity, 3 permutations with cycle type $1^1 2^1$ that flip two vertices and 2 permutations of type 3^1 that rotate the vertices. **Lemma 5** allows us to compute

$$H_{\bigodot \mathbb{C}X^{*G}} = \frac{1}{3!} \left(\frac{1}{(1-T)^3} + 3\frac{1}{(1-T)(1-T^2)} + 2\frac{1}{1-T^3} \right) =$$

$$\frac{1}{(1-T)(1-T^2)(1-T^3)}$$

6 Hironaka decomposition, efficient algorithm

In this section, we will state a strong structure theorem about a certain kind of rings and describe how it can be used to distinguish graphs. We refer to [PS08] for the proofs of **theorems 6**.

Definition 2. *Let R be a $\mathbb{C}$-algebra. Then the Krull dimension of R is defined as the maximal number of algebraically independent elements of R.*

Definition 3. *Let R be a graded ring with Krull dimension n. A set $\{\theta_1, \dots, \theta_n\} \subset R$ of algebraically independent homogeneous elements is a* ***homogeneous system of parameters*** *if R is finitely generated as a $\mathbb{C}[\theta_1, \dots, \theta_n]$-module.*

As an example of **definition 3**, we will prove that

Theorem 5. *Let X_n be the set of all undirected edges of n-vertex graphs under S_n acting on vertices. Then the graded $\mathbb{C}[x_{\{1,2\}}, \dots, x_{\{n-1,n\}}]$-subalgebra $\bigodot \mathbb{C}X_n^{*G}$ has Krull dimension $\binom{n}{2}$ and a homogeneous system of parameters.*

Proof. Applying the Reynolds operator R_{S_n} to $x_{\mathbf{\{1,2\}}}, x^2_{\mathbf{\{1,2\}}}, \ldots, x^{\binom{n}{2}}_{\mathbf{\{1,2\}}}$, we obtain the first $\binom{n}{2}$ elementary power-sum symmetric polynomials which are known to be algebraically independent. This implies that the Krull dimension of $\bigodot \mathbb{C}X_n^{*G}$ is at least $\binom{n}{2}$. But $\bigodot \mathbb{C}X_n^{*G}$ is a sub-algebra of $\mathbb{C}[x_{\mathbf{\{1,2\}}}, \ldots, x_{\mathbf{\{n-1,n\}}}]$ from which it follows that the Krull dimension of $\bigodot \mathbb{C}X_n^{*G}$ is at most $\binom{n}{2}$. □

Definition 4. *Let R be a graded $\mathbb{C}$-algebra. If for any homogeneous system of parameters $\theta_1, \ldots, \theta_n$ there exists a finite sequence of homogeneous $\eta_1, \ldots, \eta_t$ such that*

$$R = \bigoplus_{i=1}^{t} \eta_i \mathbb{C}[\theta_1, \ldots, \theta_n]$$

*then R is a **Cohen-Macaulay** algebra.*

Given a decomposition $\{\theta_1, \theta_2, \ldots, \theta_n, \eta_1, \ldots, \eta_t\}$ of R the elements $\theta_i, i = 1, 2, \ldots n$ are called *primary invariants* and the elements $\eta_i, i = 1, 2, \ldots, t$ are called the *secondary invariants.*

We refer to [PS08] for the proof of the following central theorem:

Theorem 6. *For a finite group G acting on a finite-dimensional vector space V, the invariant ring*

$$\bigodot V^{*G}$$

is Cohen-Macaulay.

The next theorem shows how Cohen-Macaulayness is related to Molien series and generators of invariant rings by comparing Molien series of

$$\bigoplus_{i=1}^{t} \eta_i \mathbb{C}[\theta_1, \ldots, \theta_n] \text{ and } \bigodot V^{*G}$$

Theorem 7. *Let G be a finite group acting on an n-dimensional vector space V. Assume that $\theta_1, \ldots, \theta_n$ is a homogeneous system of parameters of $\bigodot V^{*G}$. Let $d_i = \deg \theta_i$ for $i = 1 \ldots, n$. Then the number t and degrees of the corresponding homogeneous secondary invariants are the exponents $e_1, e_2, \ldots, e_t$ of the generating function*

$$H_{\bigodot V^{*G}}(T) \cdot \prod_{i=1}^{n} (1 - T^{d_i}) = T^{e_1} + T^{e_2} + \ldots + T^{e_t}$$

Proof. By Cohen-Macaulayness of $\bigodot V^{*G}$ there exists a set $\{\eta_1, \ldots, \eta_j\}$ of homogeneous secondary invariants. The Molien series of $\bigoplus_{i=1}^{t} \eta_i \mathbb{C}[\theta_1, \ldots, \theta_n]$

can be computed by **lemma 4** to

$$H_{\bigoplus_{i=1}^{k}\eta_i\mathbb{C}[\theta_1,\ldots,\theta_n]}(T) = \sum_{i=1}^{k}\frac{T^{\deg\eta_i}}{(1-T^{d_1})(1-T^{d_2})\ldots(1-T^{d_n})} = \frac{\sum_{i=1}^{k}T^{\deg\eta_i}}{\prod_{i=1}^{n}(1-T^{d_i})}$$

It follows that

$$H_{\bigodot V^{*G}}(T)\cdot\prod_{i=1}^{n}(1-T^{d_i}) = T^{\deg\eta_1}+T^{\deg\eta_2}+\ldots+T^{\deg\eta_k}$$

and equating coefficients we see that $k = d$ is the number of secondary invariants and $\deg\eta_i = e_i$ is the degree of η_i. □

6.1 Example: graphs

By **theorems 6, 7**, we can find the number and degrees of some system of secondary invariants given a set of primary invariants and the Molien series. **Theorem 5** implies that

$$\boldsymbol{\theta} = \left\{R_{S_n}(x^k_{\mathbf{\{1,2\}}})\big|k = 1,2\ldots,\binom{n}{2}\right\}$$

is a homogeneous system of parameters for the invariant ring of the set of undirected graph edges under S_n action.

We apply **theorem 7**. The degrees of a corresponding sequence of secondary invariants can be retrieved as the exponents of the generating function

$$T^{e_1}+T^{e_2}+\ldots+T^{e_d} = H_{\bigodot\mathbb{C}X_n{}^{*G}}(T)\cdot\prod_{i=1}^{\binom{n}{2}}(1-T^i) \tag{4}$$

The S_n action is a permutation action, therfore **lemma 5** can be applied provided the cycle types of $\sigma\in S_n$ acting on X_n. To be able to compute cycle types, we first prove that it is enough to compute the cycle type of one representative from each conjugacy class of S_n. In **section 6.2**, we describe a general method of computing cycle types for the S_n action on X_n.

Cycle types are constant on conjugacy classes

The S_n-action on X_n induces a homomorphism $\phi : S_n \longrightarrow S_{X_n}$. The cycle type of $\sigma\in S_n$ acting on X_n is defined as the cycle type of $\phi(\sigma)\in S_{X_n}$. Formulated this way, it is easy to see that this cycle type is conjugation-invariant: indeed, let $\pi\in S_n$ be any other permutation. Then $\phi(\pi\sigma\pi^{-1}) = \phi(\pi)\phi(\sigma)\phi(\pi)^{-1}$ is a conjugate of $\phi(\sigma)$ and therfore has the same cycle type.

6.2 Cycle index of S_n action on X_n

In this section, we demonstrate how to compute the cycle index of S_4 acting on X_4 and outline a solution for general n.

Assume σ is a permutation of S_4 and $\{a, b\} \in X_4$ is an edge. We distinguish between two cases:

- a, b are in different cycles of σ. Let the cycles have length $|C_a| = l_a,\ |C_b| = l_b$. Then there is no k for which $\sigma^k(a) = b$, and $\sigma^k \bullet \boldsymbol{\{a, b\}} = \boldsymbol{\{a, b\}}$ if and only if $\sigma^k(a) = a, \sigma^k(b) = b$. This is equvalent to $k \mid l_a,\ k \mid l_b$ and it follows that the cycle length of a, b is $\operatorname{lcm}(l_a, l_b)$.

- a, b are in the same cycle of σ. Assume that $a = x_1,\ b = x_l$ and the cycle is $(x_1, x_2, \ldots, x_d)$. Then $\sigma^k \bullet \boldsymbol{\{a, b\}} = \boldsymbol{\{a, b\}}$ can happen when $\sigma^k(a) = a,\ \sigma^k(b) = b$ and also if $\sigma^k(a) = b,\ \sigma^k(b) = a$. The last case occurs only if d is even and $l - 1 = \frac{d}{2}$.

S_4 has 4 conjugacy classes. The identity permutation leaves every edge invariant and has cycle type 1^6. A permutation with cycle type $(**)(*)(*)$ partitions the vertices into a set of 2 vertices forming the 2-cycle and the 2 vertices in 1-cycles. There are $2 \cdot 2 = 4$ edges going between the cycles. Each of these edges is part of a cycle of 2 elements. The 2 edges with both endpoints within one of the sets are invariant.

The cycle type of each element from the conjugacy class $(**)(*)(*)$ is therefore $1^2 2^2$.

Similar calculations for other cycle types lead to the following table:

cycle type	conjugacy class size	cycle type in X_4
(●)(●)(●)(●)	1	(●)(●)(●)(●)(●)(●)
(●●)(●)(●)	6	(●●)(●●)(●)(●)
(●●)(●●)	3	(●●)(●●)(●)(●)
(● ● ●)(●)	8	(● ● ●)(● ● ●)
(● ● ●●)	6	(●●)(● ● ●●)

Chapter 14 of [BVS76] explains how to compute the cycle types for general S_n acting on X_n using the same techniques as this section.

6.2.1 Example: graphs, $n = 4$

We will now demonstrate the computation on undirected graph edges X_4 under S_4-action.

By **Lemma 5** and the cycle type table, the Molien series of $\bigodot \mathbb{C}{X_4}^{*S_4}$ is

$$\begin{aligned} H_{\bigodot \mathbb{C}X_4^{*S_4}}(T) = \\ \frac{1}{4!}\Big(\frac{1}{(1-T)^6} + 9\frac{1}{(1-T)^2(1-T^2)^2} + \\ +8\frac{1}{(1-T^3)^2} + 6\frac{1}{(1-T^2)(1-T^4)}\Big) = \\ \frac{T^8 - T^7 + T^6 + T^4 + T^2 - T + 1}{(T-1)^6(T+1)^2\,(T^2+1)\,(T^2+T+1)^2} = \\ 1 + T + 3T^2 + 6T^3 + 11T^4 + 18T^5 + 32T^6 + 48T^7 + 75T^8 + \ldots \end{aligned} \tag{5}$$

By **equation 4** a homogeneous system of parameters $\eta_1, \eta_2, \ldots, \eta_d$ corresponding to $\boldsymbol{\theta} = \left\{ R_{S_n}(x^k_{\mathbf{\{1,2\}}}) \middle| k = 1, 2 \ldots, \binom{n}{2} \right\}$ would by **equation 5** have exponent sequence

$$\begin{aligned} H_{\bigodot \mathbb{C}X_4^{*S_4}}(T) \cdot \prod_{i=1}^{\binom{4}{2}} (1 - T^i) = \\ \frac{T^8 - T^7 + T^6 + T^4 + T^2 - T + 1}{(T-1)^6(T+1)^2\,(T^2+1)\,(T^2+T+1)^2} \cdot \prod_{i=1}^{\binom{4}{2}} (1 - T^i) = \\ T^{15} + T^{13} + 2T^{12} + 2T^{11} + 2T^{10} + 4T^9 + 3T^8 + \\ 3T^7 + 4T^6 + 2T^5 + 2T^4 + 2T^3 + T^2 + 1 \end{aligned}$$

We see that there are 30 corresponding secondary invariants invariants with degrees ≤ 15.

6.3 Algorithm: separating set for graphs

We describe an algorithm that is relatively efficient in practice for computing a separating set of polynomials for the set of graphs with n vertices.

1. Compute the Molien series $H_{\bigodot \mathbb{C}X_n^{*S_n}}(T)$ using **lemma 5** with cycle indices computed using the technique in **section 6.2**.

2. Use **theorem 5** to choose $\boldsymbol{\theta} = \left\{ R_{S_n}(x^k_{\mathbf{\{1,2\}}}) \middle| k = 1, 2 \ldots, \binom{n}{2} \right\}$ to be a homogeneous system of polynomials.

3. Compute the polynomial

$$\begin{aligned} H_{\bigodot \mathbb{C}X_n^{*S_n}}(T) \cdot \prod_{i=1}^{\binom{n}{2}} (1 - T^i) = \\ c_1 T^{e_1} + c_2 T^{e_2} + \ldots + c_r T^{e_r} \end{aligned}$$

4. Let $E = \emptyset$ be the set that will contain the η_k. For each $i = 1, 2, \ldots, r$, let $S = \emptyset$, M a sorted list of all monomials of $\left[\mathbb{C}[x_{\boldsymbol{\{1,2\}}}, \ldots, x_{\boldsymbol{\{n-1,n\}}}]\right]_{e_i}$ and iterate through the monomials m_j of M doing the following:
   (a) If $|S| = c_i$, we have found sufficiently many generators of degree e_i. We add S to E and continue with $i + 1$.
   (b) Let m be the current monomial. Use the algorithm in **section B.4** to check if $R_{S_n}(m)$ is contained in the $\mathbb{C}[\theta_1, \ldots, \theta_{\binom{n}{2}}]$-module that E generates. If m is contained in that model, continue with next m. Otherwise, add $R_{S_n}(m)$ to S and continue with next m.
5. At this step, $E = \{\eta_1, \ldots, \eta_d\}$ is a sequence of secondary invariants.
6. Output the set $\theta_1, \theta_2, \ldots, \theta_{\binom{n}{2}}, \eta_1, \eta_2, \ldots, \eta_d$.

Theorem 8. *The algorithm described above correctly finds a sequence $\{\eta_1, \eta_2, \ldots, \eta_s\}$ of secondary invariants that together with the chosen sequence $\boldsymbol{\theta}$ form a Hironaka decomposition of*

$$\bigodot \mathbb{C}X_n^{*S_n}$$

and returns a valid separating set for solving the graph isomorphism problem.

Proof. We first prove that at each step the E is a free basis for

$$\bigoplus_{\eta\in E} \eta\mathbb{C}[\theta_1, \theta_2, \ldots, \theta_{\binom{n}{2}}]$$

This is done by induction on $|E|$. Assume we add a new element $R_G(m)$ to E. If there are polynomials $p_\eta(\theta_1, \ldots, \theta_{\binom{n}{2}}), e \in E$ with $\sum_{\eta\in E} k_\eta \eta p_\eta(\theta_1, \ldots, \theta\binom{n}{2}) = 0$, we can assume $k_\eta = 1$ and rearrange $R_G(m) = \sum_{\eta\in E\setminus\{R_G(m)\}} \eta p_\eta(\theta_1, \ldots, \theta\binom{n}{2}))$ which contradicts the choice of $R_G(m)$ in **section B.4**.

Next, we prove that the algorithm finds exactly c_i polynomials of degree d_i for each $i = 1, \ldots, r$. Assume that the algorithm finishes with less that the full amount of secondary invariants. Then consider a modification of the algorithm where we remove step $4.(a)$. Then the algorithm would check every monomial up to degree e_r.

We claim that the modified version cannot produce less than the whole set of invariants and that the original version terminates with the same end result that the modified one.

The modified version tests each $R_G(m)$ where m is a monomial with total degree $\leq e_i$ for membership and adds it if it is not a member. In the end, every monomial with degree $\leq m$ will have $R_G(m) \in \bigoplus_{\eta \in E} \eta\mathbb{C}[\theta_1, \dots, \theta_{\binom{n}{2}}]$. But this implies that

$$\left[\bigodot \mathbb{C}X_n^{*S_n}\right]_d \subseteq \bigoplus_{\eta \in E} \eta\mathbb{C}[\theta_1, \dots, \theta_{\binom{n}{2}}]$$

because of the projection property of $R_G(m)$.

Let E' be a hypothetical set of secondary invariants that is guaranteed to exists by **theorem 6**.

By **theorem 7**, we know the degree sequence of E'. In particular, we know that for each $\eta' \in E'$, there is a degree d such that $\eta' \in \left[\bigodot \mathbb{C}X_n^{*S_n}\right]_d$

This implies

$$\eta' \in \bigoplus_{\eta \in E} \eta\mathbb{C}[\theta_1, \dots, \theta_{\binom{n}{2}}]$$

and therefore

$$\bigoplus_{\eta' \in E'} \eta'\mathbb{C}[\theta_1, \dots, \theta_{\binom{n}{2}}] \subseteq \bigoplus_{\eta \in E} \eta\mathbb{C}[\theta_1, \dots, \theta_{\binom{n}{2}}]$$

Strict inclusion is impossible, as we cannot have generated a larger ring than the whole invariant ring. Therefore the inclusion is equality, and E is a full system of secondary invariants.

Now consider the difference between the original algorithm and the modified version with $4.(a)$ omitted. By **theorem 7**, the degrees of $\eta \in E$ are uniquely determined. It follows that the algorithm will never find another $R_G(m)$ to add to S in step 4 after S already has the maximal size and we can add $4.(a)$ for increasing efficiency.

The set $\boldsymbol{\theta} \cup E$ of primary and secondary invariants is clearly a generating set of $\bigodot \mathbb{C}X_n^{*S_n}$ by the Hironaka decomposition. □

6.3.1 Mathematica implementation

We demonstrate steps $2, 3$ and 4 of the algorithm in `mathematica`

```
(* We investigate X_4 again*)
n = 4

(* theta is the Binomial[n, 2] power-sum symmetric ↩
   polynomials *)
theta = Table[
   reynolds[
    (variables[x, n])[[1]]^k,
    x,
    n], {k, Length[variables[x, n]]}
   ];

(* Standard vector space monomial basis of the graded component [C[var]]_d*)
dimDComponent[var_, d_, n_] :=
 Table[
  Times @@ MapThread[#1^#2 &,
    {variables[x, 4],
     expVector}
    ],
  {expVector,
   FrobeniusSolve[
    ConstantArray[1, Binomial[n, 2]], d]}
  ]

(* Ex, Binomial[4, 2] = 6 variables, all deg-2 monomials *)
dimDComponent[x, 2, n]

eta = {reynolds[x[1, 2] x[2, 3], x, n]};

hironakaRingMember[eta, theta,
 reynolds[x[1, 2] x[2, 3], x, n],
 variables[x, 4]
 ]
```

This program computes a Gröbner basis for an ideal of 7 polynomials in a ring with 13 variables. It turned out to be too many polynomials and variables for `mathematica` and as a consequence, the algorithm does not finish in a reasonable amount of time.

6.4 Analysis

The algorithm in **theorem 8** is in practice much more efficient than the algorithm from **section 3.2**. It seems to be hard to prove concrete results on the running time, amount and monomials considered to find the secondary invariants [Thi00].

Nevertheless, we can still prove that the algorithm is a major theoretical improvement over **section 3.2**. The example in **section 6.2.1** for X_4 has maximal degree of secondary invariants equal to 15. This means that the algorithm will do at most $\binom{\binom{n}{2}+15}{6}$ Gröbner basis calculations instead of the number $\binom{\binom{n}{2}+24}{6}$ from Noether's theorem.

We can in general derive a better bound than the $\Omega\left(\binom{\binom{n}{2}+n!}{\binom{n}{2}}\right)$ that follows from Noether's theorem by inspecting the degree of the Molien series.

Our chosen set of primary invariants has degrees $1, 2, \dots \binom{n}{2}$. By **theorem** 7, the maximal degree of a secondary invariant is

$$
\begin{aligned}
\deg H_{\bigodot \mathbb{C}X_n^*{}^{S_n}}(T) \cdot \prod_{i=1}^{\binom{n}{2}} (1 - T^i) &= \\
\deg \frac{1}{n!} \sum_{\sigma \in S_n} \frac{1}{\det(I_{\binom{n}{2}} - TM_\sigma)} \cdot \prod_{i=1}^{\binom{n}{2}} (1 - T^i) &= \\
\frac{\binom{n}{2}\left(\binom{n}{2}+1\right)}{2} - \binom{n}{2} \leq \binom{n}{2}^2 &
\end{aligned}
$$

There are less than $\binom{\binom{n}{2} + \binom{n}{2}^2}{\binom{n}{2}}$ monomials that the algorithm in **theorem** 8 considers, which is asymptotically an improvement over $\binom{\binom{n}{2} + n!}{\binom{n}{2}}$ Gröbner basis computations in the algorithm derived from Noether's theorem.

There is an even larger practical improvement that is difficult to account for arising from the fact that the algorithm jumps to another degree after having found the correct amount of monomials without a need to test all.

7 Results and discussion

This work mainly constitutes an introduction to invariant theory with isomorphism of graphs as a main example. For further reading, we suggest the survey [Sta79] by Richard P. Stanley.

The problem of finding separating sets by studying the invariant ring is very hard, making it unsurprising that our simple implementation did not completely characterize the invariant ring. The following is a quote on this topic made by Nicolas Thiery[Thi00]:

> ...there is a combinatorial explosion in the computations involved and, to our knowledge, the ring J_n has only been completely described for $n \leq 4$.

In conclusion, I would say that while this project may not have resulted in a useful algorithm for efficiently solving the graph isomorphism problem, it has made me familiar with several topics of beautiful mathematics.

A Symmetric tensor power construction

This section contains the definitions and lists basic proofs of the symmetric tensor power $\bigodot^n V$ and the symmetric tensor algebra for vector spaces V.

Definition of the tensor product $V \otimes W$ for vector spaces

Let V, W be a vector spaces over a field $\mathbb{K}$. For a set S, denote $\mathbb{K}\langle S \rangle$ to be the free vector space on S defined as all finite sums $\sum_{s \in T} k_s \mathbf{s}$ for finite subsets $T \subset S$ with addition $k_1\mathbf{s} + k_2\mathbf{s} = (k_1 + k_2)\mathbf{s}$.

On $\mathbb{K}\langle V \times W \rangle$, form the subspace U generated by all

$$\begin{aligned} k(v, w) - (kv, w) \\ (kv, w) - (v, kw) \\ (v_1 + v_2, w) - (v_1, w) + (v_2, w) \\ (v, w_1 + w_2) - (v, w_1) + (v, w_2) \end{aligned}$$

for all $v, v_1, v_2 \in V, w, w_1, w_2 \in W, k \in \mathbb{K}$. Then define

$$V \otimes W := \mathbb{K}\langle V \times W \rangle /_U$$

Let $\otimes : V \times W \longrightarrow V \otimes W$ be the restriction of the corresponding projection operator and denote $\otimes(v, w) = v \otimes w$ From the definition of the subspace it follows that

$$\begin{aligned} k(v \otimes w) = kv \otimes w = v \otimes kw \\ (v_1 + v_2) \otimes w = (v_1 \otimes w) + (v_2 \otimes w) \\ v \otimes (w_1 + w_2) = (v \otimes w_1) + (v \otimes w_2) \end{aligned}$$

A.1 Iterating the procedure

We define the $n : th$ tensor power $V^{\otimes n}$ as

$$V^{\otimes n} = \underbrace{V \otimes (V \otimes (\ldots V) \ldots)}_{n \text{ times}}$$

Due to the symmetry of the construction, there is an isometry $V \otimes (V \otimes V) \cong (V \otimes V) \otimes V$ defined by $v_1 \otimes (v_2 \otimes v_3) \mapsto (v_1 \otimes v_2) \otimes v_3$. Parentheses are therefore superfluous as the order of taking products is irrelevant.

Construction of symmetric tensor power $\bigodot^n V$

We let S_n act on the set $V^{\otimes n}$. S_n acts on the generators $v_1 \otimes \ldots \otimes v_n$ by

$$\sigma \bullet v_1 \otimes \ldots \otimes v_n := v_{\sigma(1)} \otimes \ldots \otimes v_{\sigma(n)}$$

We define $\bigodot^n V$ to be the quotient of $V^{\otimes n}$ and the space generated by all $(\mathbf{x} - \sigma \bullet \mathbf{x})$ for all; $x \in V^{\otimes n}$ and $\sigma \in S_n$. In $\bigodot^n V$, we denote $v_1 \otimes v_2 \otimes \ldots \otimes v_n$ to be the equivalence class of $v_1 \otimes v_2 \otimes \ldots \otimes v_n \in V^{\otimes n}$. Then for each $v_1 \otimes v_2 \otimes \ldots \otimes v_n \in \bigodot^n V$, and $\sigma \in S_n$,

$$v_1 \otimes v_2 \otimes \ldots \otimes v_n = v_{\sigma(1)} \otimes v_{\sigma(2)} \otimes \ldots \otimes v_{\sigma(n)}$$

because the difference $v_1 \otimes v_2 \otimes \ldots \otimes v_n - v_{\sigma(1)} \otimes v_{\sigma(2)} \otimes \ldots \otimes v_{\sigma(n)}$ lies in the space generated by all $\mathbf{x} - \sigma\mathbf{x}$.

Relation to monomials

Let X be a basis of V. Then any $\mathbf{x} \in V^{\otimes n}$ is (using linearity) a finite linear combination of terms $\bigotimes_{i=1}^n x_i$ with $x_i \in X$. Since projection is surjective, the same is true for $\bigodot^n V$: every $\mathbf{x} \in \bigodot^n V$ is a finite linear combination of terms $\bigotimes_{i=1}^n x_i$ with $x_i \in X$.

The following theorem shows which of these terms construct a basis:

Theorem 9. *Let V be a vector space over $\mathbf{K}$, X be a basis for V. For each selection T of n elements of X with repetition, construct $x_T = \bigotimes_{t \in T} t \in \bigodot^n V$ (order of summation is irrelevant since $\bigotimes$ is commutative in $\bigodot^n V$). Then the set of all x_T forms a basis for $\bigotimes^n V$.*

The symmetric algebra

Theorem 10. *If $\boldsymbol{X}$ is a basis for V, then*

$$\bigoplus_{n \geq 0} \left(\bigodot^n V \right)$$

is a graded algebra isomorphic to the free commutative graded algebra $\mathbb{K}[\boldsymbol{X}]$ under concatenation of tensor products.

Proof. This follows from the construction and from **theorem** 9. □

Omitted proofs and detailed discussions can be found in [Fuh11] and [FH91].

B Gröbner bases

The aim of this section is to describe two algorithms: one is the algorithm for computing Gröbner bases. We do not give a detailed implementation, only prove that a reduced Gröbner basis exist and sketch a version of the algorithm. The other one is relevant to us and is used in **section** 4. It uses Gröbner bases and computes the following: given polynomial generators $F_1, \dots, F_k$ that generate a subring $R \subseteq \mathbb{C}[x_1, \dots, x_n]$ determines whether a given polynomial $p \in \mathbb{C}[x_1, \dots, x_n]$ belongs to R. This section is mostly based on **chapter 11** of [GCL92] and **chapter 1.2** of [PS08]

B.1 Definitions and basic properties

Definition 5. *A monomial order is a total order of monomials $\prec$ such that*

- $a \prec b \implies ac \prec bc$ *for all nonconstant monomials* a, b, c
- $1 \prec a$ *for all nonconstant monomials* a.

An example of a monomial order of monomials in $x_1 \dots, x_n$ is the lexicographic order on the exponent vectors. It is clear that it is total and that it satisfies both conditions.

Definition 6. *For each $p \in \mathbb{C}[x_1, \dots, x_n]$ and a given monomial order $\prec$ on $x_1, \dots, x_n$, we define the* ***leading term of*** *p* ***with respect to*** *$\prec$, $LT_\prec(p)$ to be the largest monomial of p under $\prec$.*

Fix a monomial order $\prec$ and abbreviate $LT_\prec(p) = LT(p)$. For any ideal $I \subset \mathbb{C}[x_1, \dots, x_n]$, we define the ***leading ideal*** *$LT(I)$ as the ideal generated by all leading monomials:*

$$LT(I) = \langle LT(p) | p \in I \rangle$$

E.g. if $\prec$ is lexicographic order with $z \prec y \prec x$, then $LT_\prec(2x^2y + 3xy^2 + 4y^2z) = 2x^2y$.

Definition 6 is enough to give a useful definition of Gröbner bases:

Definition 7. *Let I be an ideal of $\mathbb{C}[x_1, \dots, x_n]$ and $\prec$ some monomial order. Then a finite set $\mathcal{G}$ of elements in I, $\mathcal{G} = \{g_1, \dots, g_k\} \subset I$, is called a* ***Gröbner basis of*** *I if the set $LT(\mathcal{G}) = \{LT(g_1), \dots, LT(g_k)\}$ generates the ideal $LT(I)$.*

If $LT(g_j)$ does not divide any monomial in g_i for all $i \neq j$ in $\{1, \dots, k\}$, we call $\mathcal{G}$ a ***reduced Gröbner basis of*** *I.*

We will now prove the existence of Gröbner bases for every ideal and monomial order and prove properties that will lead to algorithms later.

Lemma 6. *Every ideal M in $\mathbb{C}[x_1, \dots, x_n]$ generated by a (possibly infinite) set of monomials is generated by a finite number of monomials.*

Proof. There are countably many monomials, so we can enumerate all monomials in M by $\{m_1, \dots, m_k, \dots\}$. The ring $\mathbb{C}[x_1, \dots, x_n]$ is Noetherian, so the chain of ideals

$$M_1 = \langle m_1 \rangle \subset M_2 = \langle m_1, m_2 \rangle \subset \dots \subset M_k = \langle m_1, m_2, \dots, m_k \rangle \subset \dots$$

is finite and

$$M = \bigcup_{i=1}^{\infty} M_i = M_r$$

for some index r. It follows that $m_1, \dots, m_r$ are generators of M. □

Now we can give a simple proof for the existence of Gröbner bases for every ideal.

Theorem 11. *Every ideal I has a Gröbner basis $\mathcal{G}$ for every monomial order $\prec$.*

Proof. By the theorem above, $LT(I)$ has a finite monomial basis $m_1, \dots, m_s$. There are polynomials $g_i \in I$ with $g_i = LT(m_i)$ for each I. The set $\mathcal{G}$ of all such polynomials is a Gröbner basis for I in the sense of **Definition 7**

□

Nothing in the definition tells us that a Gröbner basis generates its ideal. But this is an important motivation for studying and calculating Gröbner bases.

Theorem 12. *A Gröbner basis $\mathcal{G}$ generates its ideal.*

Lemma 7. *A monomial order $\prec$ is a well-ordering of monomials.*

Proof. Assume there is an infinite decreasing chain $m_1 \succ m_2 \succ \dots$. We know that the ideal of the m_i is generated by the first k m_i for some k. It follows that every later m_i is a multiple of some m_j, $j \leq k$. But that is a contradiction. □

Now we prove the main theorem.

Proof. Assume that $I \setminus \langle \mathcal{G} \rangle$ is nonempty. Take $p \in I \setminus \langle \mathcal{G} \rangle$ with minimal leading term. $p \in I$, so $LT(p) \in LT(I)$. But we know that $LT(I) = LT(\mathcal{G})$, so $LT(p)$ is a multiple of some $LT(g_i)$. We can write $LT(p) = mLT(g_i)$ for a monomial m. Consider $q = p - \lambda m g_i$ with $\lambda \in \mathbb{C}$. $q \notin \langle \mathcal{G} \rangle$ because otherwise $p \notin \langle \mathcal{G} \rangle$. We choose λ so that the leading term of p cancels out. We claim $LT(q) \prec LT(p)$. This is true because $LT(p) = LT(m g_i)$ and $LT(a + b) \preceq \max\{LT(a), LT(b)\}$ and $LT(q) \neq LT(p - \lambda m g_i)$. we get $LT(q) \preceq LT(p - \lambda m g_i) \preceq \max\{LT(p), LT(\lambda m g_i)\} = LT(p)$. Since $q \in I \setminus \langle \mathcal{G} \rangle$, we have a contradiction. □

Now, we define *reduction* modulo a Gröbner basis and show uniqueness of the reduction.

B.2 Reduction and spanning property

There is an algorithm for reducing a polynomial to a unique remainder modulo an ideal I associated with a Gröbner basis of I. We call a monomial m **standard** if $m \notin LT(I)$ and **nonstandard** if $m \in LT(I)$. We claim that

Lemma 8. *The equivalence classes of the standard monomials are linearly independent over $\mathbb{C}$ and span $\mathbb{C}[x_1, \dots, x_n] /_I$ as a $\mathbb{C}$-vector space.*

This implies that every polynomial p has a unique corresponding sum of standard monomials called remainder modulo I

Proof. Let $p \in \mathbb{C}[x_1, \dots, x_n]$. We are to find $q \in \mathbb{C}[x_1, \dots, x_n]$ with $q = \sum_i \lambda_i m_i$ where $\lambda_i \in \mathbb{C}$, m_i are standard monomials and $p \sim q$, equivalently $p = g + q$ with $g \in I$.

Let $m = HNST(p)$ be the highest nonstandard monomial in p. Then $m \in LT(I)$, equivalently there is $r \in I$ with $m = m' LT(r)$. We then write $p' = p - HNST(p) \cdot m' LT(r)$. The polynomial p' has $HNST(p') \prec HNST(p)$, therefore repeating this operation decreases until every monomial is standard. At each step, we subtract elements of I, thus the final result q has $p \sim q \mod I$. □

This proves that the standard monomials span

$$\mathbb{C}[x_1, \dots, x_n] /_I$$

To prove linear independence in the quotient ring, assume $\sum_i \lambda_i m_i \in I$ where m_i are standard. But the m_i are by definition $\notin I$, so the only way the sum $\sum_i \lambda_i m_i$ is in I is when the sum is 0. But then the m_i are linearly independent over $\mathbb{C}$ and so every $\lambda_i = 0$ and we are done.

Reduced Gröbner basis

The Gröbner basis functions in various software packages always return reduced Gröbner bases. These Gröbner bases are *unique* for a given monomial order and can be used to compute the unique remainder, which is proven in [GCL92]. For instance in `Mathematica`, the function `GroebnerBasis` expects a list of generators of an ideal and an ordering and returns the reduced Gröbner basis for that ideal whilst the function `PolynomialReduce`, when supplied a reduced Gröbner basis, a polynomial and an ordering returns the unique remainder from **Lemma** 8. The algorithms used are explained in detail and proven correct in [GCL92]

B.3 Ring membership algorithm

The algorithm that utilizes repetitive division from **Lemma** 8 is suitable for checking membership and remainder of a polynomial over an *ideal*, but can not be used right away to test membership in a *ring*.

The problem is as follows: We have a polynomial p and polynomials $f_1, \ldots, f_n$. We want to check if there exists a q such that $p = q(f_1, \ldots, f_n)$ and find q if it exists. It turns out that we can formulate this as a reduction problem for ideals and reuse the result **Lemma** 8

Theorem 13. *Let p be a polynomial in variables $x_1, \ldots, x_m$. Let R be a subring of $\mathbb{K}[x_1, \ldots, x_m]$ generated by $f_1, \ldots, f_n$.*

In the ring $\mathbb{K}[x_1, \ldots, x_m, y_1, \ldots, y_n]$, form the ideal $I = (f_1 - y_1, \ldots, f_n - y_n)$.

Assume that $\prec$ is a monomial order with $y_j \prec x_i$ for every i, j (an example is the lexicographic order).

*Then let q be the reduction of p with respect to I and $\prec$ defined in **Lemma** 8.*

Then the following holds:

- *q is a polynomial in $y_1, \ldots, y_n$ only and not $x_1, \ldots, x_m$ if and only if $p \in R$.*
- *If p is in the subring, and then q evaluated in $f_1, \ldots, f_n$ satisfies*

$$q(f_1, \ldots, f_n) = p.$$

Proof. First, observe that

$$q(x_1, \ldots, x_m, f_1, \ldots, f_n) = p$$

To verify this, write

$$p(x_1, \ldots, x_m) = A_1(f_1 - y_1) + \ldots + A_n(f_n - y_n) + q(x_1, \ldots, x_m, y_1, \ldots, y_n)$$

The left hand side is an expression only in $x_1 \ldots x_m$, so this has to be valid for all values of y_i. We substitute $y_i = f_i$. Then we get

$$p(x_1, \ldots, x_m) = 0 + \ldots + 0 + q(x_1, \ldots, x_m, f_1, \ldots, f_n)$$

Assume that q is a polynomial only over the y_i. Then we know that $p = q(f_1, \dots, f_n)$ because of the above.

We will prove that q doesn't have any x_i:s. Assume that $p = r(f_1, \dots, f_n)$ for some r. Then each monomial $\mathbf{f}^{\boldsymbol{\alpha}}$ can be written as

$$\mathbf{f}^{\boldsymbol{\alpha}} = ((f_1 - y_1) + y_1)^{\alpha_1} \dots ((f_n - y_n) + y_n)^{\alpha_n} = A_1(f_1 - y_1) + \dots + A_n(f_n - y_n) + r(y_1, \dots, y_n)$$

and it follows that the whole of p can be expressed that way in at least one way. Thus we have

$$p = A_1(f_1 - y_1) + \dots + A_n(f_n - y_n) + r(y_1, \dots, y_n)$$

for some r. It follows that $p \sim r(y_1, \dots, y_n)$ modulo I. Every monomial $\boldsymbol{y}^{\boldsymbol{\alpha}}$ is non-standard since the $y_i \prec x_i$ and $LT(I)$ contains no monomials in $\boldsymbol{y}$. □

B.3.1 Mathematica implementation

We implement the ring membership algorithm in `mathematica`:

```
ringMembership[generators_, poly_, vars_] :=
 Block[{nvars, groebB, allvars, remainder, z},
  nvars = Table[z[i], {i, Range[Length[generators]]}];
  allvars = Join[vars, nvars];
  groebB = GroebnerBasis[
    MapThread[
     Subtract,
     {generators, nvars}
     ],
    allvars
    ];
  remainder = PolynomialReduce[poly, groebB, allvars][[2]]
    /. Table[var -> 1, {var, nvars}];
  Internal`RealValuedNumericQ[remainder]
  ]

(* example, check if 'x+y' lies in ring C[x,y] *)
ringMembership[{x, y}, x + y, {x, y}]
ringMembership[{}, 0, {x, y}]
ringMembership[{x^2, y}, x^2 y + y^2, {x, y}]
```

This algorithm may be used to reduce the size of a generating set of a ring. If $\mathcal{G}$ generates the ring $\mathbb{C}[\mathcal{G}]$, we can iterate through $\mathcal{G}$ and check if a later member already lies in the subring generated by the previous ones. In `mathematica`, it would look like this:

```
(* If newPoly lies in C[oldB], returns oldB. Otherwise,↩
    returns oldB + {newPoly} *)
step[oldB_, newPoly_, vars_] :=
 If[ringMembership[oldB, newPoly, vars],
  oldB,
  Join[oldB, {newPoly}]
  ]

(* Reduce a large list 'genList' of generators in ↩
   variables 'vars' to a hopefully smaller one *)
reduceGenerators[genList_, vars_] :=
 Fold[
  step[#1, #2, vars] &,
  {},
  genList
  ]

(* example on the generators for the graph invariant ↩
   with n=3 *)
reduceGenerators[generators[x, 3], variables[x, 3]]
```

Running this cone on the 84 generators of the graph invariant ring from **section 3.2**, we end up with the answer

$$\left\{\frac{1}{3}x_{1,2}+\frac{1}{3}x_{1,3}+\frac{1}{3}x_{2,3},\frac{1}{3}x_{1,2}^2+\frac{1}{3}x_{1,3}^2+\frac{1}{3}x_{2,3}^2,\frac{1}{3}x_{1,2}^3+\frac{1}{3}x_{1,3}^3+\frac{1}{3}x_{2,3}^3\right\}$$

That are the 3 power-sum symmetric polynomials that generate the algebra of symmetric polynomials of 3 variables.

B.4 Hironaka ring membership algorithm

We describe a modification of **algorithm B.3.1** to test whether a polynomial $p(\boldsymbol{x})$ belongs to a ring given as a set E, ***theta*** of primary and secondary generators for Hironaka decomposition $\bigoplus_{\eta\in E}\eta\,\mathbb{C}[\theta_1,\ldots,\theta_n]$.

The set of $E\cup\boldsymbol{\theta}$ is a generating set of $\bigoplus_{\eta\in E}\eta\,\mathbb{C}[\theta_1,\ldots,\theta_n]$, therefore the ordinary ring membership algorithm from **section B.3.1** applies.

A `mathematica` implementation follows:

B.4.1 Mathematica implementation

```
(* checks if poly lies in the ring with hironaka ↩
   decomposition eta, theta. Everything is a polynomial↩
    in vars *)
hironakaRingMember[eta_, theta_, poly_, vars_] :=
 ringMembership[Join[eta, theta], poly, vars]
```

References

[BL83] László Babai and Eugene M. Luks. Canonical labeling of graphs. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, STOC '83, pages 171–183, New York, NY, USA, 1983. ACM.

[BVS76] B.B Belov, E.M. Vorobev, and V.E. Shatalov. *Teoriya grafov.* Vyshaya shkola, 1976.

[FH91] W. Fulton and J. Harris. *Representation Theory: A First Course.* Graduate Texts in Mathematics / Readings in Mathematics. Springer New York, 1991.

[Fuh11] P.A. Fuhrmann. *A Polynomial Approach to Linear Algebra.* Universitext. Springer, 2011.

[GCL92] K.O. Geddes, S.R. Czapor, and G. Labahn. *Algorithms for Computer Algebra.* Kluwer Academic, 1992.

[PS08] P. Paule and B. Sturmfels. *Algorithms in Invariant Theory.* Texts & Monographs in Symbolic Computation. Springer, 2008.

[Sch87] Uwe Schöning. Graph isomorphism is in the low hierarchy. In FranzJ. Brandenburg, Guy Vidal-Naquet, and Martin Wirsing, editors, *STACS 87*, volume 247 of *Lecture Notes in Computer Science*, pages 114–124. Springer Berlin Heidelberg, 1987.

[Sta79] Richard P. Stanley. Invariants of finite groups and their applications to combinatorics. *BULLETIN OF AMER. MATH. SOC*, 1(3):475–511, 1979.

[Thi00] Nicolas M. Thiéry. Algebraic invariants of graphs; a study based on computer exploration. *SIGSAM Bull.*, 34(3):9–20, September 2000.