

THE IDEAL OF RELATIONS FOR THE RING OF INVARIANTS OF n POINTS ON THE LINE

BENJAMIN HOWARD, JOHN MILLSON, ANDREW SNOWDEN AND RAVI VAKIL*

ABSTRACT. The ring of projective invariants of n ordered points on the projective line is one of the most basic and earliest studied examples in Geometric Invariant Theory. It is a remarkable fact and the point of this paper that unlike its close relative the ring of invariants of n unordered points this ring can be completely and simply described. In 1894 Kempe found generators for this ring, thereby proving the First Fundamental Theorem for it (in the terminology introduced by Weyl). In this paper we compute the relations among Kempe's invariants, thereby proving the Second Fundamental Theorem (again in the terminology of Weyl), and completing the description of the ring 115 years later.

This paper introduces a number of new tools to the problem, and uses the graphical algebra formalism to intermediate between representation-theoretic arguments (for symmetric and Lie groups), and the symmetry-breaking of the Speyer-Sturmfels degeneration.

CONTENTS

1.	Introduction	1
2.	The invariant ring R_L	5
3.	The toric degenerations $\text{gr}_{\Xi} S_L$ and $\text{gr}_{\Xi} R_L$	6
4.	The toric ideal is generated by quadratics and toric generalized Segre cubics	11
5.	The ideal is generated by quadratics and generalized Segre cubics	16
6.	The structure of V_L and its tensor powers	19
7.	Retrogeneration of the ideal	26
8.	The ideal is generated by quadratics	33
9.	The quadratics are generated by the simplest binomials	39
	References	41

1. INTRODUCTION

We consider the ring of invariants of n points on the projective line, defined as the projective coordinate ring of the GIT quotient of $(\mathbb{P}^1)^n$ by the group $\text{SL}(2)$. This is a classical archetype of GIT and a common first example in the theory. To form the quotient, one must choose weights $w = (w_1, \dots, w_n)$; it is then given by

$$(\mathbb{P}^1)^n \dashrightarrow (\mathbb{P}^1)^n //_w \text{SL}(2) = \text{Proj } R_w$$

where $R_w = \bigoplus_{k=0}^{\infty} R_w^{(k)}$, with $R_w^{(k)} = \Gamma((\mathbb{P}^1)^n, \mathcal{O}(kw_1, \dots, kw_n))^{\text{SL}(2)}$. (Note that we use the notation $R^{(k)}$ for the k piece of a graded ring R .) The ring of invariants R_w turns out to be generated in its lowest nonzero degree, so the GIT quotient has a natural projective embedding. We denote the weight $(1, 1, \dots, 1)$ by 1^n .

Theorem 1.1 (Main Theorem, informal version). *If $w \neq 1^6$, the ideal of relations between lowest degree invariants is generated by quadratics.*

Detailed motivation and background for this problem are given in the announcement [HMSV2]. We describe there how small cases have long been known to yield beautiful classical geometry. In this paper, we show that this rich structure extends to any number of points with any weighting: the relations for the

Date: December 11, 2009.

*B. Howard was supported by NSF fellowship DMS-0703674. J. Millson was supported by the NSF grant DMS-0405606, the NSF FRG grant DMS-0554254 and the Simons Foundation. This research was partially conducted during the period A. Snowden was employed by the Clay Mathematics Institute as a Liftoff Fellow. A. Snowden was also partially supported by NSF fellowship DMS-0902661. R. Vakil was partially supported by NSF grant DMS-0801196.

moduli space are generated by a particularly simple type of quadratic, with the single exception of the Segre cubic for 6 points. In a precise sense, the ideal is cut out by essentially one equation, inherited from the $n = 8$ case. If Kempe’s theorem is the analogue of Weyl’s “First Main Theorem” for SL_2 (see [W]), then this is the analogue of his “Second Main Theorem.”

This is the culmination of a number of papers on this topic, so we wish to be clear about its relationship to the others. [HMSV1] (subsuming two earlier arXiv preprints) opened up this question: through blunt toric degeneration and explicit local calculation, hints that the relations were simple quadrics were observed. It was shown there that certain quadrics (not as simple as in this paper) cut out the space. The other important paper (in preparation) is on the interplay of the geometry and representation theory in the modular fivefold parametrizing eight points on the line. (The remaining papers are less important; one on S_6 describes some pretty algebra that fell out of our analysis. [HMSV3] will be subsumed into the representation-theory part of the modular fivefold paper. [HMSV2] is the announcement of the results in this paper and in the forthcoming fivefold paper. [HMSV4] is a technical companion to this paper, extending the results to more general base rings.)

The current paper is the main one of the series. It introduces fundamentally new techniques to the problem, and is not simply a refinement of our preliminary ideas of [HMSV1]. The quadric here is simpler than those speculated about in [HMSV1]. The argument uses the formalism of graphical algebras to balance on one hand representation theory arguments (of both symmetric and Lie groups), with on the other hand information from symmetry-breaking (from a better toric degeneration, discovered by Speyer and Sturmfels from tropical motivation).

1.1. Integrality issues. For simplicity, we prove our results over \mathbb{Q} , but our results apply more generally, as we show in [HMSV4]. In preparation for this, we prove intermediate results over more general base rings, which we hope will not distract readers interested in characteristic 0. In particular, this paper proves Theorem 1.1 over $\mathbb{Z}[\frac{1}{12!}]$, but the precise version (Theorem 1.2) only over \mathbb{Q} .

1.2. The graphical formalism. We use a graphical interpretation of the ring of invariants, which allows us to deal effectively with both the \mathfrak{S}_n -symmetries and the broken symmetries of toric degenerations. To a directed graph Γ on vertices labeled 1 through n (in bijection with the points), with valence vector $kw = (kw_1, \dots, kw_n)$, we associate an invariant in $R_w^{(k)}$:

$$X_\Gamma = \prod_{\vec{ij}} (x_i y_j - x_j y_i).$$

Here the product is taken over the edges \vec{ij} of Γ and x_i and y_i are projective coordinates for \mathbb{P}^1 . Note that if Γ has a *loop* — that is, an edge with the same source and target — then $X_\Gamma = 0$. (Throughout this paper, graphs are allowed to have loops and multiple edges between the same vertices.) A fundamental result from classical invariant theory states that the X_Γ span $R_w^{(k)}$. Kempe showed that for $w = 1^n$ and n even the ring R_w is generated in degree 1 (Theorem 2.1). Thus, in this situation, R_w is generated by those X_Γ where Γ is a *matching*, that is, a graph in which each vertex belongs to precisely one edge. A similar result holds for any n and w .

1.3. Relations. The theorem of Kempe mentioned above provides generators for the ring R_w . The purpose of this paper is to determine the relations between these generators. A number of obvious relations exist; we catalog some of them below. An important phenomenon shows itself already in these simple examples: relations on n points can be extended to give relations on more than n points. This is a key theme in our treatment and is discussed in detail in §7.1.

The most obvious relation is the *sign relation*: we have $X_{\vec{ab}} = -X_{\vec{ba}}$. This can be regarded as a relation between two invariants with $n = 2$ and $w = 1^2$. It extends to invariants with arbitrary n and w as follows: if Γ is any graph and Γ' is obtained from Γ by switching the direction of a single edge then $X_\Gamma = -X_{\Gamma'}$.

The next most simple relation appears when $n = 4$ and $w = 1^4$:

$$\begin{array}{c} \nearrow \\ \searrow \end{array} = \begin{array}{c} \uparrow \\ \uparrow \end{array} + \begin{array}{c} \longrightarrow \\ \longleftarrow \end{array}$$

This is the classical *Plücker relation*. It extends to each $R_w^{(k)}$ as well: given a graph Γ with valence vector kw and two edges \vec{ab} and \vec{cd} of Γ we have the identity

$$(1.1) \quad X_\Gamma = X_{\Gamma_1} + X_{\Gamma_2}$$

where Γ_1 and Γ_2 are the graphs obtained from Γ by replacing $\{\vec{ab}, \vec{cd}\}$ with $\{\vec{ad}, \vec{cb}\}$ and $\{\vec{ac}, \vec{bd}\}$ respectively. We will use the phrase “to Plücker two edges \vec{ab} and \vec{bc} of a graph Γ ” to mean “to replace X_Γ by $X_{\Gamma_1} + X_{\Gamma_2}$.” The sign and Plücker relations are both linear relations; in fact, they span all linear relations.

Some higher degree relations are conveniently thought of in terms of colored graphs. By a *k-colored graph* we mean a graph in which each edge has been assigned one of k colors. By a *multi-matching* of degree k we mean a k -colored graph in which each vertex appears in precisely one edge of each color. Let Γ be a multi-matching of degree k on n vertices. We define the element X_Γ of $(R_{1^n}^{(1)})^{\otimes k}$ to be the pure tensor $\otimes X_{\Gamma(i)}$ where the product is over the colors i and $\Gamma(i)$ is the subgraph of Γ on the edges of color i . In terms of colored graphs, the map $(R_{1^n}^{(1)})^{\otimes k} \rightarrow R_{1^n}^{(k)}$ “forgets the color.”

The *Segre cubic relation* is described with colored graphs as follows:

$$(1.2) \quad \begin{array}{c} \text{Diagram 1} \\ \text{Diagram 2} \end{array} = \begin{array}{c} \text{Diagram 3} \\ \text{Diagram 4} \end{array}$$

(For those readers for whom color is not available, we will follow the convention that the solid lines are green, the dashed red and the dotted blue.) Each edge should be directed in the same way on both sides of the equation. This relation holds because the graphs on each side have the same set of edges — only the colors are different. The Segre cubic relation extends to cubic relations on $R_{1^n}^{(1)}$ for any even $n > 6$. For example, we have the following relation on eight points starting from the Segre cubic on six points:

$$(1.3) \quad \begin{array}{c} \text{Diagram 1} \\ \text{Diagram 2} \end{array} = \begin{array}{c} \text{Diagram 3} \\ \text{Diagram 4} \end{array}$$

There is a “new” relation for $n = 8$, binomial and quadratic:

$$(1.4) \quad \begin{array}{c} \text{Diagram 1} \\ \text{Diagram 2} \end{array} = \begin{array}{c} \text{Diagram 3} \\ \text{Diagram 4} \end{array}$$

As with all previously discussed relations, the above relation extends to relations on more points. One way to extend the above relation is to add some number of doubled edges to each side; we call such relations the *simplest binomial relations*. Examples are given in [HMSV2, §4].

1.4. The main theorem. We now state Theorem 1.1 more precisely in the main case of n even and unit weights.

Theorem 1.2 (Main theorem, main case). *For even $n \neq 6$ the ideal I_{1^n} of relations (the kernel of $\text{Sym } R_{1^n}^{(1)} \rightarrow R_{1^n}$) is generated by the simplest binomial relations. The symmetric group \mathfrak{S}_n acts transitively on these relations, so any one of them generates I_{1^n} as an \mathfrak{S}_n -ideal.*

(The ideal I_{1^6} is principal and generated by the Segre cubic relation (1.2) over \mathbb{Z} .)

Remark 1.3. The discussion of [HMSV1, §2.17] explains how to reduce the case of arbitrary weight to the “main case.” Thus as a corollary we have Theorem 1.1, and more precisely, the quadratics are explicitly given by “clumping vertices” (loc. cit.). Thus for the remainder of the paper, we will deal only with this “main case” of $w = 1^n$ and n even.

We essentially conjectured Theorems 1.1 and 1.2 in [HMSV1, §1.5]. We saw the two main theorems of that paper as evidence: first, that a class of relations called the “simple binomial relations,” containing the simplest binomial relations, cuts out the quotient scheme-theoretically, and second, that the ideal is generated by relations of degree at most four. These two results are subsumed by Theorem 1.2 in the main case, and Theorem 1.1 (in its more precise form, Remark 1.3) in general.

Remark 1.4. Given that Theorem 1.2 states that I_{1^n} , for even $n \neq 6$, is generated by a single quadratic up to \mathfrak{S}_n -symmetry, one may wonder whether the same holds when n is odd. A short representation-theoretic argument shows that if n is 5, 7 or 9 then this is indeed the case (see [HMSV2, Fig. 3(h)] for a simple generator with $n = 5$), but if n is odd and at least 11 then the space of quadratic relations is not a cyclic \mathfrak{S}_n -module, so I_{1^n} is not principal as an \mathfrak{S}_n -ideal.

1.5. Representation-theoretic description. Here is a representation-theoretic description of the quadratics (in the main case) in characteristic 0 that is both striking and relevant. The \mathfrak{S}_n -representation on $R_{1^n}^{(1)}$ is irreducible and corresponds to the partition $n/2 + n/2$. The representation $\text{Sym}^2 R_{1^n}^{(1)}$ is multiplicity free and contains those irreducibles corresponding to partitions with at most four parts, all even (Proposition 6.5). The space of quadratic relations is the subspace of $\text{Sym}^2 R_{1^n}^{(1)}$ spanned by those irreducibles corresponding to partitions with *precisely* four parts. Being multiplicity-free, this is necessarily a cyclic \mathfrak{S}_n -module. The reader may wonder why we privilege a particular generator; the answer is that this relation is in some imprecise sense forced upon us by the graphical formalism.

1.6. Outline of proof. We now outline the proof, noting where the arguments are ad hoc or less satisfactory. The challenge is to relate three structures which often operate at cross purposes: the generation of new relations from relations on fewer points; the action of \mathfrak{S}_n on everything; and the graphical description of the algebra, including the use of colored graphs to describe relations.

In §2, we set the stage by giving our preferred description of the invariant ring. We replace the integer n by a finite set L of cardinality n , as this makes many constructions more transparent. In §3–4 we use a Speyer-Sturmfels toric degeneration to get some control on the degrees and types of generators, temporarily breaking the \mathfrak{S}_L -symmetry. We show that the degenerated ring is generated in degree one and that the relations between the degree one generators are generated by quadratic relations and certain explicit cubic relations. In §5 we lift these explicit cubic relations to the original invariant ring. Having deduced that the ideal of relations is generated by quadratics and these particular (“small generalized Segre”) cubics, we are done with the toric degeneration. Our next goal is to show that the particular cubics lie in the ideal Q_L generated by quadratic relations.

In order to take advantage of the \mathfrak{S}_L -action, in §6 we study the tensor powers of the degree one invariants as representations, introducing a useful “partition filtration.” Our results will (for example) allow us to write invariants and relations in terms of highly disconnected graphs, which is the key to our later inductive arguments. The last result of this section is disappointing: it is the only place in the article where computer calculation is used. However, the calculations are quite mild — they concern cubic invariants on six points and amount to simple linear algebra problems in vector spaces of dimension at most 35 — and we feel that a dedicated human being could perform them in a matter of hours.

In §7, we show that for $n \geq 10$ all relations are induced from those on fewer points (precisely, $n - 2$, $n - 4$ or $n - 6$ points), modulo quadratic relations. The cases $n \geq 12$ are direct and structural, but the case $n = 10$ is ad hoc and inelegant because “the graphs are too small” to apply the structural techniques. As a consequence, we find that if the ideal is generated by quadratics for $n - 2$, $n - 4$ and $n - 6$ points then it is for n points as well.

In §8, we show that the ideal is generated by quadratics when $n \neq 6$. This implies Theorem 1.1 by Remark 1.3. Thanks to the previous section, showing generation by quadratics reduces by induction to showing the result for the three “base cases” where n is 8, 10 and 12. We accomplish this by further reducing the 10 and 12 point cases to the 8 point case and then appealing to earlier work for the 8 point case. The reduction from the 10 and 12 point case to the 8 point case is one of the most difficult and least conceptual parts of the paper, so the reader may wish to skip this section on a first reading. One might hope that these results, being finite computations, could be relegated to a computer, but the computations are large enough so that this is not possible with current technology using naive algorithms.

Finally, in §9 we show that the quadratic relations are spanned by the simplest binomial relations using the representation theory of \mathfrak{S}_L , completing the proof of Theorem 1.2.

1.7. The projective coordinate ring of $X^n//G$. The third author has observed that many of the formal concepts in this paper — such as outer multiplication and the simple binomial relations — in fact apply to the study of the projective coordinate ring of $X^n//G$ for any projective variety X with an action of a group G . He has constructed a formalism for dealing with the resulting structures and formulated a few general finiteness conjectures. One of the more surprising realizations is that the graphical formalism discussed

above applies to any variety, in a certain sense. This general point of view may even shed more light on the present case: in this formalism, Theorem 1.2 can be reinterpreted as stating that a certain “ring,” made up of the rings R_{1^n} with varying n , is finitely presented. If such finitely presented “rings” were coherent (a weakening of the noetherian property) then one would immediately obtain universal degree bounds on syzygies, as we have for relations. This work will appear in a forthcoming paper.

1.8. Notation and conventions. We follow some conventions in an attempt to make the notation less onerous. Throughout, L will be a finite set. By an *even set* we mean finite set of even cardinality. Semigroups will be in script, e.g., $\mathcal{G}, \mathcal{S}, \mathcal{R}$. Graphs are denoted by uppercase Greek letters, e.g., Γ, Δ . Trivalent trees will be denoted by Ξ . An edge of a directed (resp. undirected) graph from vertex x to y is denoted \overrightarrow{xy} (resp. xy). In general, S (in various fonts) will refer to constructions involving general directed graphs, and R will refer to regular graphs. We work over \mathbb{Z} (see §1.1) in general.

Acknowledgments. We thank Shrawan Kumar, Chris Manon and Lawrence O’Neil for useful discussions.

2. THE INVARIANT RING R_L

In §2 we define the ring of invariants R_L for a finite set L , and give some of its properties.

2.1. The semi-group \mathcal{G}_L and the rings S_L and R_L . Let L be a finite set. Denote by \mathcal{G}_L the set of directed graphs on L . Give \mathcal{G}_L the structure of a semi-group by defining $\Gamma \cdot \Gamma'$ to be the graph on L whose edge set is the disjoint union of the edge sets of Γ and Γ' . For an element Γ of \mathcal{G}_L we denote the corresponding element of the semi-group algebra $\mathbb{Z}[\mathcal{G}_L]$ by X_Γ . (Readers interested in characteristic 0 may freely replace any occurrence of \mathbb{Z} or $\mathbb{Z}[1/n]$ by \mathbb{Q} throughout.)

For $a, b \in L$ let \overrightarrow{ab} denote the graph in \mathcal{G}_L with a single directed edge from a to b . Clearly \mathcal{G}_L is the free commutative semi-group on the \overrightarrow{ab} , and $\mathbb{Z}[\mathcal{G}_L]$ is the polynomial ring on the $X_{\overrightarrow{ab}}$.

Define the ring S_L as the quotient of $\mathbb{Z}[\mathcal{G}_L]$ by the following three types of relations, described in §1.2.

- *Loop relation:* If Γ has a loop then $X_\Gamma = 0$.
- *Sign relation:* If Γ is obtained from Γ' by reversing the direction of an edge then $X_\Gamma = -X_{\Gamma'}$.
- *Plücker relation:* If a, b, c and d are elements of L , then $X_{\overrightarrow{ab}}X_{\overrightarrow{cd}} = X_{\overrightarrow{ad}}X_{\overrightarrow{cb}} + X_{\overrightarrow{ac}}X_{\overrightarrow{bd}}$.

The sign relation implies the loop relation when 2 is inverted. We still write X_Γ for the image of X_Γ in the ring S_L .

Recall that a graph Γ on L is said to be *regular of degree d* if each vertex of Γ belongs to precisely d edges. Define R_L to be the subgroup of S_L generated by the X_Γ with Γ regular. Clearly R_L is a subring of S_L . Grade R_L by declaring X_Γ to be of degree d if Γ is regular of degree d . If $|L|$ is odd, then every regular graph on L has even degree, so R_L is concentrated in even degrees. (The “first and second main theorems of invariant theory,” mentioned implicitly in §1.2, imply that S_L is the SL_2 -invariant part of the Cox ring of $(\mathbb{P}^1)^L$ and that R_L is the SL_2 -invariant part of the projective coordinate ring of $(\mathbb{P}^1)^L$ with respect to the line bundle $\mathcal{O}(1)^{\boxtimes L}$.)

These constructions are functorial in L : A map of sets $\phi : L \rightarrow L'$ induces a homomorphism $\mathcal{G}_L \rightarrow \mathcal{G}_{L'}$ of semi-groups and thus a homomorphism $\mathbb{Z}[\mathcal{G}_L] \rightarrow \mathbb{Z}[\mathcal{G}_{L'}]$ of semi-group rings. This ring homomorphism respects the sign and Plücker relations and so induces a ring homomorphism $S_L \rightarrow S_{L'}$. If the fibers of L all have the same cardinality then R_L is mapped into $R_{L'}$, and we thus obtain a map $R_L \rightarrow R_{L'}$. As a special case, we see that $\mathfrak{S}_L = \mathrm{Aut}(L)$ acts on R_L .

2.2. Translation from directed graphs to undirected graphs. To avoid confusion with signs, it will often be convenient to translate from directed graphs to undirected graphs. Let L be an even set. We denote by \mathcal{M}_L the set of directed matchings on L . An *orientation* on L is defined as a map $\epsilon : \mathcal{M}_L \rightarrow \{\pm 1\}$ satisfying $\epsilon(\sigma\Gamma) = \mathrm{sgn}(\sigma)\epsilon(\Gamma)$ for $\sigma \in \mathfrak{S}_L$ and $\Gamma \in \mathcal{M}_L$. There are two orientations on L .

Fix an orientation of L . For an *undirected* matching Γ we put $Y_\Gamma = \epsilon(\tilde{\Gamma})X_{\tilde{\Gamma}}$ where $\tilde{\Gamma}$ is any directed matching with underlying undirected matching Γ . The Y_Γ span the space $R_L^{(1)}$ and satisfy the loop relation and the “undirected Plücker relation”

$$Y_{\Gamma_1} + Y_{\Gamma_2} + Y_{\Gamma_3} = 0,$$

whenever Γ_2 and Γ_3 are obtained by modifying two edges of Γ_1 appropriately. We often prefer to work with the Y_Γ instead of the X_Γ since there are no directions to keep track of. However, one must keep in mind that the action of the symmetric group on the Y_Γ is twisted by the sign character from the most obvious

action: $\sigma Y_\Gamma = \text{sgn}(\sigma) Y_{\sigma\Gamma}$ for $\sigma \in \mathfrak{S}_L$. We often use the Y_Γ without explicitly mentioning the choice of an orientation.

2.3. Kempe’s generation theorem. The purpose of this paper is to give a presentation for the ring R_L . To do this we must first find a set of generators. This problem was solved (for the main case) by Kempe [Ke].

Theorem 2.1 (Kempe). *Let L be an even set. Then the ring R_L is generated in degree one. Equivalently, the Y_Γ (or X_Γ) with Γ a matching generate R_L .*

Proof. We recall the proof of [HMSV1, Theorem 2.3], which is simpler than Kempe’s original proof, to motivate later arguments. For any regular graph Γ on L , we express Y_Γ as a polynomial in elements of the form $Y_{\Gamma'}$ with Γ' a matching. Partition L arbitrarily into two sets of equal cardinality, one called “positive” and the other “negative.” We then have three types of edges: positive (both vertices positive), negative (both negative) and neutral (one positive and one negative). After applying the Plücker relation to a positive and a negative edge, one is left with only neutral edges:

$$\begin{array}{c} + & - \\ | & | \\ + & - \end{array} = \begin{array}{c} + & - \\ \diagdown & / \\ + & - \end{array} + \begin{array}{c} + & - \\ \hline \hline \\ + & - \end{array}$$

(We have neglected signs in the above identity.) As Γ is regular, it has a positive edge if and only if it has a negative one. Thus by repeatedly applying Plücker relations to positive and negative edges we end up with an expression $Y_\Gamma = \sum \pm Y_{\Gamma_i}$ where the Γ_i have only neutral edges, and are hence bipartite. Hall’s marriage theorem states that in a regular bipartite graph one can find a matching. Thus each Γ_i can be factored into matchings, which completes the proof. \square

2.4. Kempe’s basis theorem. Fix an embedding of L into the unit circle in the plane. We say that a graph Γ on L is *non-crossing* if no two of its edges cross when drawn as chords. The following well-known theorem of Kempe (also from [Ke]) will be used in the proofs of Proposition 3.1 and Lemma 9.3.

Theorem 2.2 (Kempe). *The X_Γ with Γ non-crossing span S_L . The only linear relations among these elements are the sign and loop relations. Thus if one chooses for each undirected loop-free non-crossing graph a direction on the edges then the corresponding X_Γ form a basis for S_L . The same is true for R_L if one considers regular non-crossing graphs.*

In fact, there is a procedure called the *straightening algorithm* which expresses an arbitrary X_Γ in terms of the non-crossing basis. The algorithm is simple: take any pair of edges in Γ which cross and Plücker them. The algorithm terminates because the total lengths of the edges in each the two graphs resulting from a Plücker operation is less than that in the original graph. This nearly proves the theorem; for details see [HMSV1, Propositions 2.5, 2.6].

2.5. Some definitions. We now define some notation that will be used constantly:

- R_L is the ring of invariants, as defined above.
- V_L is the first graded piece $R_L^{(1)}$ of R_L ; it is spanned by the Y_Γ with Γ a matching.
- I_L is the ideal of relations, that is, the kernel of the map $\text{Sym}(V_L) \rightarrow R_L$.
- Q_L is the ideal of $\text{Sym}(V_L)$ generated by $I_L^{(2)}$; it is a sub-ideal of I_L .

3. THE TORIC DEGENERATIONS $\text{gr}_\Xi S_L$ AND $\text{gr}_\Xi R_L$

In §3, we discuss toric degenerations of the rings S_L and R_L . These were first described in [SS], and one was used in [HMSV1]. By a “toric ring” we mean a ring isomorphic to a semi-group algebra, where the semi-group is the set of lattice points in a strictly convex rational polyhedral cone; by a “toric degeneration” of a ring we mean a toric ring obtained as the associated graded of a filtration on the original ring. The main points of §3 are the following:

- (1) To each trivalent tree Ξ with leaf set L we give a toric degeneration of the rings S_L and R_L , denoted $\text{gr}_\Xi S_L$ and $\text{gr}_\Xi R_L$ respectively.
- (2) We give a presentation of the ring $\text{gr}_\Xi S_L$.

- (3) We discuss the theory of weightings on trivalent trees.
- (4) We identify the rings $\text{gr}_{\Xi} S_L$ and $\text{gr}_{\Xi} R_L$ with semi-group algebras of “weightings.”

3.1. Trivalent trees. By a *trivalent tree* we mean a connected undirected graph Ξ , without cycles, all of whose vertices have valence one or three. We call vertices of valence three *trinodes*, and vertices of valence one *leaves*. We say leaves x and y form a *matched pair* if they share a neighbor (trinode). We say that a trivalent tree is *matched* if it has more than four vertices and every leaf belongs to a (necessarily unique) matched pair. A matched trivalent tree necessarily has an even number of leaves.




FIGURE 1. A trivalent tree. Vertices c, d and f are trinodes; the rest are leaves. The vertices a and b form a matched pair, as do the vertices g and h . The tree is not matched because e does not belong to a matched pair.

3.2. The toric rings $\text{gr}_{\Xi} S_L$ and $\text{gr}_{\Xi} R_L$. Let Ξ be a trivalent tree with leaf set L . For a graph Γ on L , define the *level* of Γ (relative to Ξ) as

$$\text{lev}_{\Xi} \Gamma = \sum_{\overrightarrow{ab}} (\text{the distance from } a \text{ to } b \text{ in } \Xi)$$

where the sum is over the edges of Γ , and distance is the number of edges in the “geodesic.” Clearly lev_{Ξ} induces a semi-group morphism $\text{lev}_{\Xi} : \mathcal{G}_L \rightarrow \mathbb{Z}_{\geq 0}$. Define an increasing filtration F_{Ξ} on $\mathbb{Z}[\mathcal{G}_L]$ by letting $F_{\Xi}^i \mathbb{Z}[\mathcal{G}_L]$ be the subspace of $\mathbb{Z}[\mathcal{G}_L]$ spanned by the X_{Γ} with $\text{lev}_{\Xi} \Gamma \leq i$. (The notation lev_{Ξ} will not be used further.) Let $F_{\Xi}^i S_L$ be the image of $F_{\Xi}^i \mathbb{Z}[\mathcal{G}_L]$ under the surjection $\mathbb{Z}[\mathcal{G}_L] \rightarrow S_L$, giving a filtration of the ring S_L . Let $\text{gr}_{\Xi} S_L$ denote the associated graded ring. We will show that $\text{gr}_{\Xi} S_L$ is a toric ring (Proposition 3.3). For a graph Γ of level n , let \overline{X}_{Γ} denote the image of X_{Γ} in $F_{\Xi}^n S_L / F_{\Xi}^{n-1} S_L$. Clearly the \overline{X}_{Γ} span $\text{gr}_{\Xi} S_L$.

Let $F_{\Xi}^i R_L$ be the filtration on R_L induced from its inclusion into S_L . Let $\text{gr}_{\Xi} R_L$ be the associated graded ring. It is naturally the subring of $\text{gr}_{\Xi} S_L$ spanned by the \overline{X}_{Γ} for which Γ is regular.

3.3. Presentation of the ring $\text{gr}_{\Xi} S_L$. Let J_{Ξ} denote the ideal in $\mathbb{Z}[\mathcal{G}_L]$ generated by the following (cf. §2.1):

- *Loop relation:* If Γ has a loop then $X_{\Gamma} = 0$.
- *Sign relation:* If Γ' is obtained from Γ by reversing the direction of an edge then $X_{\Gamma} = -X_{\Gamma'}$.
- *Toric Plücker relation:* If a, b, c and d are elements of L satisfying:

$$(3.1) \quad \begin{aligned} & \text{the path from } a \text{ to } b \text{ in } \Xi \text{ meets the path from } c \text{ to } d, \\ & \text{and the path from } a \text{ to } c \text{ meets the path from } b \text{ to } d \text{ (see Figure 2)} \end{aligned}$$

$$\text{then } X_{\overrightarrow{ab}} X_{\overrightarrow{cd}} = X_{\overrightarrow{ac}} X_{\overrightarrow{bd}}.$$

(The notation J_{Ξ} is only used in §3.3.)



FIGURE 2. The toric Plücker relation $X_{\overrightarrow{ab}} X_{\overrightarrow{cd}} = X_{\overrightarrow{ac}} X_{\overrightarrow{bd}}$. Both pairs of geodesics overlap on the horizontal edge of the trivalent tree Ξ . Note that \overrightarrow{ad} and \overrightarrow{bc} do not overlap.

The purpose of this section is to prove the following, used in §3.5 to identify $\text{gr}_{\Xi} S_L$ with a semi-group algebra of “weightings.”

Proposition 3.1. *The map $\mathbb{Z}[\mathcal{G}_L] \rightarrow \text{gr}_\Xi S_L$ given by $X_\Gamma \mapsto \overline{X}_\Gamma$ is surjective with kernel J_Ξ .*

The kernel I of the map $\mathbb{Z}[\mathcal{G}_L] \rightarrow S_L$ is generated by the sign and Plücker relations. The kernel of the map $\mathbb{Z}[\mathcal{G}_L] \rightarrow \text{gr}_\Xi S_L$ is the ideal generated by the leading terms of all elements of I . In general, of course, this is not the same as the ideal generated by the leading terms of a generating set of I . Proposition 3.1 says that in this situation, however, this is the case.

Proof. The map is clearly surjective and contains the sign relation in its kernel. We now check that the toric Plücker relation lies in its kernel as well. Let a, b, c and d belong to L and satisfy (3.1). The equation

$$X_{\overrightarrow{ab}} X_{\overrightarrow{cd}} = X_{\overrightarrow{ac}} X_{\overrightarrow{bd}} + X_{\overrightarrow{ad}} X_{\overrightarrow{bc}}$$

holds in S_L (the normal Plücker relation). The two graphs $\overrightarrow{ab} \cdot \overrightarrow{cd}$ and $\overrightarrow{ac} \cdot \overrightarrow{bd}$ have the same level, say n , since when drawn in Ξ they use the same edges with the same multiplicity (see Figure 2). The remaining graph $\overrightarrow{ad} \cdot \overrightarrow{bc}$ has level less than n (again, see Figure 2). Thus all terms in the above relation lie in $F_\Xi^n S_L$. Reducing modulo $F_\Xi^{n-1} S_L$ we obtain

$$\overline{X}_{\overrightarrow{ab}} \overline{X}_{\overrightarrow{cd}} = \overline{X}_{\overrightarrow{ac}} \overline{X}_{\overrightarrow{bd}}$$

which shows that the toric Plücker relation lies in J_Ξ .

We now show that non-crossing graphs span $\mathbb{Z}[\mathcal{G}_L]/J_\Xi$. Embed L into the unit circle in the plane in such a way that Ξ can be drawn inside the circle without any crossings. Let Γ be a graph on L which contains crossing edges \overrightarrow{ab} and \overrightarrow{cd} . The paths ab and cd then meet in Ξ . The same reasoning as in the non-toric case now applies: applying the toric Plücker relation to this pair of edges yields an identity $\overline{X}_\Gamma = \overline{X}_{\Gamma'}$ where the total lengths of edges in Γ' is less than that of Γ (here length is computed as distance in the plane, not the trivalent tree). Continuing in this manner, we get an expression $\overline{X}_\Gamma = \overline{X}_{\Gamma'}$ where Γ' is non-crossing.

It is now formal to conclude that $\mathbb{Z}[\mathcal{G}_L]/J_\Xi \rightarrow \text{gr}_\Xi S_L$ is an isomorphism. We elaborate on this. Choose a set Z of directed non-crossing graphs such that for each undirected non-crossing graph Γ there is a unique way to direct the edges of Γ such that the resulting graph belongs to Z . The previous paragraph shows that if Γ is any graph of level n we can find $\Gamma' \in Z$ such that $X_\Gamma = \pm X_{\Gamma'} + Y$ holds in S_L , where $Y \in F_\Xi^{n-1} S_L$. Applying this result to Y repeatedly, we find that the X_Γ with Γ in Z and level at most n form a basis of $F_\Xi^n S_L$ (we already know they are linearly independent). It thus follows that the \overline{X}_Γ with Γ in Z are linearly independent in $\text{gr}_\Xi S_L$. Since the surjection $\mathbb{Z}[\mathcal{G}_L]/J_\Xi \rightarrow \text{gr}_\Xi S_L$ takes a spanning set to a set of linearly independent vectors, it must be an isomorphism. \square

3.4. Weightings. A *weighting* ξ on a trivalent tree Ξ is an assignment of a non-negative integer to each edge of Ξ . Define the *weight triple* of a trinode v of Ξ to be the weights of the three edges connected to v . (We write it as an ordered triple even though it is not ordered.) Consider the following two equivalent conditions on weight triples (a, b, c) :

- (W1) The triple (a, b, c) satisfies the triangle inequalities (a, b and c can form the sides of a triangle) and $a + b + c$ is even.
- (W2) There exists a triple (x, y, z) of non-negative integers such that $a = x + y$, $b = x + z$ and $c = y + z$.

Note that the triple (x, y, z) is uniquely determined by (a, b, c) . We say that a weighting is *admissible* if the weight triple at each trinode satisfies these conditions.

Let Γ be an undirected graph on the leaves of Ξ . We say that an edge e of Ξ *meets* an edge \overline{ij} of Γ if e occurs in the geodesic joining i and j . We define a weighting ξ_Γ of Ξ by assigning to an edge of Ξ the number of edges of Γ which it meets. We call ξ_Γ the weighting of Ξ *associated* to Γ (see Figure 3).

Define (*toric*) *Plücker equivalence* to be the equivalence relation \sim on the semi-group of undirected graphs $\mathcal{G}_L^{\text{un}}$ on L generated by the following two conditions:

- Given a, b, c, d in L satisfying (3.1), $\overline{ab} \cdot \overline{cd} \sim \overline{ac} \cdot \overline{bd}$.
- If $\Gamma \sim \Gamma'$ and Γ'' is any graph then $\Gamma \cdot \Gamma'' \sim \Gamma' \cdot \Gamma''$.

The following result gives a correspondence between graphs and weightings.

Proposition 3.2. *Associating a weighting of Ξ to a graph induces a bijection between the Plücker classes of undirected graphs on the leaves of Ξ and the admissible weightings of Ξ .*





FIGURE 3. A graph on the leaves of a trivalent tree and its associated weighting.

Proof. Let Γ be a graph and let ξ be its associated weighting. We first show that ξ is admissible. Let v be a trinode of Ξ and let (a, b, c) be its weight triple. Consider the picture




Here a is the number of edges of Γ that meet edge \overline{pv} ; b and c are defined similarly. We let x be the number of edges of Γ which, when drawn as geodesics on Ξ , go through both p and q ; we let y and z be the analogous quantities. Clearly $a = x + y$, $b = x + z$, and $c = y + z$. Thus we have shown that the weighting ξ satisfies condition (W2) at each trinode and is therefore admissible.

If we apply a toric Plücker relation to Γ then its associated weighting does not change: the two pairs of edges in the toric Plücker relation contain the same edges when drawn as geodesics. Thus associating a weighting to a graph yields a well-defined map from the Plücker classes of graphs to the set of admissible weightings. We now show that this map is bijective.

We first prove that it is surjective. We are given an admissible weighting ξ and we must produce a graph Γ such that ξ is its associated weighting. To do this it suffices to prove the following: given an admissible weighting ξ there exist two leaves i and j of Ξ such that when the geodesic joining i and j is subtracted from ξ the resulting weighting is still admissible. For, if this is the case, then we can let \overline{ij} be an edge of Γ , subtract the geodesic joining i and j from ξ and proceed by induction.

Thus let ξ be an admissible weighting of Ξ . Let i be any leaf of Ξ for which ξ does not vanish on the edge containing i . We produce the vertex j by the following greedy algorithm. Put $v_0 = i$ and let v_1 be the unique trinode connected to i . Assume now that we have defined v_0 through v_k . If v_k is a leaf then stop and put $j = v_k$. Otherwise v_k is connected to two vertices other than v_{k-1} . Let v_{k+1} be the one for which the corresponding edge has higher weight; if the two edges have the same weight then pick v_{k+1} arbitrarily. We have thus produced a pair of leaves i and j ; note that the v_k are the trinodes in geodesic joining i and j . We must show that when this geodesic is subtracted from ξ the resulting weighting is still admissible.

Let ℓ be lie strictly between 0 and k so that v_ℓ is a trinode. We have the picture




By definition of v_{l+1} we have $b \geq c$. We know (a, b, c) satisfies (W1). Clearly $(a - 1, b - 1, c)$ still satisfies the parity condition. We must show that it still satisfies the triangle inequalities, which amounts to proving $c \leq a + b - 2$. This inequality could fail in two ways: 1) $a = 0$ and c equals b or $b - 1$; or 2) $a = 1$ and $b = c$. The first case is ruled out by the way we selected v_{l-1} and an easy induction argument. The second

case is ruled out since $a + b + c$ is even. This proves that the greedy algorithm indeed works and completes the proof that our map is surjective. Note that in building the graph from the weighting there are many arbitrary choices.

We now prove that the map is injective, i.e., that if Γ and Γ' have the same weighting ξ then they are Plücker equivalent. To show this we show that Γ is Plücker equivalent to any of the graphs constructed out of ξ by using the greedy algorithm. It suffices to prove that if v_0, \dots, v_k is a sequence coming out of the greedy algorithm then we can apply toric Plücker relations to Γ so that v_0 and v_k are connected in Γ . For then we may remove this edge from Γ and the corresponding path in ξ and proceed by induction.

Thus let v_0, \dots, v_k come out of the greedy algorithm. We prove by induction on ℓ that Γ is Plücker equivalent to a graph containing an edge which passes through v_0 and v_ℓ (we say that an edge of Γ passes through two vertices of Ξ if its corresponding geodesic does). This is clear for $\ell = 1$. Thus assume it is true for ℓ and we show that it is true for $\ell + 1$. We have the picture



Here x_i is the unique vertex connected to v_i besides v_{i-1} and v_{i+1} . We have assumed that Γ contains an edge e passing through v_0 and v_ℓ ; we must show that Γ is Plücker equivalent to a graph containing an edge passing through v_0 and $v_{\ell+1}$. Now, e itself either passes through $v_{\ell+1}$ or x_ℓ . In the former case we are done. Thus we may assume that e passes through x_ℓ .

By the definition of the greedy algorithm we have

$$\xi(\overline{v_\ell x_\ell}) \leq \xi(\overline{v_\ell v_{\ell+1}}), \quad \xi(\overline{v_i v_{i+1}}) \neq 0.$$

Thus there exists an edge of Γ passing through v_ℓ and $v_{\ell+1}$. If every edge which passed through v_ℓ and $v_{\ell+1}$ also passed through x_ℓ then, by the inequality, every edge which passed through x_ℓ and v_ℓ would also pass through $v_{\ell+1}$; it would follow that no edge could pass through $v_{\ell-1}$ and v_ℓ . However, this would imply $\xi(\overline{v_{\ell-1} v_\ell}) = 0$, a contradiction. Thus there exists an edge e' of Γ which passes through $v_{\ell+1}$ and v_ℓ but *not* through x_ℓ . If v_0 is a vertex of e' then we are done. Otherwise, applying the toric Plücker relation to e and e' yields a graph containing an edge passing through v_0 and $v_{\ell+1}$. This completes the proof. \square

3.5. The rings $\text{gr}_\Xi S_L$ and $\text{gr}_\Xi R_L$ as semi-group algebras of weightings. Let Ξ be a trivalent tree with leaf set L . Embed L into the unit circle in such a way that Ξ can be drawn inside the unit circle without crossings. Choose a total order on L which is compatible with its embedding into the circle in the sense that if $a \leq b \leq c$ then one encounters b when traveling clockwise from a to c . For $a, b \in L$ define ϵ_{ab} to be 1 if $a < b$, -1 if $a > b$ and 0 if $a = b$. For a directed graph Γ on L define ϵ_Γ to be the product of the ϵ_{ab} over the edges \overrightarrow{ab} of Γ . We write Γ^{un} for the undirected graph associated to Γ . Finally, let \mathcal{S}_Ξ denote the set of admissible weights on Ξ . It is a semi-group since the sum of two admissible weightings is again admissible. We can now prove:

Proposition 3.3. *There is a unique isomorphism of rings $\text{gr}_\Xi S_L \rightarrow \mathbb{Z}[\mathcal{S}_\Xi]$ mapping \overline{X}_Γ to $\epsilon_\Gamma \xi_{\Gamma^{\text{un}}}$.*

Proof. We define an auxiliary ring by modifying the sign relation in the presentation of $\text{gr}_\Xi S_L$ given in §3.3. Define the ideal J_Ξ^{un} of $\mathbb{Z}[\mathcal{G}_L]$ by the following types of relations:

- *Loop relation:* If Γ has a loop then $X_\Gamma = 0$.
- *Modified sign relation:* If Γ' is obtained from Γ by reversing the direction of an edge then $X_\Gamma = X_{\Gamma'}$.
- *Toric Plücker relation:* As in §3.3.

Denote the image of X_Γ in $\mathbb{Z}[\mathcal{G}_L]/J_\Xi^{\text{un}}$ by $\overline{X}_\Gamma^{\text{un}}$. Note that $\overline{X}_\Gamma^{\text{un}}$ makes sense for an *undirected* graph Γ , and that if Γ and Γ' are Plücker equivalent undirected graphs then the toric Plücker relation implies $\overline{X}_\Gamma^{\text{un}} = \overline{X}_{\Gamma'}^{\text{un}}$. By comparing the definition of J_Ξ^{un} to the presentation of $\text{gr}_\Xi S_L$ given in Proposition 3.1, we find that the map

$$\text{gr}_\Xi S_L = \mathbb{Z}[\mathcal{G}_L]/J_\Xi \rightarrow \mathbb{Z}[\mathcal{G}_L]/J_\Xi^{\text{un}}, \quad \overline{X}_\Gamma \mapsto \epsilon_\Gamma \overline{X}_\Gamma^{\text{un}}$$

is well-defined and an isomorphism. From the equivalence of weightings and Plücker classes of undirected graphs given in Proposition 3.2, we find that the map

$$\mathbb{Z}[\mathcal{G}_L]/J_\Xi^{\text{un}} \rightarrow \mathbb{Z}[\mathcal{S}_\Xi], \quad \overline{X}_\Gamma^{\text{un}} \mapsto \xi_{\Gamma^{\text{un}}}$$

is a well-defined isomorphism. The proposition now follows. □

We translate this result to the regular case. Call a weighting ξ on Ξ *regular of degree d* if for each leaf v we have $\xi(e_v) = d$, where e_v is the unique edge containing v . Let \mathcal{R}_Ξ be the semi-group of admissible regular weightings on Ξ . We then have:

Corollary 3.4. *There is a unique isomorphism of graded rings $\text{gr}_\Xi R_L \rightarrow \mathbb{Z}[\mathcal{R}_\Xi]$ which takes \overline{X}_Γ to $\epsilon_\Gamma \cdot \xi_{\Gamma^{\text{un}}}$.*

Proposition 3.3 and Corollary 3.4 show that $\text{gr}_\Xi S_L$ and $\text{gr}_\Xi R_L$ are semi-group algebras and therefore toric rings. (It is not difficult to see that \mathcal{S}_Ξ and \mathcal{R}_Ξ are the set of lattice points in a strictly convex rational polyhedral cone, but this will not be of importance to us.)

3.6. Reduced weightings. We close §3 with a discussion of reduced weightings, which we use in §4–5. Let Ξ be a trivalent tree. By a *reduced weighting* on Ξ we simply mean a weighting on Ξ — the terms are synonymous but used to distinguish the usage of “admissible.” We say that a reduced weighting is *admissible* if it satisfies the triangle inequality (as in (W1)) at each trinode; the parity condition is not enforced. We say that a reduced weighting ξ is *regular of degree $\leq d$* if $\xi(e_v) \leq d$ for all leaves v , where e_v is the edge meeting leaf v . Let $\overline{\mathcal{R}}_\Xi$ be the set of all ordered pairs (ξ, d) with d a non-negative integer and ξ an admissible reduced weighting which is regular of degree $\leq d$. We define the *degree* of (ξ, d) to be d .

Let Ξ be a matched trivalent tree and let Ξ^- be the trivalent tree obtained by deleting the leaves of Ξ and the edges that they touch. We call Ξ^- the *truncation* of Ξ .

Proposition 3.5. *With notation as above, there is a canonical isomorphism of semi-groups $\mathcal{R}_\Xi \rightarrow \overline{\mathcal{R}}_{\Xi^-}$ preserving degree. The image of a weighting ξ on Ξ is the pair (ξ', d) where d is the degree of ξ and ξ' is the weighting on Ξ^- given by $\xi'(x) = \frac{1}{2}\xi(x)$.*

The proof is easy. See Figure 4 for an illustration.




FIGURE 4. An illustration of Proposition 3.5. The admissible weighting ξ on the left is regular of degree 1, the associated reduced admissible weighting ξ^- on the right is regular of degree ≤ 1 .


4. THE TORIC IDEAL IS GENERATED BY QUADRATICS AND TORIC GENERALIZED SEGRE CUBICS

We have described a family of toric degenerations of R_L , depending on a choice of trivalent tree. The purpose of §4 is to choose a specific family of trees (the Y-trees) to ensure the degenerated ring (i) is generated in degree one, (ii) has relations generated in degrees two and three, and (iii) is such that there is a precise description of the degree three relations:

Theorem 4.1. *Let Ξ be a Y-tree (defined in §4.1). Then $\text{gr}_\Xi R_L$ is generated in degree one, and the relations between degree one elements are generated by quadratic relations and the generalized toric Segre cubic relations (defined in §4.4).*

(Manon [M] independently solved the presentation problem for a large class of weighted trivalent trees, including this one; however, we will need the form of (iii).)

4.1. The Y- and caterpillar trees. For $r \geq 3$, define the r th Y-tree as:



There are r “Y’s” in the tree, and $2r$ leaves. We call the vertex at i the i th *base vertex*. We call an edge between two base vertices a *base edge*. By a *stalk* we mean one of the internal edges in one of the Y’s. By

the i th stalk, for $2 \leq i \leq r-1$, we mean the one above the i th base vertex. We call the two remaining stalks at either end of the tree the 1st and r th stalk. The Y-tree is a matched tree (§3.1). Sometimes we bend the first and last horizontal edges so that all the Y's are in a row, as in Figure 4. The r th *caterpillar tree* is the truncation of the r th Y-tree:



We use the same terminology (base edges, stalks, etc.) for caterpillar trees. Note that the third caterpillar tree is just a trinode.

4.2. Generators of $\text{gr}_{\Xi} R_L$. We now prove the first half of Theorem 4.1.

Proposition 4.2. *Let Ξ be a Y-tree. Then the ring $\text{gr}_{\Xi} R_L$ is generated in degree one.*

Proof. (This proof essentially applies to any matched tree.) We follow the spirit of our proof of Kempe's theorem (Theorem 2.1). Assign to each leaf of Ξ a sign in such a way that in each matched pair of leaves one leaf is positive and the other is negative. Suppose Γ is a regular undirected graph of degree k . Each edge of Γ is then either positive, negative or neutral. We will show that Γ is Plücker equivalent (§3.4) to a graph with only neutral edges. As in the proof of Theorem 2.1, this implies that it is a product of matchings, expressing \overline{X}_{Γ} as a product of degree one elements. This proof is more difficult than that of Theorem 2.1 because we are now only allowed to apply the Plücker relation to pairs of edges of Γ which meet (i.e., whose geodesics meet) in the tree Ξ .

Step 1: Given a positive and negative edge of Γ which meet in Ξ , apply the toric Plücker relation. In the resulting graph the two new edges are both neutral and so the number of non-neutral edges has decreased. We can thus continue this process until we have reached a graph which is Plücker equivalent to Γ and contains no overlapping positive and negative edges.

Step 2: Suppose there are still non-neutral paths, but no two of opposite parity meet each other. Let e be the leftmost base edge of Ξ which separates non-neutral edges of Γ of opposite parity and let e' be the edge to the left of e . The edge e' will be a base edge unless e is the first base edge, in which case it will be the first stalk. All non-neutral edges of Γ which are entirely to the left of e are of the same type, say positive. All edges of Γ which pass through e are neutral.

There is a positive edge \overline{ab} of Γ which contains e' . We claim there is an edge \overline{cd} of Γ which contains e and where d is positive and to the right of e . Assuming the claim (which we prove below), then \overline{ab} and \overline{cd} must meet (either at e' or at the stalk between e' and e); now apply the toric Plücker relation to this pair. The resulting positive path contains e . This brings the leftmost cluster of positive paths closer to the negative paths. By continuing this process, we will eventually cause the two clusters to meet, at which point we return to Step 1 and reduce the number of non-neutral edges. Continuing in this manner, we will eventually remove all non-neutral edges.

Proof of claim: We now prove the claim that there must exist an edge \overline{cd} of Γ containing e and for which d is positive and to the right of e . Suppose there are $2m$ leaves to the left of e . Let A (resp. B , C) be the number of positive (resp. negative, neutral) edges of Γ entirely to the left of e . Let D (resp. E) be the number of neutral edges containing e for which the positive (resp. negative) leaf is to the left of e . Then

$$2A + C + D = mk, \quad 2B + C + E = mk$$

as the first (resp. second) counts k times each positive (resp. negative) leaf to the left of e , and there are m such leaves. We thus find $2A + D = 2B + E$ and as $B = 0$ and $A \neq 0$ by assumption we conclude $E \neq 0$, as claimed. \square

4.3. Relations in a semi-group algebra. We now turn our attention to relations in $\text{gr}_{\Xi} R_L$. We begin with a general discussion of relations in a semi-group algebra. Let \mathcal{R} be a semi-group equipped with a homomorphism $\text{deg} : \mathcal{R} \rightarrow \mathbb{Z}_{\geq 0}$, so that the semi-group algebra $\mathbb{Z}[\mathcal{R}]$ is graded. Assume that \mathcal{R} is generated by $\mathcal{V} = \text{deg}^{-1}(1)$, so that the natural map $\text{Sym}(\mathbb{Z}\langle \mathcal{V} \rangle) \rightarrow \mathbb{Z}[\mathcal{R}]$ is a surjection (here $\text{Sym}(\mathbb{Z}\langle \mathcal{V} \rangle)$ is the polynomial ring in indeterminates \mathcal{V} while $\mathbb{Z}[\mathcal{R}]$ is the semi-group algebra of \mathcal{R}). Call the kernel I of this surjection the *ideal of relations* of $\mathbb{Z}[\mathcal{R}]$.

Let $\underline{\xi} = (\xi_1, \dots, \xi_n)$ and $\underline{\xi}' = (\xi'_1, \dots, \xi'_n)$ be two elements of \mathcal{V}^n such that $\sum \xi_i = \sum \xi'_i$. We then write $\underline{\xi} \sim \underline{\xi}'$ and say that $\underline{\xi}$ and $\underline{\xi}'$ are *related*. For an element ξ of \mathcal{V} let $[\xi]$ denote the corresponding element

of $\text{Sym}(\mathbb{Z}\langle \mathcal{V} \rangle)$; for an element $\underline{\xi}$ of \mathcal{V}^n let $[\underline{\xi}]$ denote the monomial $[\xi_1] \cdots [\xi_n]$. Given related $\underline{\xi}$ and $\underline{\xi}'$ the element $[\underline{\xi}] - [\underline{\xi}']$ belongs to I . We call such relations *binomial relations*. One easily verifies that I is generated by binomial relations.

We say that a relation $\underline{\xi} \sim \underline{\xi}'$ has *degree* $\leq k$ if $\xi_i = \xi'_i$ holds for all but k indices. We say that $\underline{\xi} \sim \underline{\xi}'$ has *essentially degree* $\leq k$ if there exists a sequence of relations

$$\underline{\xi} = \underline{\xi}^{(0)} \sim \underline{\xi}^{(1)} \sim \dots \sim \underline{\xi}^{(p)} = \underline{\xi}'$$

for which each $\underline{\xi}^{(i)} \sim \underline{\xi}^{(i+1)}$ has degree $\leq k$. This notation is only relevant to us when k is 2 or 3 and we then use the terms “essentially quadratic” and “essentially cubic.” If $\underline{\xi} \sim \underline{\xi}'$ has essential degree $\leq k$ then the element $[\underline{\xi}] - [\underline{\xi}']$ of I lies in the ideal generated by the k th graded piece of I .

Note that we consider our tuples as ordered, so that $(\xi_1, \xi_2) = (\xi_2, \xi_1)$ constitutes a non-trivial relation. However, as the symmetric group is generated by transpositions, it follows that if $\underline{\xi}'$ is a permutation of the tuple $\underline{\xi}$ then the relation $\underline{\xi} \sim \underline{\xi}'$ is essentially quadratic.

4.4. The toric generalized Segre cubic relation. We have the following relation between degree one reduced weightings on the third caterpillar tree:

$$(4.1) \quad \begin{array}{c} 0 \\ | \\ -1-1- \end{array} \quad \begin{array}{c} 1 \\ | \\ -0-1- \end{array} \quad \begin{array}{c} 1 \\ | \\ -1-0- \end{array} = \begin{array}{c} 1 \\ | \\ -1-1- \end{array} \quad \begin{array}{c} 1 \\ | \\ -1-1- \end{array} \quad \begin{array}{c} 0 \\ | \\ -0-0- \end{array}$$

(We will omit $+$ signs in such equations, interpreting them as binomial relations in $\mathbb{Z}[\mathcal{R}]$.) One obtains this relation by converting the usual graphical Segre cubic (1.2) into a relation between weightings on the third Y-tree and then passing to the associated reduced weighting on the third caterpillar tree. One may verify by hand that this single relation generates all relations among admissible reduced weightings on the third caterpillar tree.

We now introduce a class of toric relations that generalize (4.1). We call a reduced weighting of Ξ^- of degree ≤ 1 a *reduced matching*. Recall that the set \mathcal{V}_Ξ of reduced matchings generates \mathcal{R}_Ξ (Proposition 4.2). Let X, Y and Z be reduced matchings on the r th caterpillar tree such that X and Y take value 1 on the r th stalk and Z takes value 0 on the r th stalk. Let X', Y' and Z' be reduced matchings on the s th caterpillar tree such that X' and Y' take value 1 on the first stalk and Z' takes value 0 on the first stalk. We then have the relation

$$= \begin{array}{c} 1 \\ | \\ \boxed{X} - 1 - 1 - \boxed{X'} \end{array} \quad \begin{array}{c} 1 \\ | \\ \boxed{Y} - 1 - 1 - \boxed{Y'} \end{array} \quad \begin{array}{c} 0 \\ | \\ \boxed{Z} - 0 - 0 - \boxed{Z'} \end{array} \\ = \begin{array}{c} 1 \\ | \\ \boxed{X} - 1 - 0 - \boxed{Z'} \end{array} \quad \begin{array}{c} 1 \\ | \\ \boxed{Z} - 0 - 1 - \boxed{X'} \end{array} \quad \begin{array}{c} 0 \\ | \\ \boxed{Y} - 1 - 1 - \boxed{Y'} \end{array}$$

We call these *toric generalized Segre cubic* relations.

4.5. The type vector of a triple. Let $\underline{\xi} = (\xi_1, \xi_2, \xi_3)$ be a triple of reduced matchings on the r th caterpillar tree. Define the *type* of $\underline{\xi}$ at the i th base vertex to be one of A, B or \emptyset , as follows. We call $\underline{\xi}$ type A at i if it looks like

$$\begin{array}{c} 1 \\ | \\ -1-1- \end{array} \quad \begin{array}{c} 1 \\ | \\ -1-1- \end{array} \quad \begin{array}{c} 0 \\ | \\ -0-0- \end{array}$$

at the i th vertex (the order of the triple is irrelevant). We call $\underline{\xi}$ type B at i if it looks like

$$\begin{array}{c} 1 \\ | \\ -1-0- \end{array} \quad \begin{array}{c} 1 \\ | \\ -0-1- \end{array} \quad \begin{array}{c} 0 \\ | \\ -1-1- \end{array}$$

there. In all other cases we call $\underline{\xi}$ type \emptyset at i . We define the *type vector* of $\underline{\xi}$, denoted $t(\underline{\xi})$, to be the ordered tuple of the $r - 2$ types of $\underline{\xi}$. The type of a triple is a quadratic invariant — if $\underline{\xi} \sim \underline{\xi}'$ is an essentially quadratic relation between triples, then $t(\underline{\xi}) = t(\underline{\xi}')$. The toric generalized Segre cubic changes the type.

4.6. Reformulation of Theorem 4.1. Rather than proving the statement about relations in Theorem 4.1 directly, we will prove the following:

Proposition 4.3. *Let Ξ be the r th caterpillar tree ($r \geq 3$). Then the ideal of relations I_Ξ of $\text{gr}_\Xi R_L = \mathbb{Z}[\mathcal{R}_\Xi]$ is generated by relations of degree two and three. Furthermore, “type” is the only quadratic invariant on \mathcal{V}_Ξ^3 , that is, if $\underline{\xi} \sim \underline{\xi}'$ is a relation with $\underline{\xi}, \underline{\xi}' \in \mathcal{V}_\Xi^3$ then $\underline{\xi} \sim \underline{\xi}'$ is essentially quadratic if and only if $t(\underline{\xi}) = t(\underline{\xi}')$.*

Theorem 4.1 follows easily from this, since generalized toric Segre cubics allow one to switch the type of a triple between A and B at any base vertex.

We now give an overview of the proof of Proposition 4.3. First of all we split relations into two types defined in §4.8, “breakable” and “unbreakable.” Breakable means that some degree one piece of the relation vanishes at some base edge. The first idea in the proof is that, given a relation, one can chop the tree into pieces such that the relation is unbreakable on each piece. One can then use a gluing argument to reduce to the unbreakable case. We then prove that an unbreakable relation is essentially quadratic. We do this by introducing a normal form for monomials and showing that any monomial can be changed into normal form by a series of quadratic relations. We use the notion of a “balanced” monomial, a tuple of matchings which assume roughly the same value on each edge of the tree. A key result (Proposition 4.4) is that any tuple can be balanced by quadratic relations.

4.7. Balancing. We say that a tuple of integers (x_i) is *balanced* if $|x_i - x_j|$ is always 0 or 1. We say that a tuple $\underline{\xi} \in \mathcal{V}^n$ is *balanced* if for each base edge e of Ξ the tuple of integers $(\xi_i(e))$ is balanced.

Proposition 4.4. *Given any tuple $\underline{\xi}$ in \mathcal{V}^n there exists an essentially quadratic relation $\underline{\xi} \sim \underline{\xi}'$ with $\underline{\xi}'$ balanced.*

Proof. It suffices to prove the proposition for $n = 2$, as one can repeatedly balance pairs of integers to balance a set of integers. Thus suppose that $n = 2$. First we show that the proposition holds for the third caterpillar tree. In the general case of the r th caterpillar, we break Ξ into its trinodes, balance the weightings on each of these separately and then “glue.”

Third caterpillar: We indicate a reduced matching ξ on the third caterpillar tree by a triple (a, b, c) where a is the weight of the first (left) stalk, b is the weight of the second (vertical) stalk and c is the weight of the third (right) stalk. The triple (a, b, c) satisfies the triangle inequalities and has $b \leq 1$. If $b = 0$ then $a = c$. If $b = 1$ then one of a or c is nonzero, and $|a - c| \leq 1$. Suppose $\xi_1 = (a_1, b_1, c_1)$ and $\xi_2 = (a_2, b_2, c_2)$ are reduced matchings and $|a_1 - a_2| \geq 2$ or $|c_1 - c_2| \geq 2$. Without loss of generality we take $a_1 + 2 \leq a_2$. Since $c_1 \leq a_1 + 1$ and $c_2 \geq a_2 - 1$ we have $c_1 \leq c_2$. We know that $c_2 \geq 1$ since $a_2 \geq 2$.

If $c_1 < c_2$ then define $\xi'_1 = (a_1 + 1, b_1, c_1 + 1)$ and $\xi'_2 = (a_2 - 1, b_2, c_2 - 1)$. Then $\xi'_1 + \xi'_2 = \xi_1 + \xi_2$ so $(\xi'_1, \xi'_2) \sim (\xi_1, \xi_2)$. The resulting pair (ξ'_1, ξ'_2) is now closer to being balanced since $|(a_1 + 1) - (a_2 - 1)| = |a_1 - a_2| - 2$ and $|(c_1 + 1) - (c_2 - 1)| \leq |c_1 - c_2|$.

If $c_1 = c_2$ then $a_2 = a_1 + 2$ and $c_1 = c_2 = a_1 + 1$. Define $\xi'_1 = (a_1 + 1, b_1, c_1)$ and $\xi'_2 = (a_2 - 1, b_2, c_2)$. Then ξ'_1 and ξ'_2 are reduced admissible matchings and $(\xi'_1, \xi'_2) \sim (\xi_1, \xi_2)$. Again, we get strictly closer to a balanced pair with such an assignment.

A finite number of steps as above will related the original pair (ξ_1, ξ_2) to a balanced pair.

Breaking and gluing: Suppose now that we have a pair (ξ_1, ξ_2) of reduced admissible matchings on the r th caterpillar. Break the caterpillar up into an ordered tuple of $r - 2$ caterpillar trees with three vertices by cutting each base edge in two. These are arranged from left to right and indexed as $2, \dots, r - 1$. The matchings ξ_i on the original caterpillar define matchings on each copy of the third caterpillar. Let e_j, f_j and g_j denote the first, second and third stalk on the j th trinode and define $\xi_{i,j} = (\xi_i(e_j), \xi_i(f_j), \xi_i(g_j))$ for $i = 1, 2$ and $2 \leq j \leq r - 1$. Now apply our previous result on the third caterpillar — each pair $(\xi_{1,j}, \xi_{2,j})$ is equivalent to a balanced pair $(\xi'_{1,j}, \xi'_{2,j})$.

Define $t_2 = 1$. Now, because the pairs are balanced, for $3 \leq j \leq r - 1$ we have either $c'_{1,j} = a'_{1,j+1}$ and $c'_{2,j} = a'_{2,j+1}$, or $c'_{1,j} = a'_{2,j+1}$ and $c'_{2,j} = a'_{1,j+1}$. In the former case set $t_j = t_{j-1}$ while in the latter case set $t_j = 3 - t_{j-1}$. We will use the tuple (t_2, \dots, t_{r-1}) of 1's and 2's to glue these balanced pairs on individual trinodes to obtain a balanced pair on the r th caterpillar.

We now define a pair of admissible weightings (ξ'_1, ξ'_2) on the r th caterpillar by a gluing procedure. Let the edges of the r th caterpillar be labeled as $s_1, s_2, b_2, s_3, b_3, \dots, b_{r-2}, s_{r-1}, s_r$, where s_j is the j th stalk and b_k is the base edge between the k th and $(k + 1)$ st base vertices. To begin, we define ξ'_1 (resp. ξ'_2) on

s_1, s_2, b_2 to agree with $\xi'_{1,2}$ (resp. $\xi'_{2,2}$). For $3 \leq j \leq r-3$ set

$$(\xi'_1(b_{j-1}), \xi'_1(s_j), \xi'_1(b_j)) = \xi'_{t_j}, \quad (\xi'_2(b_{j-1}), \xi'_2(s_j), \xi'_2(b_j)) = \xi'_{2-t_j}.$$

Finally define

$$(\xi'_1(b_{r-2}), \xi'_1(s_{r-1}), \xi'_1(s_r)) = \xi'_{t_{r-1}}, \quad (\xi'_2(b_{r-2}), \xi'_2(s_{r-1}), \xi'_2(s_r)) = \xi'_{2-t_{r-1}}.$$

The above assignments are well-defined, (ξ'_1, ξ'_2) is balanced and $(\xi'_1, \xi'_2) \sim (\xi_1, \xi_2)$. \square

4.8. Reduction of Proposition 4.3 to the unbreakable case. We say that a matching ξ on Ξ is *breakable* if there exists a base edge e with $\xi(e) = 0$, and *unbreakable* otherwise. We say that a tuple of matchings $\underline{\xi} \in \mathcal{V}_{\Xi}^n$ is *unbreakable* if each ξ_i is. In §4.9 we will prove the following proposition:

Proposition 4.5. *Let $\underline{\xi} \sim \underline{\xi}'$ be a relation with $\underline{\xi}, \underline{\xi}' \in \mathcal{V}_{\Xi}^n$ and $\underline{\xi}$ unbreakable. Then $\underline{\xi} \sim \underline{\xi}'$ is essentially quadratic.*

In this section we prove the following:

Proposition 4.6. *Proposition 4.5 implies Proposition 4.3.*

Proof. As remarked in §4.4, Proposition 4.3 is true for the third caterpillar tree. This will be the base case of an inductive argument.

Let $x \sim y$ be a relation of length n . Using quadratic relations, we may assume that both x and y are balanced. If each x_i is unbreakable then each y_i is as well (since x and y are balanced) and we are done. Assume then that there is a base edge e for which x_i is breakable at e for some i .

Cut the edge e in half to produce two new trees Ξ' and Ξ'' . We regard e as an edge of both of these trees. Both of these trees can be regarded as smaller caterpillar trees. Also, giving a weighting on Ξ is equivalent to giving weightings on Ξ' and Ξ'' which agree at e .

Now, let x' and x'' be the restrictions of x to Ξ' and Ξ'' (and similarly for y). The key point is that because e is breakable and x and y are balanced these restricted weightings are *matchings*, that is, they assign e either 0 or 1. In other words, we have $x', y' \in \mathcal{V}_{\Xi'}^n$, and $x'', y'' \in \mathcal{V}_{\Xi''}^n$.

Proceeding by induction, we can assume that all relations are essentially cubic on Ξ' and Ξ'' , or essentially quadratic if the types agree. We can then pick a sequence of cubic (resp. quadratic) relations between x' and y' and between x'' and y'' and concatenate them to form a sequence of cubic (resp. quadratic) relations between x and y . By “concatenate,” we mean that one should first order the tuples so that those taking value zero at the edge e should be glued together (the order does not matter), and those taking value one at e should be glued together (the order does not matter). In the final step — that is, after x' has been replaced with y' (up to permutation) and x'' has been replaced with y'' (up to permutation) — one can finally permute the y'' matchings taking value zero at e , and permute those taking value one at e , and lastly permute the concatenated matchings, so that the result is equal to y . (Recall that permutations are essentially quadratic since they are generated by 2-cycles.) \square

4.9. Proof of Proposition 4.5. Let ξ be an unbreakable matching. There are four possibilities for ξ at one of the internal trinodes. We label them as follows:

$$\begin{array}{c} | \\ \hline 0 \\ \hline n \quad n \\ A_n \end{array} \quad \begin{array}{c} | \\ \hline 1 \\ \hline n \quad n \\ B_n \end{array} \quad \begin{array}{c} | \\ \hline 1 \\ \hline n \quad n+1 \\ C_n \end{array} \quad \begin{array}{c} | \\ \hline 1 \\ \hline n+1 \quad n \\ D_n \end{array}$$

In all cases n is non-zero. We order them as follows: $A_n \leq B_n \leq C_n \leq D_n \leq A_{n+1}$. There are also four possibilities for ξ at the trinodes on the left end:

$$\begin{array}{c} 1 \\ | \\ \hline 1 \\ \hline 0 \end{array} \quad \begin{array}{c} 0 \\ | \\ \hline 1 \\ \hline 1 \end{array} \quad \begin{array}{c} 1 \\ | \\ \hline 1 \\ \hline 1 \end{array} \quad \begin{array}{c} 1 \\ | \\ \hline 2 \\ \hline 1 \end{array}$$

$E \qquad F \qquad G \qquad H$

The right end is similar (flip each over). We order these: $E \leq F \leq G \leq H$.

We can express an unbreakable matching as a string using the above alphabet and some obvious rules. Let ξ and ξ' be two unbreakable matchings. We define a partial ordering on unbreakable matchings, by $\xi \leq \xi'$ if this is the case when restricted to each trinode. Equivalently, $\xi \leq \xi'$ if the following holds:

- Both ends of ξ are at most the respective ends of ξ' .
- For each base edge e we have $\xi(e) \leq \xi'(e)$.
- If e and e' are consecutive base edges for which $\xi(e) = \xi'(e)$ and $\xi(e') = \xi'(e')$ then $\xi(s) \leq \xi'(s)$ where s is the stalk in between e and e' .

We say that an unbreakable tuple of matchings $\underline{\xi} \in \mathcal{V}_{\Xi}^n$ is in *normal form* if $\underline{\xi}$ is balanced and $\underline{\xi}_1 \leq \underline{\xi}_2 \leq \dots \leq \underline{\xi}_n$. Proposition 4.5 now follows from the following proposition:

Proposition 4.7. *Every unbreakable element of \mathcal{V}_{Ξ}^n is related to a unique normal form, and this relation is essentially quadratic.*

The following is the key lemma:

Lemma 4.8. *Let $(\xi, \xi') \in \mathcal{V}_{\Xi}^2$ be unbreakable and balanced. Then there exists $(\eta, \eta') \in \mathcal{V}_{\Xi}^2$ balanced and unbreakable such that $\xi + \xi' = \eta + \eta'$ and for each internal edge e we have*

$$\eta(e) = \min(\xi(e), \xi'(e)), \quad \eta'(e) = \max(\xi(e), \xi'(e)).$$

Proof. The idea is the same as in the proof of Proposition 4.4: break apart at a base edge (or at an end), permute, and glue back together to achieve the desired order. \square

We now prove Proposition 4.7.

Proof of Proposition 4.7. Let $\underline{\xi} \in \mathcal{V}_{\Xi}^n$ be a given unbreakable tuple. We may assume that $\underline{\xi}$ is balanced. By repeatedly using Lemma 4.8 we find that $\underline{\xi}$ is quadratically related to an unbreakable element $\underline{\xi}' \in \mathcal{V}_{\Xi}^n$ which has the property that $\xi'_i(e) \leq \xi'_j(e)$ for $i \leq j$ and all internal edges e .

For uniqueness, it suffices to show uniqueness for third caterpillars and for the ends. Consider the third caterpillar at a base vertex. Suppose the sum of the n weightings (a_i, b_i, c_i) , $1 \leq i \leq n$, on the third caterpillar is equal to (a, b, c) (left, stalk, right). Suppose these weightings are increasing in the order we have defined. This means the a_i 's are increasing, starting as the floor of a/n and ending as the ceiling of a/n . This determines the value of each a_i . Similarly the value of each c_i is determined (beginning with floor of c/n , ending with ceiling of c/n). Wherever $a_i \neq c_i$ we must have $b_i = 1$, so these b_i 's are determined. However if $a_i = c_i$, then b_i could be either 0 or 1. The set $\{i \mid a_i = c_i\}$ consists of at most two intervals I, J . These intervals are determined by the a_i, c_i (which are determined by the normality condition). Within the interval I , the b_i 's must be increasing. Similarly in the interval J , the b_i 's must increase. Thus the b_i 's are all determined by the normality condition. The argument for the ends is similar. \square

5. THE IDEAL IS GENERATED BY QUADRATICS AND GENERALIZED SEGRE CUBICS

In §5, we lift the generalized Segre toric cubics to the ring R_L and prove the following:

Theorem 5.1. *For any even set L the ideal I_L is generated over \mathbb{Z} by quadratics and the small generalized Segre cubic relations.*

We will introduce the generalized Segre cubics, and the small generalized Segre cubics, shortly. Theorem 5.1 will follow easily from our toric results once we make these definitions. In one of our ad hoc arguments in the 10 point case we will need a slightly refined version of Theorem 5.1 given in Remark 5.3. The proof of this theorem is the only place we use the toric results.

5.1. Brief additional comments on colored graphs. We will use the language of colored graphs, introduced in §1. We consider both directed and undirected colored graphs. If Γ is a directed multi-matching on L whose edges have been colored with colors from the set C then X_{Γ} is defined as an element of $V_L^{\otimes C}$ as in §1. If L is oriented and Γ is a regular undirected multi-matching one can also make sense of Y_{Γ} as an element of $V_L^{\otimes C}$. Clearly the X_{Γ} (or Y_{Γ}) span $V_L^{\otimes C}$. Furthermore, the X_{Γ} satisfy the sign and ‘‘colored’’ Plücker relations, and these generate all the linear relations among them. (The colored Plücker relation is just the usual Plücker relation on a pair of edges, with the restriction that these two edges be of the same color.) We have thus given a description of $V_L^{\otimes C}$ in terms of colored graphs. There is a similar description for $\text{Sym}^k(V_L)$. The only difference is that in $\text{Sym}^k(V_L)$ the particular color of an edge is not relevant. What matters is whether two edges have the same color — two colored graphs represent the same element of $\text{Sym}^k(V_L)$ if one is obtained from the other by permuting the colors.

5.2. Generalized Segre cubic data. Let L be an even set. By a *generalized Segre datum* we mean a pair $\Sigma = (\Gamma, \mathcal{U})$ where Γ is an undirected graph on L whose edges have been colored one of red, green or blue and $\mathcal{U} = \{U_R, U_G, U_B\}$ is a partition of L into three even subsets (called *parts*), such that

- Every vertex of Γ has valence one for each of the three colors.
- Between any two parts there are either two edges or no edges. If there are two edges then these edges have the color of the “opposite” part. For instance, any edge from U_G to U_B must be red.

See Figure 5 for a schematic presentation. We call the edges between the parts *special*. We call Σ *small* if U_R, U_G or U_B has cardinality two.




FIGURE 5. Schematic presentation of a generalized Segre datum.

Let Σ be a generalized Segre datum. Suppose e and e' are edges of Γ of the same color and have a vertex in a common part. Given such a pair, let $Y_\Gamma + Y_{\Gamma'} + Y_{\Gamma''} = 0$ be the colored Plücker relation on e and e' . Then it is easily verified that $\Sigma' = (\Gamma', \mathcal{U})$ and $\Sigma'' = (\Gamma'', \mathcal{U})$ are both generalized Segre data. Define the *space of generalized Segre data* to be the \mathbb{Z} -module spanned by generalized Segre data modulo relations of the form $\Sigma + \Sigma' + \Sigma'' = 0$.

5.3. Generalized Segre cubic relations. Let $\Sigma = (\Gamma, \mathcal{U})$ be a generalized Segre datum. Let $\tilde{\Gamma}$ be a directed colored graph with underlying undirected graph Γ . We let $\epsilon(\tilde{\Gamma})$ be the product of the $\epsilon(\tilde{\Gamma}_i)$, where $\tilde{\Gamma}_i$ is the directed matching of color i in $\tilde{\Gamma}$. (Here ϵ is a chosen orientation on L , see §2.2.) Form a new directed colored graph $\tilde{\Gamma}'$ as follows. The graph $\tilde{\Gamma}'$ will be a recoloring of $\tilde{\Gamma}$, so we just specify a new color for each edge. We use the colors purple and black. The red (resp. blue, green) edges of $\tilde{\Gamma}$ in U_R (resp. U_B, U_G) are black in $\tilde{\Gamma}'$ and all other edges of $\tilde{\Gamma}$ are purple in $\tilde{\Gamma}'$. It is clear that every vertex has black valence one and thus purple valence two.

Define $Y_\Sigma \in \text{Sym}^3(V_L)$ and $Y'_\Sigma \in R_L^{(1)} \otimes R_L^{(2)}$ by

$$Y_\Sigma = Y_\Gamma = \epsilon(\tilde{\Gamma})X_{\tilde{\Gamma}}, \quad Y'_\Sigma = \epsilon(\tilde{\Gamma}) \left(X_{\tilde{\Gamma}'_{\text{black}}} \otimes X_{\tilde{\Gamma}'_{\text{purple}}} \right).$$

These only depend on Σ and not the choice of $\tilde{\Gamma}$. There is a well-defined map $R_L^{(1)} \otimes R_L^{(2)} \rightarrow \text{Sym}^3(V_L)/Q_L^{(3)}$ given by writing the element of $R_L^{(2)}$ in terms of degree one elements and then formally multiplying to get an element of $\text{Sym}^3(V_L)$. (This is only defined modulo quadratic relations Q_L , because of the choice of how to write $R_L^{(2)}$ in terms of degree one elements.) We may therefore regard both Y_Σ and Y'_Σ as elements of $\text{Sym}^3(R_L^{(1)})/Q_L$. Define

$$\text{Rel}(\Sigma) = Y_\Sigma - Y'_\Sigma,$$

regarded as an element of $\text{Sym}^3(V_L)/Q_L^{(3)}$. As the two terms in $\text{Rel}(\Sigma)$ are recolorings of the same graph, $\text{Rel}(\Sigma)$ is a relation, that is, it maps to zero in R_L . We call such relations *generalized Segre cubic relations*. We use the same name for lift of some $\text{Rel}(\Sigma)$ to $I_L^{(3)} \subset \text{Sym}^3(V_L)$. A generalized Segre cubic relation is shown in Figure 6. Note that for $\sigma \in \mathfrak{S}_L$ we have $\sigma \text{Rel}(\Sigma) = \text{sgn}(\sigma) \text{Rel}(\sigma\Sigma)$.

One easily verifies that Rel gives a homomorphism

$$\text{Rel} : \{\text{the space of generalized Segre cubic data}\} \rightarrow I_L^{(3)}/Q_L^{(3)}.$$

This allows us to interpret relations as graphs with extra structure, which has two advantages. First, it gives a source of linear relations between (also thought of as operations we can perform on) generalized Segre relations: Plücker relations on the generalized Segre data which respect the extra structure. Second, one can read off certain properties of the generalized Segre relation from the original graph. For instance, if Σ is a generalized Segre cubic datum for which Γ is disconnected, then the associated generalized Segre relation arises from a relation on fewer points (see Proposition 7.3 for a precise statement).




FIGURE 6. A generalized Segre relation on six points. Here $U_G = \{1, 2\}$, $U_R = \{3, 4\}$ and $U_B = \{5, 6\}$. This is equal to the image of the usual Segre cubic relation in I_L/Q_L .

5.4. **Degenerate Segre cubic relations.** We call a generalized Segre datum Σ *degenerate* if:

- (1) one of the pairs of special edges is missing; or
- (2) in one of the parts the two pairs of special edges do not connect.

By (2) we mean that one part, say U_G , can be partitioned into two pieces A and B such that no edges go between A and B , the blue special edges go into A and the red special edges go into B . A generalized Segre relation is called degenerate if it comes from a degenerate generalized Segre datum.

Proposition 5.2. *A degenerate generalized Segre relation lies in Q_L .*

Proof. First consider a degenerate datum satisfying (1). Say that the red special edges are missing. We show how to move from Y_Σ to Y'_Σ using quadratic relations. First switch the red edges and the green edges of Γ which occur in U_G (this relation lies in Q_L). Then switch the red edges and blue edges of Γ in U_B . The resulting graph now looks like Γ' if we make the red edges black and the other edges purple.

Now say that Σ is a degenerate datum satisfying (2). Say that in U_G the red and blue special edges do not connect and let $U_G = A \amalg B$ be a decomposition as described above. Let Γ' be the graph obtained from Γ by switching the red and green colors in A , let $U'_G = B$, let $U'_R = U_R \cup A$ and let $U'_B = U_B$. Let Σ' be the generalized Segre datum $(\Gamma', \{U'_R, U'_G, U'_B\})$. Then Σ' is degenerate of type (1) and so $\text{Rel}(\Sigma') \in Q_L$. Consider the difference

$$\text{Rel}(\Sigma) - \text{Rel}(\Sigma') = (Y_\Sigma - Y_{\Sigma'}) + (Y'_\Sigma - Y'_{\Sigma'}).$$

Now $Y'_\Sigma = Y'_{\Sigma'}$. Also, $Y_\Sigma - Y_{\Sigma'} \in Q_L$ (the blue subgraphs of Y_Σ and $Y_{\Sigma'}$ are identical, and the remainder of the graph gives a quadratic relation). Thus $\text{Rel}(\Sigma) \in Q_L$. \square

5.5. **Proof of Theorem 5.1.** It follows from Theorem 4.1 that I_L is generated by quadratics and lifts of the generalized toric Segre cubics. Thus to prove Theorem 5.1, it suffices to show that the generalized Segre cubics are lifts of the generalized toric Segre cubics, which is what we now do.

Let L be an even set. Pick a Y-tree Ξ with leaf set L and let Ξ^- be the associated reduced (caterpillar) tree. Consider a generalized toric Segre cubic relation on Ξ^- :

$$(5.1) \quad \begin{array}{ccc} \begin{array}{c} 1 \\ | \\ \boxed{X} - 1 - 0 - \boxed{Z'} \\ | \\ p \end{array} & \begin{array}{c} 1 \\ | \\ \boxed{Z} - 0 - 1 - \boxed{X'} \\ | \\ p \end{array} & \begin{array}{c} 0 \\ | \\ \boxed{Y} - 1 - 1 - \boxed{Y'} \\ | \\ q \end{array} \\ = & \begin{array}{c} 1 \\ | \\ \boxed{X} - 1 - 1 - \boxed{X'} \\ | \\ p \end{array} & \begin{array}{c} 1 \\ | \\ \boxed{Y} - 1 - 1 - \boxed{Y'} \\ | \\ q \end{array} & \begin{array}{c} 0 \\ | \\ \boxed{Z} - 0 - 0 - \boxed{Z'} \\ | \\ q \end{array} \end{array}$$

Implicit in this diagram is a decomposition of Ξ^- into three pieces:



The labels indicate how these three pieces are glued together (after gluing, p and q disappear). Corresponding to this decomposition is a decomposition of L into three pieces, the left piece P , the right piece P' and the center piece, which has two elements x and y (the two vertices of the Y-tree Ξ which connect to the stalk of the trinode in the center of the above diagram). We regard X , Y and Z as reduced degree one weightings on the left caterpillar tree, where X and Y take value 1 on p while Z takes value 0 there. Similarly, we regard X' , Y' and Z' as reduced weightings on the right caterpillar tree, where X' and Y' take value 1 on q while Z' takes value 0 there.

Now, we can regard X , Y and Z as non-reduced degree one weightings on the Y-tree with vertex set $P \cup \{p^+, p^-\}$. Similarly, we can regard X' , Y' and Z' as non-reduced degree one weightings on the Y-tree with vertex set $P' \cup \{q^+, q^-\}$. We lift these six weightings to matchings, which we denote by \tilde{X} , etc. In the graph \tilde{X} there are two elements of P , say a and b , which connect to p^+ and p^- . We let \tilde{X}_0 be the restriction of \tilde{X} to $P \setminus \{a, b\}$. We let c and d be the two vertices in \tilde{Y} connecting to p^+ and p^- and let \tilde{Y}_0 be the restriction of \tilde{Y} to $P \setminus \{c, d\}$. In \tilde{Z} the vertices p^+ and p^- are connected to each other. We similarly define primed versions.

We now define a small generalized Segre cubic datum Σ . The partition is given by $U_B = P'$, $U_R = \{x, y\}$ and $U_G = P$. The graph Γ is given as follows:

- The blue graph is the union of \tilde{X}_0 , \tilde{Z}' and the edges \overline{ax} and \overline{by} .
- The green graph is the union of \tilde{Z} , \tilde{X}'_0 and the edges $\overline{a'x}$ and $\overline{b'y}$.
- The red graph is the union of \tilde{Y}_0 , \tilde{Y}'_0 and the edges \overline{xy} , $\overline{cc'}$ and $\overline{dd'}$.

These three matchings are lifts for the three reduced weightings appearing on the left side of (5.1). It is clear that this is a valid small generalized Segre cubic datum.

We must now show that the generalized Segre cubic relation $\text{Rel}(\Sigma)$ associated to the generalized Segre cubic datum defined above lifts the relation (5.1). Now, the relation associated to the datum is

$$\begin{aligned} & \left(\tilde{X}_0 \cdot \overline{ax} \cdot \overline{by} \cdot \tilde{Z}' \right) \left(\tilde{Z} \cdot \overline{a'x} \cdot \overline{b'y} \cdot \tilde{X}'_0 \right) \left(\tilde{Y} \cdot \overline{cc'} \cdot \overline{dd'} \cdot \overline{xy} \cdot \tilde{Y}' \right) \\ &= \left(\tilde{X}_0 \cdot \tilde{Y}_0 \cdot \overline{ax} \cdot \overline{a'x} \cdot \overline{by} \cdot \overline{b'y} \cdot \overline{cc'} \cdot \overline{dd'} \cdot \tilde{X}'_0 \cdot \tilde{Y}'_0 \right) \left(\tilde{Z} \cdot \overline{xy} \cdot \tilde{Z}' \right). \end{aligned}$$

Here the parentheses should each be interpreted as single graphs — for example, the first parenthetical is the concatenation of \tilde{X}_0 , \overline{ax} , \overline{by} and \tilde{Z}' . The quadratic term Δ on the right side would be the purple subgraph in the graphical notation. Now, the above relation belongs to $I_L^{(3)}/Q_L^{(3)}$. To get an element of $I_L^{(3)}$, we must rewrite Δ in terms of degree one elements. Because the toric ring $\text{gr}_{\Xi} R_L$ is generated in degree one, we can write

$$\Delta = \sum_{i=1}^n \Phi_i \Phi'_i$$

where each Φ_i and Φ'_i has degree one, $\Phi_1 \Phi'_1$ has toric weight equal to that of Δ and $\Phi_i \Phi'_i$ has toric weight strictly less than that of Δ for $i > 1$. We may take $\Phi_1 = \tilde{X}_0 \cdot \overline{ax} \cdot \overline{a'y} \cdot \overline{bb'} \cdot \tilde{X}'_0$ and $\Phi'_1 = \tilde{Y}_0 \cdot \overline{cx} \cdot \overline{c'y} \cdot \overline{dd'} \cdot \tilde{Y}'_0$, since the product of these two graphs is equal to Δ in the toric ring.

It now follows that $\text{Rel}(\Sigma)$ is represented by the relation

$$\left(\tilde{X}_0 \cdot \overline{ax} \cdot \overline{by} \cdot \tilde{Z}' \right) \left(\tilde{Z} \cdot \overline{a'x} \cdot \overline{b'y} \cdot \tilde{X}'_0 \right) \left(\tilde{Y} \cdot \overline{cc'} \cdot \overline{dd'} \cdot \overline{xy} \cdot \tilde{Y}' \right) = \left(\sum_{i=1}^n \Phi_i \Phi'_i \right) \left(\tilde{Z} \cdot \overline{xy} \cdot \tilde{Z}' \right).$$

The leading term (in terms of the grading) of this relation is

$$\left(\tilde{X}_0 \cdot \overline{ax} \cdot \overline{by} \cdot \tilde{Z}' \right) \left(\tilde{Z} \cdot \overline{a'x} \cdot \overline{b'y} \cdot \tilde{X}'_0 \right) \left(\tilde{Y} \cdot \overline{cc'} \cdot \overline{dd'} \cdot \overline{xy} \cdot \tilde{Y}' \right) = (\Phi_1 \Phi'_1) \left(\tilde{Z} \cdot \overline{xy} \cdot \tilde{Z}' \right).$$

The two sides are the same as the two sides of (5.1). This shows that the leading term of our generalized Segre relation $\text{Rel}(\Sigma)$ is equal to the generalized toric Segre relation we started with, which proves Theorem 5.1.

Remark 5.3. It follows from the proof that, if we totally order L , then I_L is generated by quadratic relations and those small generalized Segre cubic relations coming from data for which $U_G < U_R < U_B$ and $|U_R| = 2$. We will use this stronger form of Theorem 5.1 in the proof of Proposition 7.9.

6. THE STRUCTURE OF V_L AND ITS TENSOR POWERS

In §6, we study the partition filtration on $\otimes^n V_L$, $\text{Sym}^n(V_L)$ and $\wedge^n V_L$ for $n \leq 3$ and the \mathfrak{S}_L -action on the associated graded pieces. This provides us with essential structural properties of the Plücker relation; for example: the elements $(Y_\Gamma)^2$ with Γ a matching span $\text{Sym}^2(V_L)$ (see the discussion following Proposition 6.1) and useful generalizations.

6.1. The partition filtration. Let L be an even set. By a *partition* of L we mean a collection of non-empty disjoint subsets of L whose union is L . We say that a partition is *into even parts* if each of the subsets is even. We similarly speak of partitions of $|L|$ into even parts. Given a partition \mathcal{U} of L we denote by $|\mathcal{U}|$ the corresponding partition of $|L|$. We partially order the set of partitions of L and $|L|$ by refinement. For example $2 + 2 + 2 + 2$ is smaller than $4 + 2 + 2$, but $4 + 4$ and $6 + 2$ are not comparable.

Given a regular colored graph Γ on L we obtain a partition \mathcal{U}_Γ of L by taking the vertex sets of the connected components of Γ . The partition \mathcal{U}_Γ necessarily has even parts. We use this to define a filtration, indexed by the partitions of $|L|$ into even parts, on any \mathbb{Z} -module U which is spanned by graphs. We denote this filtration, which we call the *partition filtration*, by $F_p U$. (Warning: this filtration is indexed by a partially ordered set, not a totally ordered set.) For instance, if p is a partition of $|L|$ then $F_p \text{Sym}^k(V_L)$ is the span of the X_Γ (or Y_Γ) for which $|\mathcal{U}_\Gamma| \leq p$. We denote the associated graded subquotient by $\text{gr}_p U$. To be precise, $\text{gr}_p U$ is the quotient of $F_p U$ by the span of the $F_{p'} U$ with $p' < p$. The partition filtration is preserved by the action of \mathfrak{S}_L , so this group naturally acts on $\text{gr}_p U$.

6.2. Degree one spaces. The $n = 1$ case is easy: the only non-zero piece of the partition filtration on V_L occurs for $p = 2 + \dots + 2$ and then $\text{gr}_p = V_L$ is the irreducible representation of \mathfrak{S}_L corresponding to the partition $n/2 + n/2$ (§1.5).

6.3. Degree two spaces. We now study the partition filtration on $V_L^{\otimes 2}$, $\text{Sym}^2(V_L)$ and $\bigwedge^2 V_L$.

Proposition 6.1. *We have, over $\mathbb{Z}[\frac{1}{2}]$:*

$$\text{gr}_p(V_L^{\otimes 2}) = \begin{cases} \text{Sym}^2(V_L) & \text{if } p = 2 + \dots + 2, \\ \bigwedge^2 V_L & \text{if } p = 4 + 2 + \dots + 2, \\ 0 & \text{otherwise.} \end{cases}$$

The space $V_L^{\otimes 2}$ is spanned by regular 2-colored graphs. Such graphs are disjoint unions of cycles of even size. Proposition 6.1 says that $V_L^{\otimes 2}$ is spanned by graphs which are unions of 2-cycles and at most one 4-cycle. Furthermore, $\text{Sym}^2(V_L)$ is spanned by graphs which are unions of 2-cycles: the elements $(Y_\Gamma)^2$ with Γ a matching span $\text{Sym}^2(V_L)$.

We begin our proof of Proposition 6.1 with the following result:

Lemma 6.2. *The space $V_L^{\otimes 2}$ is spanned over $\mathbb{Z}[\frac{1}{2}]$ by graphs which are unions of 2-cycles and 4-cycles.*

Proof. It suffices, by induction, to show that every regular 2-colored graph on at least six vertices can be written as a sum of graphs which are not connected. In other words, letting U_L be the subspace of $V_L^{\otimes 2}$ spanned by disconnected graphs, it suffices to show that $V_L^{\otimes 2}/U_L$ is zero. Thus let Γ be a graph on L , which we can assume to be connected (otherwise it already belongs to U_L). We must show $Y_\Gamma = 0$ in $V_L^{\otimes 2}/U_L$.

Pick four consecutive vertices a, b, c and d of Γ . The Plücker relation on the edges \overline{ab} and \overline{cd} is

$$0 = \begin{array}{c} b \quad c \\ \diagdown \quad \diagup \\ a \quad d \end{array} + \begin{array}{c} c \quad b \\ \diagdown \quad \diagup \\ a \quad d \end{array} + \begin{array}{c} \overline{bc} \\ \overline{ad} \end{array}$$

The rightmost term belongs to U_L . We therefore have $Y_\Gamma = -Y_{\Gamma'}$ in $V_L^{\otimes 2}/U_L$ where Γ' is obtained by transposing two consecutive vertices in Γ . In words, we may transpose consecutive vertices at the cost of a sign, modulo U_L .

Now consider six consecutive vertices a, b, c, d, e and f in Γ . Applying the Plücker relation on the edges \overline{ab} and \overline{ef} we find

$$0 = \begin{array}{c} c \quad d \\ \diagdown \quad \diagup \\ b \quad e \\ \diagup \quad \diagdown \\ a \quad f \end{array} + \begin{array}{c} d \quad c \\ \diagdown \quad \diagup \\ e \quad b \\ \diagup \quad \diagdown \\ a \quad f \end{array} + \begin{array}{c} \overline{cd} \\ \overline{ef} \end{array}$$

Again, the rightmost term belongs to U_L and thus can be discarded. This shows that we may pick four consecutive vertices and reverse their order at the cost of picking up a sign. However, we may now move the affected vertices back to the original position using six transpositions, *not* introducing a new sign. We thus find that $Y_\Gamma = -Y_\Gamma$ in $V_L^{\otimes 2}/U_L$, which shows that $2Y_\Gamma$ belongs to U_L . \square

In the proof of Lemma 6.2 we only ever use six consecutive vertices. By keeping track of the graphs we discarded during the course of the proof we obtain an identity, shown in Figure 7, that we will use on a few later occasions.




FIGURE 7. A graphical identity on 6 points. More accurately, this should be regarded as a family of identities between various tensors $Y_\Gamma \otimes Y_\Delta$ in $V_L^{\otimes 2}$ for any L ; we have only drawn the edges on each side which are different.

Lemma 6.2 states that $V_L^{\otimes 2}$ is spanned by graphs with only 2- and 4-cycles. To prove Proposition 6.1, we must therefore show that two 4-cycles can be converted into two 2-cycles and one 4-cycle. This is a question about graphs on eight points; in fact, it suffices to show $\text{gr}_{4,4}(V_L^{\otimes 2}) = 0$ when L has cardinality eight. We prove this after the following lemma.

Lemma 6.3. *Let L be an even set and let Γ be a regular 2-colored graph on L . Put $p = |\mathcal{U}_\Gamma|$. Assume that Γ has a 4-cycle and let Γ' be the graph obtained by switching the colors of Γ on this 4-cycle. Then $Y_\Gamma = -Y_{\Gamma'}$ holds in $\text{gr}_p(V_L^{\otimes 2})$.*

Proof. Let Γ be given and say that the vertices in the 4-cycle are labeled $\{1, 2, 3, 4\}$. By squaring the identity

$$(-1) \begin{array}{c} 2 \quad 3 \\ \diagdown \quad \diagup \\ 1 \quad 4 \end{array} = \text{---} + \begin{array}{c} | \quad | \\ | \quad | \end{array}$$

we obtain


$$\begin{array}{c} \diagdown \quad \diagup \\ \diagup \quad \diagdown \end{array} = \text{---} + \begin{array}{c} \square \\ \square \end{array} + \begin{array}{c} \square \\ \square \end{array} + \begin{array}{c} | \quad | \\ | \quad | \end{array}$$

The middle two terms on the right side of this identity are Γ and Γ' . The other terms have a more refined partition and so map to 0 in $\text{gr}_p(V_L^{\otimes 2})$. This proves the lemma. \square

Lemma 6.4. *Let L be a set of cardinality eight. Then $\text{gr}_{4,4}(V_L^{\otimes 2}) = 0$ over $\mathbb{Z}[\frac{1}{2}]$.*

Proof. In this proof we write F_p for $F_p(V_L^{\otimes 2})$, and similarly for gr_p . Let τ be the involution on $V_L^{\otimes 2}$ which transposes factors. In terms of colored graphs, τ switches the colors. We will show that τ acts by $+1$ and by -1 on $\text{gr}_{4,4}$, which will establish the lemma.

We first show that τ acts by -1 on $\text{gr}_{4,4}$. By Lemma 6.2 we have $F_{4,4} = V_L^{\otimes 2}$ and $F_{6,2} = F_{4,2,2}$. Thus $\text{gr}_{4,4}$ can be described as the space of all degree two multi-matchings on L modulo the disconnected ones. It therefore suffices to show that τ acts by -1 on a connected graph. Consider such a graph Γ , say the following one:



By the same reasoning as in the proof of Lemma 6.2 we may switch consecutive vertices at the cost of a sign, when working modulo disconnected graphs. Thus we may pick vertex 1 and move it counterclockwise around the cycle to the position of vertex 8. This uses seven transpositions and so introduces a sign. The resulting cycle is the same as the original but with the colors switched. We thus have $\tau\Gamma = -\Gamma$ modulo disconnected graphs, which shows that τ acts by -1 on $\text{gr}_{4,4}$.

We now show that τ acts by $+1$ on $\text{gr}_{4,4}$. To do this, it suffices to show that τ fixes the graph



modulo $\text{gr}_{4,2,2}$. This follows immediately from Lemma 6.3: switching the color in one square introduces a sign modulo $\text{gr}_{4,2,2}$, so switching the color in both squares introduces no sign. \square

We now complete the proof of Proposition 6.1.

Proof of Proposition 6.1. Combining Lemma 6.2 with Lemma 6.4 shows that $\text{gr}_p(V_L^{\otimes 2}) = 0$ unless $p = 2 + \dots + 2$ or $p = 4 + 2 + \dots + 2$. We thus have a filtration

$$0 \subset F_{2+\dots+2}(V_L^{\otimes 2}) \subset F_{4+2+\dots+2}(V_L^{\otimes 2}) = V_L^{\otimes 2}.$$

Let τ be the transposition of factors on $V_L^{\otimes 2}$, as in the proof of Proposition 6.4. It is clear that τ acts as the identity on $F_{2+\dots+2}$ since this space is spanned by graphs which are unions of 2-cycles. On the other hand, Lemma 6.3 shows that τ acts by -1 on $F_{4+2+\dots+2}/F_{2+\dots+2}$. From this it follows that $\text{Sym}^2(V_L) = F_{2+\dots+2} = \text{gr}_{2+\dots+2}$ and that the quotient map $V_L^{\otimes 2} \rightarrow \text{gr}_{4+2+\dots+2}$ factors to give an isomorphism $\bigwedge^2 V_L \rightarrow \text{gr}_{4+2+\dots+2}$. \square

6.4. The action of \mathfrak{S}_L on degree two spaces. Having described the spaces $\text{gr}_p(V_L^{\otimes 2})$, we now turn to their structure as \mathfrak{S}_L -modules. Our main result is the following. Note that characteristic 0 notions and arguments from the representation theory of \mathfrak{S}_L make sense over $\mathbb{Z}[1/|L|!]$ with the obvious changes.

Proposition 6.5. *In the following table, each \mathfrak{S}_L -module is multiplicity free. The set of irreducibles it contains corresponds to the given set of partitions.*

\mathfrak{S}_L -module	Set of partitions of $ L $
$\text{Sym}^2(V_L)$	at most four parts, all even
$\bigwedge^2 V_L$	exactly four parts, all odd
$V_L^{\otimes 2}$	union of previous two sets
$R_L^{(2)}$	at most three parts, all even
$I_L^{(2)}$	exactly four parts, all even

These statements hold over $\mathbb{Z}[1/|L|!]$.

For instance, $I_L^{(2)}$ is a direct sum of those irreducible representations of \mathfrak{S}_L corresponding to partitions of $|L|$ into exactly four even parts. To prove the proposition it suffices to work over the complex numbers \mathbb{C} . We use Schur-Weyl theory: for a vector space A we have a decomposition

$$A^{\otimes L} \cong \bigoplus_{\lambda} M_{\lambda} \otimes \mathbf{S}_{\lambda}(A)$$

where the sum is over all partitions λ of $|L|$, M_{λ} denotes the irreducible representation of \mathfrak{S}_L attached to a partition λ and \mathbf{S}_{λ} is the Schur functor corresponding to λ . Here $A^{\otimes L}$ denotes the tensor product of copies of A indexed by L . The above decomposition respects the action of \mathfrak{S}_L and is functorial with respect to A .

Let P be the two-dimensional vector space over \mathbb{C} with basis $\{x, y\}$. Let $\bigwedge^2 P \rightarrow \mathbb{C}$ be the isomorphism taking $x \wedge y$ to 1. We define $\text{SL}(P) = \text{Sp}(P)$ to be the group of linear transformations of P preserving this alternating form.

Since V_L is the space of degree one invariants, it is equal, by definition, to $(P^{\otimes L})^{\text{SL}(P)}$. The action of \mathfrak{S}_L on V_L in this description is the obvious one. We thus have an \mathfrak{S}_L -equivariant isomorphism

$$V_L \otimes V_L = (P^{\otimes L} \otimes P^{\otimes L})^{\text{SL}(P) \times \text{SL}(P)} = ((P \otimes P)^{\otimes L})^{\text{SL}(P) \times \text{SL}(P)}.$$

The \mathfrak{S}_L -action is the diagonal action on the first two spaces and the usual one on the last space. The above isomorphism is also equivariant with respect to the transposition of factors τ . We now apply the Schur-Weyl decomposition to obtain

$$V_L \otimes V_L \cong \bigoplus_{\lambda} M_{\lambda} \otimes \mathbf{S}_{\lambda}(P \otimes P)^{\mathrm{SL}(P) \times \mathrm{SL}(P)}.$$

The space $P \otimes P$ has a natural symmetric inner product coming from the alternating inner product on P . This inner product is preserved by the group $\mathrm{SL}(P) \times \mathrm{SL}(P)$, and it is not hard to see that the resulting map $\mathrm{SL}(P) \times \mathrm{SL}(P) \rightarrow \mathrm{SO}(P \otimes P)$ is surjective. The transposition τ on $P \otimes P$ preserves the inner product but has determinant -1 . Thus τ and $\mathrm{SO}(P \otimes P)$ generate $\mathrm{O}(P \otimes P)$. Now, the space $\mathrm{Sym}^2(V_L)$ is just the τ invariant part of $V_L \otimes V_L$, so

$$\mathrm{Sym}^2(V_L) \cong \bigoplus_{\lambda} M_{\lambda} \otimes \mathbf{S}_{\lambda}(P \otimes P)^{\mathrm{O}(P \otimes P)}.$$

Similarly, $\bigwedge^2 V_L$ is just the subspace of $V_L \otimes V_L$ on which τ acts by -1 and so

$$\bigwedge^2 V_L \cong \bigoplus_{\lambda} M_{\lambda} \otimes \mathbf{S}_{\lambda}(P \otimes P)^{\mathrm{O}(P \otimes P), -}$$

where the minus sign means to take the subspace on which $\mathrm{O}(P \otimes P)$ acts by its sign representation. The first three lines of the table in Proposition 6.5 now follow from the $n = 4$ case of following lemma. This lemma appears as statements (1) and (2) in the proof of [Ku, Lemma 2.2].

Lemma 6.6 (S. Kumar). *Let λ be a partition and let V be a vector space of dimension n with a non-degenerate symmetric inner product.*

- $\mathbf{S}_{\lambda}(V)^{\mathrm{O}(V)}$ is one-dimensional if λ has at most n parts, all even, and is zero otherwise.
- $\mathbf{S}_{\lambda}(V)^{\mathrm{O}(V), -}$ is one-dimensional if λ has exactly n parts, all odd, and is zero otherwise.

We now turn our attention to the space $R_L^{(2)}$, the degree two invariants. We have an \mathfrak{S}_L -equivariant isomorphism $R_L^{(2)} = (\mathrm{Sym}^2(P)^{\otimes L})^{\mathrm{SL}(P)}$, so

$$R_L^{(2)} \cong \bigoplus_{\lambda} M_{\lambda} \otimes \mathbf{S}_{\lambda}(\mathrm{Sym}^2(P))^{\mathrm{SL}(P)}.$$

Use the alternating inner product on P to define a symmetric inner product on $\mathrm{Sym}^2(P)$ via

$$\langle vv', ww' \rangle = \langle v, w \rangle \langle v', w' \rangle + \langle v, w' \rangle \langle v', w \rangle.$$

The group $\mathrm{SL}(P)$ preserves this inner product and it is not hard to show that the map $\mathrm{SL}(P) \rightarrow \mathrm{SO}(\mathrm{Sym}^2(P))$ is surjective. We thus have

$$R_L^{(2)} = \bigoplus_{\lambda} M_{\lambda} \otimes \mathbf{S}_{\lambda}(\mathrm{Sym}^2(P))^{\mathrm{SO}(\mathrm{Sym}^2(P))}.$$

The fourth line of the table in Proposition 6.5 now follows from the following lemma:

Lemma 6.7. *Let λ be a partition of an even number and let V be a three dimensional vector space with a non-degenerate symmetric inner product. Then $\mathbf{S}_{\lambda}(V)^{\mathrm{SO}(V)}$ is one-dimensional if λ has at most three parts, all of which are even, and is zero otherwise.*

Proof. By the second part of Lemma 6.6 we have $\mathbf{S}_{\lambda}(V)^{\mathrm{O}(V), -} = 0$, as three odd numbers cannot have an even sum. We thus have $\mathbf{S}_{\lambda}(V)^{\mathrm{SO}(V)} = \mathbf{S}_{\lambda}(V)^{\mathrm{O}(V)}$ and the result follows from the first part of Lemma 6.6. \square

The final line of the table in Proposition 6.5 follows from $R_L^{(2)} \cong \mathrm{Sym}^2(V_L)/I_L^{(2)}$.

6.5. Degree three spaces. We now turn our attention to the cubic spaces $V_L^{\otimes 3}$ and $\mathrm{Sym}^3(V_L)$. We say that a 3-regular graph on L is a *benzene cycle* if it is a cycle in which the edges alternate between being single and doubled (Figure 8). We use this term because molecules of benzene are depicted with such graphs (on six points), as in the figure. We use the term *benzene chain* for a chain of edges which alternate between being single and doubled. A benzene 2-cycle is interpreted to mean a triple edge. The main result of this section is the following:

Proposition 6.8. *The space $\mathrm{Sym}^3(V_L)$ is spanned over $\mathbb{Z}[\frac{1}{2}]$ by graphs which are unions of benzene 2-, 4- and 6-cycles. In particular, $\mathrm{gr}_p(\mathrm{Sym}^3(V_L)) = 0$ unless the parts of p are at most 6.*




FIGURE 8. A benzene 6-cycle.

We deduce Proposition 6.8 from the following:

Proposition 6.9. *Let L be a set of cardinality at least 8 and let Δ be a fixed matching on L . Then V_L is spanned by those Y_Γ for which the graph $\Delta\Gamma$ is not connected.*

Proof. Let Γ be a given matching. Think of Δ as having blue edges and Γ as having black edges, so $\Delta\Gamma$ is a regular 2-colored graph. We must show that by using only black Plücker relations we can write it as a sum of disconnected graphs. (We assume that $\Delta\Gamma$ is connected to begin with.)

Let a, b, c and d be four consecutive vertices in $\Delta\Gamma$, where \overline{ab} is a black edge. We have the black Plücker relation

$$0 = \begin{array}{c} b \text{---} c \\ / \quad \backslash \\ a \quad \quad d \end{array} + \begin{array}{c} c \text{---} b \\ / \quad \backslash \\ a \quad \quad d \end{array} + \frac{\overline{b \text{---} c}}{\overline{a \text{---} d}}$$

The rightmost graph is disconnected. Thus, working modulo such graphs, we may transpose blue edges at the cost of a sign. Now let a, b, c, d, e and f be six consecutive vertices where \overline{ab} is a black edge. We have the black Plücker relation

$$0 = \begin{array}{c} c \text{---} d \\ / \quad \backslash \\ b \quad \quad e \\ / \quad \backslash \\ a \quad \quad f \end{array} + \begin{array}{c} d \text{---} c \\ / \quad \backslash \\ e \quad \quad b \\ / \quad \backslash \\ a \quad \quad f \end{array} + \frac{\overline{c \text{---} d}}{\overline{a \text{---} f}}$$

As before, the rightmost graph is disconnected. This shows that we may take four consecutive vertices and reverse their direction, at the cost of a sign, assuming the outer two edges are blue. In the same way, by considering eight consecutive vertices we see that a consecutive string of six vertices may be flipped at the cost of a sign, assuming that the outer two edges are blue.

Now let a, b, c, d, e and f be six consecutive vertices, the outer two edges of which are blue. We write $[a, b, c, d, e, f]$ to denote this situation. By the above, we have

$$\begin{aligned} [a, b, c, d, e, f] &= -[f, e, d, c, b, a] = [f, e, a, b, c, d] \\ &= -[b, a, e, f, c, d] = [b, a, d, c, f, e] = -[a, b, c, d, e, f] \end{aligned}$$

so $Y_\Gamma = -Y_\Gamma$ modulo matchings for which $\Delta\Gamma$ is disconnected, establishing the proposition. \square

By keeping track of the discarded graphs in the above proof we obtain a complicated identity, shown in Figure 9, that we will use on a few later occasions. As an immediate corollary of the above proposition we have the following.

Corollary 6.10. *Let Γ be a regular 3-colored graph and let Φ be a benzene cycle in Γ . Then in $V_L^{\otimes 3}$ one can write Y_Γ as a linear combination of Y_{Γ_i} 's where in each Γ_i the subgraph Φ is replaced with a union of benzene 2-, 4- and 6-cycles. This holds over $\mathbb{Z}[\frac{1}{2}]$.*

Proof. Assume Φ has at least eight vertices (otherwise there is nothing to prove). Let Δ be the doubled edges of Φ and let Φ' be the single edges of Φ , so that $\Phi = \Delta\Phi'$. Use Proposition 6.9 to rewrite $Y_{\Phi'}$ as a sum of $Y_{\Phi'_i}$ with $\Delta\Phi'_i$ disconnected. This will rewrite Y_Γ as a sum of Y_{Γ_i} where in each Y_{Γ_i} the benzene cycle Φ has been replaced with a union of smaller benzene cycles. Continuing in this way, one deduces the proposition. \square

We can now prove Proposition 6.8.

Proof of Proposition 6.8. Let Γ be a regular 3-colored graph on L , thought of as having colors red, green and blue. Consider the red-green subgraph of Γ . We may apply Proposition 6.1 to rewrite this graph as a sum of graphs which are unions of 2-cycles. Thus we may as well assume that the red-green subgraph of




FIGURE 9. A graphical identity. The blue edges here are not relevant to the identity: they are just drawn to emphasize the disconnectedness of the graphs on the right side. Thus this is a linear relation between Y_Γ 's where the Γ are matchings on 8 points. This identity can be obtained by applying the procedure of Proposition 6.9 to the term on the left side, or by applying the straightening algorithm to the final term on the right side.

Γ is made up of 2-cycles. By now considering the blue edges as well, we see that Γ is a union of benzene cycles. We may now appeal to Corollary 6.10 to break up large benzene cycles into smaller ones. \square

6.6. The action of \mathfrak{S}_L on degree three spaces. We now examine the spaces $\text{gr}_p(\text{Sym}^3(V_L))$ more closely when L is small and determine their structure as \mathfrak{S}_L -modules. We assume throughout that $|L|!$ is invertible. We denote by M_L the $\mathbb{Z}[\mathfrak{S}_L]$ -module with a basis given by the set of undirected matchings on L .

Proposition 6.11. *As an \mathfrak{S}_L -representation the space M_L is multiplicity free and contains those irreducibles corresponding to partitions of $|L|$ into even parts.*

This proposition is well-known and essentially equivalent to the decomposition of the plethysm $\text{Sym}^n \text{Sym}^2$. In any case, we only need to use this when L has cardinality four or six, where it can easily be established by hand. We now begin our study of $\text{gr}_p(\text{Sym}^3(V_L))$. For the sake of brevity, we will denote this space simply by gr_p in this section. We will also write \mathfrak{S}_n in place of \mathfrak{S}_L where $n = |L|$. Our first result is the following:

Proposition 6.12. *The space $\text{gr}_{2+\dots+2}$ is multiplicity free and contains only those irreducible representations of \mathfrak{S}_n which have an even number of parts.*

Proof. The space $\text{gr}_{2+\dots+2}$ is the image of the map $\epsilon \otimes M_L \rightarrow \text{Sym}^3(V_L)$ which takes a matching Γ to Y_Γ^3 . The result now follows from Proposition 6.11. \square

We will only need to use the above proposition for $n \leq 6$. We now determine the spaces gr_p completely for $n = 4$ and 6. As a warm-up, consider the $n = 2$ case: gr_2 is one-dimensional, and it is the sign representation of \mathfrak{S}_2 . Reason: the space $\text{Sym}^3(V_L)$ is spanned by Y_Γ where Γ is the graph on L with a tripled edge. The group \mathfrak{S}_2 fixes Γ and so acts on Y_Γ through the sign character. All other gr_p in this case are zero. We now turn to $n = 4$.

Proposition 6.13. (a) *The space $\text{gr}_{2,2}$ is free over $\mathbb{Z}[\frac{1}{4!}]$ and three-dimensional. As an \mathfrak{S}_4 -representation it decomposes into two irreducibles corresponding to the partitions $2 + 2$ and $1 + 1 + 1 + 1$.*

(b) *The space gr_4 is free over $\mathbb{Z}[\frac{1}{4!}]$ and one-dimensional. The representation of \mathfrak{S}_4 on it is trivial.*

Proof. The space gr_4 is spanned by benzene 4-cycles. We have the identity

$$0 = \begin{array}{|c|} \hline \color{red}{\square} \\ \hline \end{array} + \begin{array}{|c|} \hline \color{blue}{\square} \\ \hline \end{array} + \begin{array}{|c|} \hline \color{green}{\square} \\ \hline \end{array}$$

obtained by Plückerizing the blue edges. We thus see that if Γ is benzene 4-cycle in which \overline{ab} is a doubled edge then $Y_\Gamma + Y_{(ab)\Gamma} = 0$ in gr_4 , which can be rephrased as $(ab)Y_\Gamma = Y_\Gamma$. In particular, two benzene 4-cycles are equal in gr_4 (up to a possible sign) if they have the same doubled edges. By cubing

$$\begin{array}{|c|} \hline \color{red}{\times} \\ \hline \end{array} = \begin{array}{|c|} \hline \color{red}{\text{---}} \\ \hline \end{array} + \begin{array}{|c|} \hline \color{red}{\text{---}} \\ \hline \end{array}$$

we obtain

$$0 = \begin{array}{|c|} \hline \color{red}{\square} \\ \hline \end{array} + \begin{array}{|c|} \hline \color{blue}{\square} \\ \hline \end{array}$$

in gr_4 (all other terms belong to $F_{2,2}$). This shows that we can switch which edges are the doubled edges. We have thus shown that gr_4 is one dimensional. Since we already know that some transpositions act by the identity, it follows that \mathfrak{S}_4 acts trivially on gr_4 .

Now, $\text{Sym}^3(V_L)$ is four-dimensional. We know that gr_4 is one-dimensional and so $\text{gr}_{2,2}$ must be three-dimensional. By Proposition 6.12 we know that $\text{gr}_{2,2}$ is a quotient of the direct sum of the irreducible representations of \mathfrak{S}_4 corresponding to $2 + 2$ and $1 + 1 + 1 + 1$. Since this direct sum has dimension 3, it follows that $\text{gr}_{2,2}$ must equal it. \square

Finally we consider the case $n = 6$. As stated in §1.6, the computer calculations in the proof are very mild and could probably be done by hand in a matter of hours.

- Proposition 6.14.** (a) *The space $\text{gr}_{2,2,2}$ is free over $\mathbb{Z}[\frac{1}{6!}]$ and 15-dimensional. As an \mathfrak{S}_6 -representation it decomposes into three irreducibles, corresponding to the partitions $3 + 3$, $2 + 2 + 1 + 1$ and $1 + 1 + 1 + 1 + 1 + 1$.*
(b) *The space $\text{gr}_{4,2}$ is free over $\mathbb{Z}[\frac{1}{6!}]$, 15-dimensional and decomposes into two irreducibles: one corresponding to $5 + 1$ and the other to $4 + 1 + 1$.*
(c) *The space gr_6 is free over $\mathbb{Z}[\frac{1}{6!}]$ and five-dimensional. It is irreducible and corresponds to $3 + 3$.*
(d) *Let Q be a set of cardinality six and let c, d and e be three distinct elements of Q . Then $\text{gr}_6(\text{Sym}^3(V_Q))$ is spanned by benzene 6-cycles in which cd or ce appears as a doubled edge. (We call the set Q here, rather than L , since that is what it will be called in the one place where we apply this statement.)*

Proof. The code for the computer calculations required here can be found on the webpage [HMSV5].

(a) A computer calculation shows that $F_{2,2,2} = \text{gr}_{2,2,2}$ is 15 dimensional. By Proposition 6.12, the space $\text{gr}_{2,2,2}$ is a quotient of the direct sum of the irreducible representations of \mathfrak{S}_6 corresponding to the partitions given in the statement of this proposition. Since this direct sum is also 15 dimensional the quotient map is an isomorphism.

(b) A computer calculation shows that $F_{4,2}$ is 30 dimensional, and so $\text{gr}_{4,2} = F_{4,2}/F_{2,2,2}$ is 15 dimensional. Now observe that there is a map $F_4 \otimes F_2 \rightarrow F_{4,2}$ which takes an element of F_4 on the vertices $\{1, 2, 3, 4\}$ and appends a tripled edge on the vertices $\{5, 6\}$. One sees using this that that $\text{gr}_{4,2}$ is a quotient of $\text{Ind}_{\mathfrak{S}_4 \times \mathfrak{S}_2}^{\mathfrak{S}_6}(\text{gr}_4 \otimes \text{gr}_2)$. By the Littlewood-Richardson rule, the induction is a direct sum of the two irreducible representations of \mathfrak{S}_6 corresponding to the partitions given in the statement of the proposition. Since this sum is also 15 dimensional the quotient map is an isomorphism.

(c) We have a non-canonical \mathfrak{S}_6 -equivariant decomposition $\text{Sym}^3(V_L) \cong \text{gr}_{2,2,2} \oplus \text{gr}_{4,2} \oplus \text{gr}_6$. On the other hand, a character computation shows that $\text{Sym}^3(V_L) \cong \text{gr}_{2,2,2} \oplus \text{gr}_{4,2} \oplus M_{3+3}$ (where $M_{3,3}$ is the irreducible representation corresponding to $3 + 3$). Thus $\text{gr}_6 = M_{3+3}$.

(d) This is a straightforward computer calculation. \square

7. RETROGENERATION OF THE IDEAL

In this section, we prove that for $|L|$ sufficiently large, the ideal of relations is retrogenerated (generated by quadratics and relations on fewer points):

Theorem 7.1. *Let L be an even set of cardinality at least 10. Then $I_L = I_L^{\text{retro}}$. For $|L| \geq 12$ this holds over $\mathbb{Z}[\frac{1}{2}]$ while for $|L| = 10$ it holds over $\mathbb{Z}[\frac{1}{10!}]$.*

The ideal I_L^{retro} of retrogenerated relations will be defined in §7.1. We use general structural arguments to prove Theorem 7.1 when $|L| \geq 12$ but when $|L| = 10$ the small size of the graphs involved forces us to give an inelegant ad hoc argument. We suggest that a reader consider skipping the arguments in the $|L| = 10$ case on a first reading.

7.1. Outer multiplication and the retrogenerated ideal I_L^{retro} . Let L and L' be even sets. We have an *outer multiplication* map

$$\boxtimes : V_L^{\otimes n} \otimes V_{L'}^{\otimes n} \rightarrow V_{L \amalg L'}^{\otimes n}, \quad Y_\Gamma \otimes Y_{\Gamma'} \mapsto Y_{\Gamma \amalg \Gamma'}.$$

In words, one takes a colored graph Γ on L and a colored graph Γ' on L' , with the same set of colors, and obtains a colored graph on $L \amalg L'$ by taking the disjoint union of Γ and Γ' . We will often omit the symbol \boxtimes and write the outer product using juxtaposition.

Outer multiplication does not descend to symmetric powers, as shown by the following example:

Let L and L' be the two sets of four vertices occurring on the left sides. The two left sides above define equal elements of $\text{Sym}^2(V_L) \otimes \text{Sym}^2(V_{L'})$. However, the two right sides are different elements of $\text{Sym}^2(V_{L \amalg L'})$ — their difference is the simplest binomial relation (1.4).

One sees from the above example that outer multiplication does not descend to symmetric powers for the following reason: if Γ is an n -colored graph on L and one permutes the colors in each connected component of Γ to obtain a new n -colored graph Γ' then Y_Γ and $Y_{\Gamma'}$ do not represent the same element of $\text{Sym}^n(V_L)$ in general. Now, if Γ and Γ' are as in the previous sentence then $Y_\Gamma - Y_{\Gamma'}$ lies in the ideal Q_L of $\text{Sym}(V_L)$ generated by quadratic relations. (Reason: any permutation of colors can be obtained by successive transpositions of colors, and the relations thus arising are clearly quadratic.) We hence find that outer multiplication descends to a map

$$\boxtimes : \text{Sym}^n(V_L)/Q_L^{(n)} \otimes \text{Sym}^n(V_{L'})/Q_{L'}^{(n)} \rightarrow \text{Sym}^n(V_{L \amalg L'})/Q_{L \amalg L'}^{(n)}$$

and that $I_L^{(n)}/Q_L^{(n)} \otimes \text{Sym}^n(V_{L'})/Q_{L'}^{(n)}$ is mapped into $I_{L \amalg L'}^{(n)}/Q_{L \amalg L'}^{(n)}$ under \boxtimes , that is, the outer product of anything with a relation is still a relation. A motivating example appeared in the introduction: the Segre relation on 6 points (1.2) induces a relation on 8 points (1.3).

For a given L , we define the *ideal of retrogenerated relations*, denoted I_L^{retro} , to be the ideal of $\text{Sym}(V_L)/Q_L$ generated by the images of $I_{L'}^{(n)}/Q_{L'}^{(n)} \otimes \text{Sym}^n(V_{L''})/Q_{L''}^{(n)}$ under \boxtimes as (L', L'') varies over all partitions of L into two disjoint proper even subsets and n varies over all positive integers. We also write I_L^{retro} for the inverse image of I_L^{retro} under $\text{Sym}(V_L) \rightarrow \text{Sym}(V_L)/Q_L$. We have inclusions $Q_L \subset I_L^{\text{retro}} \subset I_L$. (Theorem 8.1 will show that these three ideals are all the same.) We say a relation (in I_L) is *retrogenerated* if it lies in I_L^{retro} .

A basic fact is that outer multiplication does not increase the “essential degree” of a relation:

Proposition 7.2. *Let $L = L' \amalg L''$ be a partition of L into two proper even subsets. Let $x \in I_{L'}^{(n)}/Q_{L'}^{(n)}$ belong to the ideal generated by relations of degree $\leq k$ and let $y \in \text{Sym}^n(V_{L''})/Q_{L''}^{(n)}$. Then the outer product $x \boxtimes y$ belongs to the ideal of $\text{Sym}(V_L)/Q_L$ generated by relations of degree $\leq k$.*

Proof. Write $x = \sum a_i r_i$ where r_i is a relation of degree $\leq k$ and a_i belongs to $\text{Sym}(V_{L'})/Q_{L'}$ and write $y = \sum y_j$ where each y_j is a product of degree one elements. The outer product of $x \boxtimes y$ is a sum of terms of the form $a_i r_i \boxtimes y_j$. Say r_i has degree $k' \leq k$ so that a_i has degree $n - k'$, and write $y_j = b_1 \cdots b_n$. Then $a_i r_i \boxtimes y_j = (a_i \boxtimes (b_1 \cdots b_{n-k'}))(r_i \boxtimes (b_{n-k'+1} \cdots b_n))$ — this is the basic compatibility between outer multiplication and usual multiplication and is trivial to verify. We have thus shown that $a_i r_i \boxtimes y_j$ is a multiple of $r_i \boxtimes (b_{n-k'+1} \cdots b_n)$, a relation of degree $k' \leq k$. \square

Outer multiplication, simple binomial relations and the retrogenerated ideal are quite formal constructions and are present when studying $X^n//G$ for any X and G . They are described nicely by the formalism of third author mentioned in §1.7. By contrast, the following two propositions are specific to the present case.

Proposition 7.3. *Let $\Sigma = (\Gamma, \mathcal{U})$ be a generalized Segre datum for which Γ is disconnected. Then $\text{Rel}(\Sigma)$ belongs to I_L^{etro} .*

Proof. Write $\Gamma = \Gamma_1 \amalg \Gamma_2$. If one of Γ_1 or Γ_2 is entirely contained within one of the parts U_R, U_G or U_B then the relation $\text{Rel}(\Sigma)$ is manifestly retrogenerated. For instance, if Γ_1 is contained within one of the parts then Γ_2 with the partition induced from \mathcal{U} forms a generalized Segre datum Σ_2 and $\text{Rel}(\Sigma)$ is the outer product of Y_{Γ_1} and $\text{Rel}(\Sigma_2)$. If neither Γ_1 nor Γ_2 is contained solely within one part then each contains a pair of special edges and the datum Σ is forced to be degenerate (of the second case given in §5.4). The relation $\text{Rel}(\Sigma)$ thus belongs to $Q_L \subset I_L^{\text{etro}}$ by Proposition 5.2. \square

The next proposition is a key point in our inductive arguments.

Proposition 7.4. *The ideal $I_L^{\text{etro}} \subset \text{Sym}(V_L)/Q_L$ is generated over $\mathbb{Z}[\frac{1}{2}]$ by the images of $I_{L'}^{(3)}/Q_{L'}^{(3)} \otimes \text{Sym}^3(V_{L''})/Q_{L''}^{(3)}$ under \boxtimes as (L', L'') varies over all partitions of L into two disjoint subsets where L'' has cardinality 2, 4 or 6.*

Proof. Proposition 7.2 shows that I_L^{etro} is generated by elements of the form $r \boxtimes Y_\Gamma$ where r belongs to $I_{L'}^{(3)}/Q_{L'}^{(3)}$ and Y_Γ belongs to $\text{Sym}^3(V_{L''})/Q_{L''}^{(3)}$, as (L', L'') varies over all partitions of L into two disjoint even subsets. Proposition 6.8 shows that we can write $Y_\Gamma = \sum a_i Y_{\Gamma_i}$ where a_i belongs to $\mathbb{Z}[\frac{1}{2}]$ and each Γ_i is a union of benzene 2-, 4- and 6-cycles. It follows that Y_{Γ_i} is itself an outer product $Y_{\Gamma_{i1}} \boxtimes \cdots \boxtimes Y_{\Gamma_{in}}$ where each $Y_{\Gamma_{ij}}$ is a benzene 2-, 4- or 6-cycle. Since \boxtimes is associative, we have $r \boxtimes Y_{\Gamma_i} = (r \boxtimes Y_{\Gamma_{i1}} \boxtimes \cdots \boxtimes Y_{\Gamma_{i(n-1)}}) \boxtimes Y_{\Gamma_{in}}$, which expresses $r \boxtimes Y_{\Gamma_i}$ as the outer product of a graph on 2, 4 or 6 fewer points with a graph on 2, 4 or 6 points. Thus $r \boxtimes Y_\Gamma$ is a sum of such graphs, which establishes the proposition. \square

7.2. Square rotation relations. To prove Theorem 7.1 we introduce the *square rotation relations*, needed only in this proof. A *square rotation datum* is a pair $\Pi = (\Gamma, U)$ where U is a subset of L of cardinality 4 and Γ is an undirected graph on L with edges colored purple and black such that:

- The vertices of $L \setminus U$ have black valence one and purple valence two.
- The vertices of U have black valence zero and purple valence one.

Let Π be a given square rotation datum. Suppose that a pair of edges e and e' in Γ have the same color. Let $\Gamma + \Gamma' + \Gamma'' = 0$ be the Plücker relation on e and e' . Then $\Pi' = (\Gamma', U)$ and $\Pi'' = (\Gamma'', U)$ are both square rotation data. We define the *space of square rotation data* to be the \mathbb{Z} -span of the square rotation data modulo the relations $\Pi + \Pi' + \Pi'' = 0$.

Let Π be a square rotation datum. We have the following quadratic relation on the four points in U :

$$(7.1) \quad \begin{array}{|c|} \hline \square \\ \hline \end{array} = \begin{array}{|c|} \hline \square \\ \hline \end{array}$$

Multiplying both sides by Γ we obtain an element of $R_L^{(1)} \otimes R_L^{(2)}$ which maps to zero in R_L . We may thus regard it as an element of $I_L^{(3)}/Q_L^{(3)}$, similar to what we did for generalized Segre relations. We call such relations *square rotation relations*. We have a linear map

$$\text{Rel} : \{\text{the space of square rotation data}\} \rightarrow I_L^{(3)}/Q_L^{(3)}$$

mapping a square rotation datum to its associated relation.

7.3. Retrogeneration of square rotation relations. The result is:

Proposition 7.5. *If L has cardinality at least 12 then any square rotation relation is retrogenerated.*

We use a lemma to prove the proposition. Let Π be a given square rotation datum. The purple subgraph of Γ has valence two everywhere except at the four vertices in U . It thus breaks up into a union of cycles and two paths terminating in U . We call these two paths the *special paths*. The *length* of a special path is the number of vertices it contains.


Lemma 7.6. (a) *A square rotation relation with both special paths of even length lies in Q_L .*
 (b) *A square rotation relation with one special path of even length is quadratic.*

Proof. (a) The way we obtain a square rotation relation from the square rotation datum is to simply append the graph Γ to the relation (7.1). Now, on each side of (7.1) there are two purple edges. Change the color of these edges to grey. The hypothesis on the lengths of the special paths says that we can color every other edge in the special paths grey (keeping the first and last edges purple) so that each vertex belongs to one grey and one purple edge. We may assume that the purple cycles occurring in Γ have even size (as we can force this using Plücker relations), so we can pick a purple-grey alternating coloring of them. Thus we have factored the purple subgraph of Γ in such a way that all the purple edges appearing in (7.1) have been colored grey. But now the relation is evidently essentially quadratic (i.e., belongs to Q_L) since it is taking place solely on the black-grey graph.

(b) If one special path has even length and one has odd length then Γ has an odd purple cycle. We can thus Plücker the odd special path and the odd cycle so that both special paths have even length. We now use part (a). \square

We now prove Proposition 7.5.

Proof of Proposition 7.5. Let a square rotation datum be given. By the above lemma, we can assume that each of the special paths has odd length. Using arguments similar to those occurring in the proof of Proposition 6.1, we can force the special paths to have lengths three. Thus each special path contains a single vertex not belonging to U . Call these two vertices x and y . We now use Proposition 6.1 to convert the purple cycles in Γ into 2-cycles. We thus have a benzene chain joining x with y . The graph Γ now looks like the following (with the possibility that there are some additional black-purple benzene cycles not pictured):



Here the numbered vertices constitute the set U . The special paths are 1- x -2 and 3- y -4. If there are in fact benzene cycles in Γ then the relation is retrogenerated (the proof of this is similar to that of Proposition 7.3). We can thus assume that there are no benzene cycles and so the graph really does look like the above one. Since we have at least 12 vertices, the benzene chain will have at least four single black edges, so we can apply Corollary 6.10 or identity of Figure 9 to break up the benzene chain and get a disconnected graph. The associated relation will therefore be retrogenerated. \square

7.4. Retrogeneration of the ideal on 12 points (Theorem 7.1 with $|L| = 12$). We begin with two lemmas.

Lemma 7.7. *Let $\Sigma = (\Gamma, \mathcal{U})$ be a generalized Segre datum, and let $U \subset U_G$ be a set of four vertices such that the blue-green graph of Γ has a 4-cycle contained in U . Let $\bar{\Gamma}$ be the graph obtained by rotating the colors in this 4-cycle and let $\bar{\Sigma} = (\bar{\Gamma}, \mathcal{U})$. Then $\text{Rel}(\Sigma) \equiv \text{Rel}(\bar{\Sigma})$ modulo quadratic and square rotation relations.*


Proof. We must show that

$$\text{Rel}(\Sigma) - \text{Rel}(\bar{\Sigma}) = Y_\Gamma - Y'_\Gamma - Y_{\bar{\Gamma}} + Y'_{\bar{\Gamma}}$$

belongs to the ideal generated by quadratic and square rotation relations. Clearly, $Y_\Gamma - Y_{\bar{\Gamma}} \in Q_L$ since the red subgraph in each is the same, while $Y'_\Gamma - Y'_{\bar{\Gamma}}$ is a square rotation relation by definition. \square

Lemma 7.8. *Let $\Sigma = (\Gamma, \mathcal{U})$ be a generalized Segre datum and let $U \subset U_G$ be a set of four vertices such that the blue-green graph of Γ has a 4-cycle contained in U . Let $\{\Gamma_i\}$ be the three graphs obtained by replacing the 4-cycle on U by two 2-cycles (there are three ways to do this). Let $\{\Sigma_i\}$ be the corresponding generalized Segre data. Then $2 \text{Rel}(\Sigma) \equiv \sum \text{Rel}(\Sigma_i)$ modulo quadratic and square rotation relations.*

Proof. Recall the identity



of the proof of Lemma 6.3. This identity holds in $V_L^{\otimes 2}$, that is, it follows from the colored Plücker relations. This shows that in the space of generalized Segre data, we have $\Sigma + \bar{\Sigma} = \sum_{i=1}^3 \Sigma_i$ where $\bar{\Sigma}$ was defined in the previous lemma. By the previous lemma, we have $\text{Rel}(\Sigma) = \text{Rel}(\bar{\Sigma})$ modulo quadratic and square rotation relations. This proves the current lemma. \square


We now complete the proof of Theorem 7.1 when $|L| \geq 12$.

Proof of Theorem 7.1 when $|L| \geq 12$. By Theorem 5.1, it suffices to show that small generalized Segre cubic relations are retrogenerated. Thus let Σ be a given small generalized Segre datum. We assume without loss of generality that U_R has cardinality two, and that Σ is non-degenerate.

We begin by considering the blue-green subgraph of U_G . This is a union of cycles and a single chain going between the two special blue edges. By using Proposition 6.1, or more accurately the identity of Figure 7, we can convert this graph into a union of 2-cycles and 4-cycles and make the chain have length three (so that the two special blue edges are joined by a single green edge). We can now use the above two lemmas to convert the 4-cycles into 2-cycles, modulo retrogenerated relations.

We have thus reduced to the case where the blue-green subgraph of Γ in U_G is made up of 2-cycles (except for the path of length three involving the two special blue edges). We now consider the red edges. Except for the special blue edges, the graph $\Gamma|_{U_G}$ is a union of benzene chains with the blue and green edges being paired. There are two incoming special red edges and two red edges connected to the two special blue edges. From each of these edges a benzene chain emanates which must terminate at one of the other edges. These four red edges are thus contained in two benzene chains. If the two special red edges are contained in the same benzene chain then the graph is degenerate and the relation lies in Q_L . It therefore suffices to consider the case where each special red edge connects to a special blue edge via a benzene chain.

Now, we may run the entire argument given above inside U_B as well. We thus conclude that $\Gamma|_{U_B}$ is a union of benzene chains with the blue and green edges being paired. The two special red edges connect with the two special green edges by benzene chains. The picture is thus something like:



Here $U_G = \{1, 2, 3, 4\}$, $U_B = \{5, 6, 7, 8, 9, 10\}$ and $U_R = \{11, 12\}$. If there are other benzene cycles present, then the relation is immediately retrogenerated. Also, the two benzene chains connecting the left and right side could be crossed; that is, the chains could go from 1 to 7 and 2 to 8 instead of as they do. However, one can always rectify this by Plückerizing two red edges such as $\overline{89}$ and $\overline{67}$ — one of the resulting graphs has the chains uncrossed while the other is degenerate. So the above graph is the only sort we need consider.

Now, we can repartition the vertices so that all the blue-green doubled edges are contained in U_G . For example, with the above graph we would repartition so that $U_G = \{1, 2, 3, 4, 5, 6, 9, 10\}$, $U_B = \{7, 8\}$ and $U_R = \{11, 12\}$. The resulting Segre datum yields the same relation as the original. We next point out that we can put all the blue-green doubled edges into a single benzene chain by Plückerizing two red edges. For example, in the above graph we would Plücker the edges $\overline{110}$ and $\overline{67}$. In one graph, $\overline{17}$ is a red edge and all the blue-green doubled edges are in a single benzene chain running between 2 and 8. The other graph is degenerate. Now, since we have at least 12 vertices, the benzene chain will have at least four single red edges. It can therefore be broken apart using Corollary 6.10 or identity of Figure 9, yielding a disconnected graph and therefore a retrogenerated relation. This completes the proof of Theorem 7.1 when $|L| \geq 12$. \square

7.5. Retrogeneration of the ideal on 10 points (Theorem 7.1 with $|L| = 10$). For the remainder of §7 we let L be a finite set of cardinality 10. To prove Theorem 7.1 in this case, we show the following:

Proposition 7.9. *The space $I_L^{(3)}/I_L^{\text{retro},(3)}$ has dimension at most two (over $\mathbb{Z}[\frac{1}{2}]$).*

We now explain why this implies Theorem 7.1 for $|L| = 10$. For the next few sentences we work over a field k of characteristic not 2, 3, 5 or 7. A character computation shows that $\text{Sym}^3(V_L)$ does not contain the trivial or alternating representation. Now, the representation $I_L^{(3)}/I_L^{\text{retro},(3)}$ is at most two-dimensional, and a summand of $\text{Sym}^3(V_L)$. As any representation of \mathfrak{S}_L of dimension at most two is made up of some combination of the trivial and alternating representations, it follows that $I_L^{(3)}/I_L^{\text{retro},(3)}$ vanishes. Thus $I_L^{(3)} = I_L^{\text{retro},(3)}$ over k . Since this holds for all k of characteristic not 2, 3, 5 or 7, we conclude $I_L^{(3)} = I_L^{\text{retro},(3)}$. Since I_L is generated by quadratics and $I_L^{(3)}$ (Theorem 5.1), this implies $I_L = I_L^{\text{retro}}$.


We now turn to proving Proposition 7.9. We prove the following result:

Proposition 7.10. *Let $\mathcal{U} = \{U_G, U_R, U_B\}$ be a fixed partition of L with $|U_R| = 2$. Consider the subspace V of $I_L^{(3)}/I_L^{\text{retro},(3)}$ spanned by all generalized Segre cubic relations coming from data with partition equal to \mathcal{U} . Then, over $\mathbb{Z}[\frac{1}{2}]$:*

- (1) *If U_G (and thus U_B) has cardinality four then $V = 0$.*
- (2) *If either U_G or U_B has cardinality six then $\dim V \leq 1$.*


We explain why Proposition 7.10 implies Proposition 7.9. Fix an order on L . Call a partition $\mathcal{U} = \{U_G, U_R, U_B\}$ of L *admissible* if each part is non-empty of even cardinality, $U_G < U_R < U_B$ in the order and $|U_R| = 2$. We call a generalized Segre cubic datum *admissible* if its partition is is; we extend the notion to relations in the obvious manner. By inspection, there are three admissible partitions of L . One of these has $|U_G| = |U_B| = 4$ while in the other two, one of U_G or U_B has cardinality six. It follows from Proposition 7.10 that the admissible Segre cubic relations span a subspace of $I_L^{(3)}/I_L^{\text{retro},(3)}$ of dimension at most two. On the other hand, we know that the admissible generalized Segre cubic relations span $I_L^{(3)}/Q_L^{(3)}$ (see Remark 5.3). As $I_L^{\text{retro},(3)}$ contains $Q_L^{(3)}$ we conclude that $I_L^{(3)}/I_L^{\text{retro},(3)}$ has dimension at most two, establishing Proposition 7.9.

We now begin proving Proposition 7.10. We first consider the case where $|U_G| = 4$ (so $|U_B| = 4$ too). Consider the blue-green subgraph of U_G . In it there are two green edges, two blue edges contained entirely in U_G and two blue edges (the special ones) going between U_G and U_R . By using identity of Figure 7 we can force there to be a blue-green doubled edge. This implies that the special blue edges are joined by one green edge. By now considering the red edges, we see that there are two possibilities for the picture in U_G :




The right case is disconnected, hence retrogenerated. We thus need only consider the left case. We now go through the same considerations in U_B as we just did in U_G and conclude that it too must look like the graph on the left, except with the colors blue and green reversed. We find that the graph as a whole must be one of the two in the statement of the following lemma. That lemma then shows that that the generalized Segre relation we are considering is retrogenerated, which completes the $|U_G| = 4$ case.

Lemma 7.11. *Let Γ be one of the following two graphs.*



Let $\mathcal{U} = \{U_R, U_G, U_B\}$ be a partition such that $U_R = \{9, 10\}$, $U_G \supset \{1, 2\}$, $U_B \supset \{3, 4\}$ and $\Sigma = (\Gamma, \mathcal{U})$ is a generalized Segre cubic datum. Then $\text{Rel}(\Sigma)$ is retrogenerated over $\mathbb{Z}[\frac{1}{2}]$.

Proof. We first note that the relation $\text{Rel}(\Sigma)$ is independent of the partition \mathcal{U} . We therefore assume that the vertices 5, 6, 7 and 8 belong to U_G . Now, let Γ be the left graph. Redraw Γ as follows:



All the red edges other than $\overline{910}$ contain at least one vertex in U_G so we are allowed to Plücker any of them together to obtain a relation of generalized Segre cubic data. We thus may apply identity of Figure 9 to the red edges. All the resulting terms either come from fewer points by outer multiplication or else are degenerate. (Precisely: the first two graphs from the first row of Figure 9 come from fewer points, and the

last two are degenerate. In the second row, all are degenerate except for the third. In the third row, all are degenerate except the first.) We conclude that $\text{Rel}(\Sigma)$ is retrogenerated.


Now let Γ be the right graph. By applying the Plücker relation to the edges $\overline{28}$ and $\overline{36}$ we get two graphs, one of which is degenerate, the other of which looks like the graph handled in the previous paragraph. Thus $\text{Rel}(\Sigma)$ for the right graph is retrogenerated as well. \square

We now establish the second case of Proposition 7.10. Thus let \mathcal{U} be a partition in which $|U_G| = 6$. In this case, both U_R and U_B have cardinality two and so the situation is symmetric with respect to these two colors. We begin with the following observation:


Lemma 7.12. *If there is a green-blue or green-red doubled edge in U_G then $\text{Rel}(\Sigma)$ is retrogenerated over $\mathbb{Z}[\frac{1}{2}]$.*

Proof. Say there is a green-blue doubled edge in U_G . Then, as in the case $|U_G| = 4$, we can apply identity of Figure 7 to force there to be a second green-blue doubled edge in U_G . The graph Γ must now look like one of the two in Lemma 7.11 and so, by that lemma, we obtain the present one. \square

Now consider a general generalized Segre cubic datum Σ with $|U_G| = 6$. By applying identity of Figure 7 repeatedly, or appealing to Proposition 6.1, we can arrange it so that the blue-green subgraph on U_G is of the form



The right graph has a blue-green doubled edge, and so the associated relation belongs to I'_L by Lemma 7.12. We therefore need only consider the left graph. We now consider the red edges. There are two that are completely contained in U_G and the two special edges, each of which has one vertex in U_G . Here are three possible graphs:



We will refer to these as type A, B and C, respectively. Of course, there are many other possibilities for what the graph could look like. However, all reduce to one of the above three types after some simple Plücker relations. For instance, one could consider the graph which is like the type A one, but where the blue and green colors are switched in the square 2-3-5-6. By applying the Plücker relation to the edges $\overline{45}$ and the red special edge containing 3, this relation is rewritten as a sum of two of those appearing above.

We now have the following result:


Lemma 7.13. *Type B and C relations are retrogenerated over $\mathbb{Z}[\frac{1}{2}]$.*

Proof. Consider the type B graph drawn above. We apply identity of Figure 7 to the red-green chain 4-5-2-3-6-*, where * is the relevant vertex of U_B . The first graph on the right side of the identity is disconnected and therefore retrogenerated. All the remaining graphs have green-red doubled edges and so are also retrogenerated by Lemma 7.12. This completes the type B case. The type C case is handled similarly: in the above labeling one applies the identity of Figure 7 to the chain 4-5-2-6-3-* and then proceeds exactly as in the type B case. \square

The following completes the proof of Proposition 7.10.

Lemma 7.14. *The type A relations span a subspace of $I_L^{(3)}/I_L^{\text{retro},(3)}$ of dimension at most one over $\mathbb{Z}[\frac{1}{2}]$.*

Proof. Put $U_G = \{1, 2, 3, 4, 5, 6\}$ and let Γ be a 3-colored graph on L whose restriction to U_G is the type A graph



Applying the Plücker relation to the edges $\overline{45}$ and $\overline{12}$ gives an expression for Γ in terms of two other graphs Γ_1 and Γ_2 , one of which (say Γ_1) is degenerate. Now apply the Plücker relation to the edges $\overline{56}$ and $\overline{23}$ of

Γ_2 . Again, this expresses Γ_2 in terms of two other graphs Γ_3 and Γ_4 , one of which (say Γ_4) is degenerate. Thus, we see that $\text{Rel}(\Sigma)$ and $\text{Rel}(\Sigma_4)$ are scalar multiples of each other. But $\text{Rel}(\Sigma_4)$ is nothing other than (25) $\text{Rel}(\Sigma)$. We thus see that Γ and (25) Γ yield the same relation (up to a scalar). By similar reasoning, we see that (14) Γ and (36) Γ give the same relation as Γ .

Now consider the Plücker relation on Γ on the edges $\overline{12}$ and $\overline{3*}$, where $*$ is the relevant vertex of U_B . This results in two graphs Γ_1 and Γ_2 . Say Γ_1 is the graph which has the edge $\overline{1*}$. Then Γ_1 is a type B graph and thus retrogenerated. Now apply the Plücker relation on the edges $\overline{25}$ and $\overline{36}$ to Γ_2 to get two new graphs Γ_3 and Γ_4 . One of these, say Γ_3 , has a green-blue doubled edge and is thus retrogenerated. The other, Γ_4 , is just (23) Γ . We thus see that Γ and (23) Γ give the same relation, up to a scalar. Similar reasoning shows that (45) Γ , (56) Γ and (12) Γ give the same relation as Γ .

From the previous two paragraphs, we see that if σ is any permutation of U_G then Γ and $\sigma\Gamma$ give the same relation, up to a scalar. Since every type A graph is of the form $\sigma\Gamma$ for some σ , this establishes the lemma. \square

8. THE IDEAL IS GENERATED BY QUADRATICS

The goal of §8 is to prove that I_L is generated by quadratics if $|L| \geq 8$:

Theorem 8.1. *If L is an even set of cardinality at least eight then $I_L = Q_L$ over $\mathbb{Z}[\frac{1}{12!}]$.*

As described in §1.6, this concludes the proof of Theorem 1.1. Thanks to our work in §7, the main remaining work is to prove three base cases:

Theorem 8.2. *The ideals I_8 , I_{10} and I_{12} are generated by quadratics over $\mathbb{Z}[\frac{1}{12!}]$.*

This implies Theorem 8.1 (and hence Theorem 1.1) by the following inductive argument. Let $n \geq 14$ and assume the theorem has been established for even sets of cardinality less than n . By Theorem 7.1 we have $I_n = I_n^{\text{retro}}$. By Proposition 7.4 we know that I_n^{retro} is generated by the outer products of cubic relations on $n-2$, $n-4$ and $n-6$ points with arbitrary graphs on 2, 4 and 6 points. The inductive hypothesis, and the fact that $n \geq 14$, ensures that any cubic relation on $n-2$, $n-4$ or $n-6$ points lies in the ideal generated by quadratics. Since outer multiplication does not increase essential degree (Proposition 7.2), we find $I_n^{\text{retro}} = Q_n$. Thus Theorem 8.1 is established by induction.


As for Theorem 8.2, quadratic generation of I_8 is proved over $\mathbb{Z}[\frac{1}{3}]$ in [HMSV3]. (Quadratic generation of I_8 over \mathbb{Q} had previously been established by computer; see e.g., [Ko] or [FS, Lemma 1.1].) We handle the 10 and 12 point cases by showing that I_{10} (resp. I_{12}) is generated by relations coming from 8 (resp. 8 and 10) points: one does not need relations coming from 6 points to generate I_L . The quadratic generation of I_8 then implies that of I_{10} and hence of I_{12} as well.

Here is an overview of the argument. We let I' be the quotient of the cubic part of the ideal by quadratic relations and relations retrogenerated from 8 (resp. 8 and 10) points. We wish to show that this space is zero. By Theorem 7.1, I' is generated as an \mathfrak{S}_{10} -module (resp. \mathfrak{S}_{12} -module) by a single relation: the outer product of the Segre cubic on 6 points by a benzene 4-cycle (resp. 6-cycle). (Hence we only need to prove that these two relations are generated by quadratics. This could in principle be checked by computer, but with current technology would probably require an algorithm of equal difficulty to our proof!) We can thus write I' as a quotient of a representation V induced from $\mathfrak{S}_6 \times \mathfrak{S}_4$ (resp. $\mathfrak{S}_6 \times \mathfrak{S}_6$), where the first factor acts on the Segre cubic on 6 points and the second on the benzene cycle. We then write down a family of elements in the kernel of $V \rightarrow I'$, which allows us to obtain an upper bound on the dimension of I' . Finally, we examine the irreducible \mathfrak{S}_{10} -modules (resp. \mathfrak{S}_{12} -modules) occurring in V and find that they all have larger dimension than the bound found for I' . This shows that I' is zero. The fact that we use the representation theory of \mathfrak{S}_{10} and \mathfrak{S}_{12} is the source of the denominator $12!$ in Theorem 8.1.


Throughout §8, L will always denote a set of cardinality 10 or 12. We handle the two cases simultaneously as much as possible. When we make different statements for the two cases, we give the 10-point statement first and the 12-point statement second.

8.1. The spaces $I'_{\mathcal{P}}$ and the A-, B- and C-relations. By a *bipartition* of L we mean an ordered pair $\mathcal{P} = (P, Q)$ with $P \sqcup Q = L$, where P has cardinality six, and Q has cardinality four if $|L| = 10$, and cardinality six if $|L| = 12$. By a graph of *type A* (resp. *type B*, resp. *type C*) with respect to a bipartition

\mathcal{P} we mean one of the form



when $|L| = 10$, or one of the form



when $|L| = 12$. In all cases the vertices of the top hexagon belong to P while the vertices in the bottom square or hexagon belong to Q .

We name some edges that we will refer to often. The most obvious distinguished edges are the two red edges joining the P -vertices and Q -vertices. We call them the *special red edges*. There is a unique blue (resp. green) edge in the top hexagon meeting one of the special red edges; we call it the (*top*) *special blue* (resp. *green*) *edge*. Finally, for type B graphs, we call the green edge in the bottom hexagon which meets the special red edges the *bottom special green edge*. We call the chain a - b - c - d , where \overline{ab} is the special blue edge, \overline{bc} is the special red edge to which it connects and \overline{cd} is the bottom special green edge, the *special chain* (only defined for type B graphs).

Let Γ be a graph of type A, B or C with respect to \mathcal{P} . Define a partition $\mathcal{U} = \{U_R, U_G, U_B\}$ of L as follows. In the type A case, U_R, U_G and U_B each contain two vertices in the top hexagon, and these vertices must be connected by a red, green and blue edge respectively. The bottom benzene cycle is in U_B . In the type B and C cases we take U_G to be the two vertices belonging to the top special green edge, we take U_B to be the two vertices belonging to the top special blue edge together with the bottom set (that is, the set Q), and we take U_R to be the leftover two vertices in the top hexagon (they form a red edge). In each case, (Γ, \mathcal{U}) is a generalized Segre cubic datum, and so has an associated generalized Segre cubic relation. In the type A case the partition \mathcal{U} was not uniquely determined, but its image in I'_L is independent of the choice of partition. We may thus in all cases speak unambiguously of the relation defined by Γ , denoted by $\text{Rel}(\Gamma)$. Note that if σ is a permutation of L then $\text{Rel}(\sigma\Gamma) = \epsilon(\sigma)\sigma(\text{Rel}(\Gamma))$.

We let $I'_{\mathcal{P}}$ be the subspace of I'_L generated by the relations attached to all type A graphs with respect to \mathcal{P} . Now, $I'_{\mathcal{P}}$ carries a natural representation of $\mathfrak{S}_P \times \mathfrak{S}_Q$. As such it is a quotient of the representation $I_P^{(3)} \otimes \text{gr}_{4,Q}^{(3)}$ (resp. $I_P^{(3)} \otimes \text{gr}_{6,Q}^{(3)}$). In particular, it is either zero or irreducible of dimension one (resp. five), as $I_P^{(3)}$ is irreducible of dimension one and $\text{gr}_{4,Q}^{(3)}$ (resp. $\text{gr}_{6,Q}^{(3)}$) is irreducible of dimension one (resp. five; see Propositions 6.13 and 6.14). The natural map (of \mathfrak{S}_L modules)

$$\bigoplus_{\mathcal{P}} I'_{\mathcal{P}} \rightarrow I'_L$$

is surjective since I_L is retrogenerated (Theorem 7.1) and there are no relations on 2 or 4 points.

Proposition 8.3. *Let \mathcal{P} be a bipartition of L . Then the type B relations with respect to \mathcal{P} belong to $I'_{\mathcal{P}}$ and span it. The same is true for the type C relations.*

We prove this proposition by showing that type A relations can be written in terms of type B relations, type B in terms of type C and finally type C in terms of type A, all with respect to the same \mathcal{P} . We accomplish this in a series of lemmas. After this proof we will have no need for the type C relations.


Lemma 8.4. *Any type A relation with respect to \mathcal{P} belongs to the space spanned by the type B relations with respect to \mathcal{P} .*

Proof. Given a type A graph, Plücker a red edge from the top hexagon and a red edge from the bottom benzene cycle. This expresses the type A relation as a sum of two type B relations. \square

Lemma 8.5. *Any type B relation with respect to \mathcal{P} belongs to the space spanned by the type C relations with respect to \mathcal{P} .*

Proof. To establish this lemma we just have to rewrite the blue/green graph on Q in terms of squares of matchings. This is exactly what we did in when proving Theorem 7.1(a) (see in particular Lemmas 7.7 and 7.8). The key point is that all square rotation relations which come up in our particular case have special paths of even length and are therefore quadratic.

Alternatively, one can argue as follows. The generalized Segre relation associated to graphs of type B is a genuine cubic relation, that is, the purple/black graph on the right side of the relation factors into a product of degree one graphs. For instance, in the 12 point case the factored relation may be written as



(the black graph on the right is colored blue while the purple graph has been factored into the green and red graphs). Now, the bottom half of each side is the same. One can therefore apply the same colored Plücker and quadratic relations on each side and rewrite the blue/green graph on Q as a sum of squares of matchings. This expresses the type B relation in terms of type C relations. \square

Lemma 8.6. *Any type C relation with respect to \mathcal{P} belongs to the space spanned by the type A relations with respect to \mathcal{P} .*

Proof. The proof of the 10 point case proceeds exactly as the proof of Lemma 7.11. Note that after applying the identity of Figure 9 to the redrawn graph, the first graph on the second line of the right side is a type A graph with respect to \mathcal{P} . All other graphs either come from 8 points or else are degenerate. To prove the 12 point case, we again use the identity of Figure 9, applying it to the bottom four red edges, that is, the two special red edges and the two red edges contained in the bottom hexagon. The first term in Figure 9 is a type A relation with respect to \mathcal{P} . All the other terms come from 8 or 10 points. \square

8.2. Relations among different $I'_{\mathcal{P}}$. In this section we demonstrate some linear dependencies between the various spaces $I'_{\mathcal{P}}$:

Proposition 8.7. *Let Γ be a type A graph with respect to a bipartition $\mathcal{P} = (P, Q)$, let a and b be two distinct elements of P and let c and d be two distinct elements of Q which are joined by a doubled edge. Then $\text{Rel}(\Gamma)$ is contained in $\sum I'_{\sigma\mathcal{P}}$, where the sum is taken over those permutations σ of $\{a, b, c, d\}$ for which $\sigma\mathcal{P} \neq \mathcal{P}$.*

We deduce this proposition from the following one:

Proposition 8.8. *Let Γ be a type B graph with respect to \mathcal{P} and let a - b - c - d be its special chain. Then*

$$\text{Rel}(\Gamma) = (ad) \text{Rel}(\Gamma) - (bc) \text{Rel}(\Gamma) + (ad)(bc) \text{Rel}(\Gamma) + \text{Rel}(\Delta)$$

where Δ is a type A graph with respect to $(bd)\mathcal{P}$. In particular, $\text{Rel}(\Gamma)$ belongs to $\sum I'_{\sigma\mathcal{P}}$, the sum taken over those permutations σ of $\{a, b, c, d\}$ for which $\sigma\mathcal{P} \neq \mathcal{P}$.

Proof of Proposition 8.7 given Proposition 8.8. Let Γ , a , b , c and d be given as in the statement of Proposition 8.7. The vertices c and d are connected by red and green edges. Now, using a quadratic relation we may recolor the P -part of Γ so that a and b are connected by a blue edge. Since \mathfrak{S}_P acts on the Segre cubic relation on P via the sign character, all blue edges in P are more or less the same so we may draw Γ as:



We now Plücker the red edge joining c and d and the unique red edge containing b . We obtain an expression $\Gamma = \Gamma_1 + \Gamma_2$ where each Γ_I is a type B graph with respect to \mathcal{P} . The special chain in Γ_1 is $a-b-c-d$ while in Γ_2 it is $a-b-d-c$. By Proposition 8.8, both $\text{Rel}(\Gamma_1)$ and $\text{Rel}(\Gamma_2)$ belong to $\sum I'_{\sigma, \mathcal{P}}$, where the sum is over those permutations σ of $\{a, b, c, d\}$ for which $\sigma\mathcal{P} \neq \mathcal{P}$. Thus $\text{Rel}(\Gamma)$ belongs to this space as well, which completes the proof. \square

We now begin proving Proposition 8.8. Consider a type B graph Γ . For convenience, we label it:



With this labeling, the special chain $a-b-c-d$ is 1-2-3-4. We let \mathcal{P} be the relevant bipartition.

By Plücker the edges $\overline{27}$ and $\overline{34}$ we obtain $\Gamma = -\Gamma_2 - \Gamma_3$ where Γ_2 and Γ_3 are:



We now Plücker the edges $\overline{12}$ and $\overline{39}$ in Γ_3 to obtain a relation $\Gamma_3 = -\Gamma_4 - \Gamma_5$ where Γ_4 and Γ_5 are given by:




Now, we have $\Gamma_4 = (23)\Gamma$. Thus (as generalized Segre data)

$$(8.1) \quad \Gamma - (23)\Gamma + \Gamma_2 = \Gamma_5.$$


Lemma 8.9. *We have $(14)\text{Rel}(\Gamma_2) = \text{Rel}(\Gamma_2) + \text{Rel}(\Delta)$ where Δ is a type A graph with respect to $(24)\mathcal{P}$.*

Proof. The graph Γ_2 has a benzene chain with three (resp. four) blue edges, beginning at vertex 1 and ending at vertex 4. The picture is thus:



Note that $U_R = \{7, 8\}$ and $U_G = \{5, 6\}$ while U_B consists of the remaining vertices.


We now Plücker the edges $\overline{81}$ and $\overline{74}$. In the first term, 8 connects to 4 and 7 to 1. The second term is a degenerate Segre datum (since $\overline{81}$ and $\overline{74}$ are the special green edges in U_B). It is thus quadratic and can be ignored. We next Plücker the edges $\overline{61}$ and $\overline{54}$, the results being similar. We have thus shown that $\text{Rel}(\Gamma_2)$ defines the same element of I'_L as the relation associated to following graph (since we used two Plücker relations no sign is introduced):



We now Plücker $\overline{12}$ and $\overline{410}$. The first term is $(14)\Gamma_2$. The second term is a type A graph Δ with respect to $(24)\mathcal{P}$. We thus find $\text{Rel}(\Gamma_2) = -\text{Rel}((14)\Gamma_2) - \text{Rel}(\Delta)$, which establishes the proposition. \square

Lemma 8.10. *We have (14) $\text{Rel}(\Gamma_5) = \text{Rel}(\Gamma_5)$.*

Proof. We redraw Γ_5 in a more convenient way.



As before, $U_R = \{7, 8\}$ and $U_G = \{5, 6\}$ while U_B consists of the remaining vertices.

We now Plücker the edges $\overline{61}$ and $\overline{54}$. In the first term, 5 connects to 1 and 6 to 4. The second term is a degenerate Segre datum and can be ignored. We next Plücker $\overline{24}$ with $\overline{81}$. In the first term, 8 connects to 4 and 2 to 1. The second term is degenerate and can be ignored. We have thus shown that $\text{Rel}(\Gamma_5)$ defines the same element of I'_L as the relation associated to the following graph:



We now Plücker the edges $\overline{19}$ and $\overline{410}$. The first term is just (14) Γ_5 . The second term is retrogenerated from 8 (resp. 10) points and therefore does not contribute in I'_L . We thus have $\text{Rel}(\Gamma_5) = -\text{Rel}((14)\Gamma_5)$, which establishes the proposition. \square

We can now complete the proof of Proposition 8.8.

Proof of Proposition 8.8. By Lemma 8.10 and (8.1), we see that $\text{Rel}(\Gamma) - \text{Rel}((23)\Gamma) + \text{Rel}(\Gamma_2)$ is invariant under (14), that is,

$$\text{Rel}(\Gamma) - \text{Rel}((23)\Gamma) + \text{Rel}(\Gamma_2) = (14) \text{Rel}(\Gamma) - (14) \text{Rel}((23)\Gamma) + (14) \text{Rel}(\Gamma_2).$$

We now apply Lemma 8.9 and write (14) $\text{Rel}(\Gamma_2) = \text{Rel}(\Gamma_2) + \text{Rel}(\Delta)$ where Δ is a type A graph with respect to (24) \mathcal{P} . The Γ_2 terms on each side of the equation cancel, and we are left with

$$\text{Rel}(\Gamma) = (14) \text{Rel}(\Gamma) - (23) \text{Rel}(\Gamma) + (14)(23) \text{Rel}(\Gamma) + \text{Rel}(\Delta).$$

This completes the proof. \square

8.3. Proof of Theorem 8.2. We now complete the proof of Theorem 8.2. As we have explained, to do this it suffices to show that $I'_L = 0$. To do this we first use Proposition 8.7 to obtain an upper bound for the dimension of I'_L , and then we use representation theory to show that this upper bound forces I'_L to be zero.

Fix an order on L . For $n = 10$ (resp. $n = 12$), we say that a bipartition $\mathcal{P} = (P, Q)$ of L is *good* if there is at most one (resp. if there are at most two) element(s) of Q which are larger than the second smallest element of P .

Proposition 8.11. *The natural map $\bigoplus I'_{\mathcal{P}} \rightarrow I'$ is surjective, where the sum is taken over all good bipartitions \mathcal{P} .*

Proof. We handle the two cases $n = 10$ and 12 separately, for the sake of clarity. We begin with the 10 point case. Let \mathcal{P} be a bipartition which is not good. Let $a < b$ be the smallest two elements of P and let $c < d$ be the largest two elements of Q . Since \mathcal{P} is not good, we have $a < b < c < d$. Now let Γ be a type A graph with respect to \mathcal{P} . Since $I_P^{(3)} \otimes \text{gr}_{4,Q}^{(3)}$ is one-dimensional, if Γ' is any other type A graph with respect to \mathcal{P} then we have $\text{Rel}(\Gamma) = \text{Rel}(\Gamma')$. Thus we may assume that c and d are connected by a doubled edge in Γ . We may then apply Proposition 8.7 to conclude that Γ belongs to $\sum I'_{\sigma\mathcal{P}}$, the sum taken over those permutations σ of $\{a, b, c, d\}$ for which $\sigma\mathcal{P} \neq \mathcal{P}$. Each of these $\sigma\mathcal{P}$ is closer to being good than \mathcal{P} (measured by how many elements of Q are larger than the second smallest element of P , for instance). By induction, we deduce the proposition.

We now handle the case where L has cardinality 12. Let \mathcal{P} be a bipartition which is not good. Let $a < b$ be the smallest two elements of P and let $c < d < e$ be the largest three elements of Q . Since \mathcal{P} is not good, we have $a < b < c < d < e$. Now let Γ be a type A graph with respect to \mathcal{P} . By Proposition 6.14(d) we can rewrite $\text{Rel}(\Gamma)$ in terms of $\text{Rel}(\Gamma_i)$, where each Γ_i is a type A graph with respect to \mathcal{P} in which either \overline{cd} or \overline{ce} appears as a doubled edge. We may as well then just assume that Γ has \overline{cd} as a doubled edge (the \overline{ce} argument is the same). We now apply Proposition 8.7 and find that $\text{Rel}(\Gamma)$ belongs to $\sum I'_{\sigma\mathcal{P}}$ as σ varies

over those permutations of $\{a, b, c, d\}$ for which $\sigma\mathcal{P} \neq \mathcal{P}$. As before, each $\sigma\mathcal{P}$ is closer to being good than \mathcal{P} and so we deduce the proposition by induction. \square

Proposition 8.12. *For $n = 10$ (resp. $n = 12$) there are 25 (resp. 112) good bipartitions of L .*

Proof. We again consider the two cases separately, and begin with the 10 point case. Identify L with $\{1, 2, \dots, 10\}$. Consider bipartitions of L for which the second smallest element of P is x . Of course, we must have $x \leq 6$. On the other hand, if $x \leq 4$ then at least two elements of Q will be larger than x and so the bipartition will not be good. Thus we must have $x = 5$ or $x = 6$.

- $x = 6$. In this case, P must contain each of 7, 8, 9 and 10. In addition, P contains one number less than 6. There are five such choices, each of which gives a good bipartition.
- $x = 5$. In this case, P must contain each of 6, 7, 8, 9 and 10 except for one (it cannot contain them all for then x would not be second smallest). There are five ways to choose what to omit. In addition, P must contain one number less than 5. There are four such choices. Thus, all told, there are 20 good bipartitions.

Thus in total we have $5 + 20 = 25$ good bipartitions.

We now handle the case where L has cardinality 12. Again, we identify L with $\{1, \dots, 12\}$ and consider bipartitions for which the second smallest element of P is x . We must have $x \leq 8$. Now, if $x \leq 5$ then Q necessarily has three elements larger than x and so the bipartition is not good. Thus we only need to consider the cases $x = 6, 7, 8$.

- $x = 8$. In this case, P must contain each of 9, 10, 11 and 12. In addition, P contains one number less than 8. There are seven choices for such a number, and all give good bipartitions.
- $x = 7$. In this case, P must contain each of 8, 9, 10, 11 and 12 except for one (it cannot contain them all for then x would not be second smallest). There are five ways to choose what to omit. Furthermore, P must contain one element smaller than 7. There are six such choices. Thus there are 30 good bipartitions in this case.
- $x = 6$. In this case, P must contain each of 7, 8, 9, 10, 11 and 12 except for two. There are 15 ways to omit two of these numbers. Furthermore, P must contain one element smaller than 6. There are five such choices. Thus there are 75 good bipartitions in this case.

Thus in total we have $7 + 30 + 75 = 112$ good bipartitions. \square

Corollary 8.13. *For $n = 10$ (resp. $n = 12$) we have $\dim I'_L \leq 25$ (resp. ≤ 560).*

Proof. We know that I'_L admits a surjection from $\bigoplus I'_{\mathcal{P}}$, the sum taken over the good bipartitions \mathcal{P} . When L has cardinality 10, there are 25 of these bipartitions and each $I'_{\mathcal{P}}$ has dimension at most one. When L has cardinality 12, there are 112 of these bipartitions and each $I'_{\mathcal{P}}$ has dimension at most five. This gives the result. \square

Proposition 8.14. *For $n = 10$ (resp. $n = 12$) the space $\bigoplus I'_{\mathcal{P}}$ (summed over all bipartitions \mathcal{P}) is either zero a direct sum of two irreducible \mathfrak{S}_L -representations of dimensions 84 and 126 (resp. three irreducible representations of dimensions 616, 1925 and 2079).*

Proof. The space $\bigoplus I'_{\mathcal{P}}$ is isomorphic as an \mathfrak{S}_L -module to $\text{Ind}_{\mathfrak{S}_P \times \mathfrak{S}_Q}^{\mathfrak{S}_L} I'_{\mathcal{P}}$ for any fixed bipartition $\mathcal{P} = (P, Q)$. We now handle the two cases separately.

First suppose $n = 10$. Then $I'_{\mathcal{P}}$ is either zero or isomorphic as an $(\mathfrak{S}_P \times \mathfrak{S}_Q)$ -module to $I_P^{(3)} \otimes \text{gr}_{4,Q}^{(3)}$. As an \mathfrak{S}_P -module $I_P^{(3)}$ corresponds to the partition $1 + 1 + 1 + 1 + 1 + 1$ (the alternating representation) while as an \mathfrak{S}_Q -module $\text{gr}_{4,Q}^{(3)}$ corresponds to the partition 4 (the trivial representation). We now use the Littlewood-Richardson rule to compute $\text{Ind}_{\mathfrak{S}_P \times \mathfrak{S}_Q}^{\mathfrak{S}_L} I'_{\mathcal{P}}$. Excluding the case where $I'_{\mathcal{P}} = 0$, we find that the induction decomposes into a direct sum of two irreducibles, corresponding to the partitions

$$5 + 1 + 1 + 1 + 1 + 1, \quad 4 + 1 + 1 + 1 + 1 + 1 + 1.$$

The hook length formula shows that these irreducibles have dimensions 126 and 84, respectively.

Now say $n = 12$. Again, $I'_{\mathcal{P}}$ is either zero or isomorphic as an $(\mathfrak{S}_P \times \mathfrak{S}_Q)$ -module to $I_P^{(3)} \otimes \text{gr}_{6,Q}^{(3)}$. The representation $I_P^{(3)}$ corresponds to the partition $1 + 1 + 1 + 1 + 1 + 1$ while $\text{gr}_{6,Q}^{(3)}$ corresponds to $3 + 3$. The

Littlewood-Richardson rule shows that the induction decomposes into three irreducibles, corresponding to the partitions

$$4 + 3 + 1 + 1 + 1 + 1 + 1, \quad 4 + 4 + 1 + 1 + 1 + 1, \quad 3 + 3 + 1 + 1 + 1 + 1 + 1 + 1.$$

These irreducibles have dimensions 2079, 1925 and 616, respectively. \square

We now complete the proof of Theorem 8.2.

Proof of Theorem 8.2. We must show $I'_L = 0$. If $n = 10$, we know that on one hand, I'_L is at most 25-dimensional, while on the other, it is a direct sum of at most two irreducibles of dimensions 84 and 126. It follows that I'_L must be zero. Similarly, when $n = 12$, we know that, on one hand, I'_L is at most 560-dimensional, while on the other, it is a direct sum of at most three irreducibles of dimensions 616, 1925 and 2079. Again, we conclude $I'_L = 0$. \square

9. THE QUADRATICS ARE GENERATED BY THE SIMPLEST BINOMIALS

In §9, we prove the following, concluding the proof of Theorem 1.2:

Theorem 9.1. *If L has even cardinality n , then the simplest binomial relations span $I_L^{(2)}$ over $\mathbb{Z}[\frac{1}{n!}]$.*

The proof uses the *simple binomial relations*. We show that the spans of the simple binomial relations and simplest binomial relations coincide. We then use the \mathfrak{S}_L -module structure of $I_L^{(2)}$ to show that simple binomial relations generate it. It suffices to prove the result over any field of characteristic 0 or greater than n , so we do this.

9.1. Simple binomial relations. Let L be an even set. A *binomial quadratic datum* is a pair $D = (\Gamma, U)$ where Γ is an undirected regular 2-colored graph on L and U is a subset of L such that all edges of Γ are contained in U or $L \setminus U$. Define Γ' to be the graph obtained by inverting the colors of the edges of Γ contained in U , and set

$$\text{Rel}(D) = Y_\Gamma - Y_{\Gamma'},$$

which is clearly a relation.

Let D be a binomial quadratic datum. Suppose that a pair of edges e and e' of Γ have the same color and both lie in U or $L \setminus U$. Let Γ' and Γ'' be the other graphs occurring in the colored Plücker relation on e and e' . Then $D' = (\Gamma', U)$ and $D'' = (\Gamma'', U)$ are binomial quadratic data. We define the *space of binomial quadratic data* to be the \mathbb{Z} -span of the binomial quadratic data modulo the (Plücker) relations $D + D' + D'' = 0$. The association $D \mapsto \text{Rel}(D)$ descends to a linear map

$$\text{Rel} : \{\text{the space of binomial quadratic data}\} \rightarrow I_L^{(2)}.$$

We call a binomial quadratic datum *simple* if U has cardinality four. We call the resulting relations *simple binomial relations*. We say that a binomial quadratic datum is *simplest* if it is simple and in addition Γ is made up of 2-cycles and two 4-cycles. The associated relations are the *simplest binomial relations* defined in the introduction; (1.4) is an example. Note that if $D = (U, \Gamma)$ is a binomial quadratic datum and Γ is a union of 2-cycles and zero or one 4-cycle then $\text{Rel}(D) = 0$. Although there are more simple binomial relations than simplest binomial relations, they span the same space:

Proposition 9.2. *Every simple binomial relation is a linear combination of simplest binomial relations.*

Thus to prove Theorem 9.1 it suffices to show that the simple binomial relations span $I_L^{(2)}$. (We only need characteristic not 2 to prove this.)

Proof. Let $D = (\Gamma, U)$ be a simple binomial quadratic datum. By Proposition 6.1 we can use colored Plücker relations in $L \setminus U$ to write $\Gamma|_{L \setminus U}$ as a sum of graphs, each of which is a union of 2-cycles and at most one 4-cycle. This expresses $\text{Rel}(D)$ in terms of simplest binomial relations. \square

Let L be an even set and let U be a subset of cardinality four. Let $L' = L \setminus U$. Define a map

$$\iota : \bigwedge^2 V_{L'} \otimes \bigwedge^2 V_U \rightarrow I_L^{(2)} \quad \text{by} \quad (Y_\Gamma \wedge Y_{\Gamma'}) \otimes (Y_\Delta \wedge Y_{\Delta'}) \mapsto Y_{\Gamma\Delta} Y_{\Gamma'\Delta'} - Y_{\Gamma'\Delta} Y_{\Gamma\Delta'}.$$

This is the simple binomial relation associated to the the simple binomial data (Φ, U) where Φ is the 2-colored graph $\Gamma'\Delta\Delta'$ in which $\Gamma\Delta$ has one color and $\Gamma'\Delta'$ the other. Note that ι is $(\mathfrak{S}_{L'} \times \mathfrak{S}_U)$ -equivariant.

Lemma 9.3. *The map ι is injective.*

Proof. Let $Y_\Delta, Y_{\Delta'}$ be a basis of V_U . It suffices to show that the map $\bigwedge^2 V_{L'} \rightarrow \text{Sym}^2(R_L^{(1)})$ given by

$$Y_\Gamma \wedge Y_{\Gamma'} \mapsto Y_{\Gamma\Delta} Y_{\Gamma'\Delta'} - Y_{\Gamma'\Delta} Y_{\Gamma\Delta'}$$

is injective. Let f (resp. f') be the map $V_{L'} \rightarrow R_L^{(1)}$ given by $Y_\Gamma \mapsto Y_{\Gamma\Delta}$ (resp. $Y_{\Gamma'} \mapsto Y_{\Gamma'\Delta'}$). Both f and f' are injective. Furthermore, the images of f and f' are linearly disjoint. To see this, use the fact that planar graphs form a basis for $R_L^{(1)}$ (Theorem 2.2): Put L on a circle so that the vertices in U are consecutive. If Γ and Γ' are planar graphs on L' and Δ and Δ' are distinct planar graphs on U then $\Gamma\Delta$ and $\Gamma'\Delta'$ are distinct planar graphs on L . The lemma now follows from the following lemma in linear algebra. \square

Lemma 9.4. *Let $f, g : V \rightarrow W$ be linear maps of vector spaces which are injective and have linearly disjoint images. Then the map*

$$\bigwedge^2 V \rightarrow \text{Sym}^2(W), \quad v \wedge w \mapsto f(v)g(w) - f(w)g(v)$$

is also injective.

Proof. Let v_i be a basis for V and assume that $v = \sum \alpha_{ij} v_i \wedge v_j$ belongs to the kernel, so

$$0 = \sum \alpha_{ij} (f(v_i)g(v_j) - f(v_j)g(v_i)) = \sum (\alpha_{ij} - \alpha_{ji}) f(v_i)g(v_j).$$

The vectors $f(v_i)g(v_j)$ are linearly independent in $\text{Sym}^2(W)$, so $\alpha_{ij} = \alpha_{ji}$ and hence $v = 0$. \square

9.2. Completion of the proof of Theorem 9.1. Let X denote the set of partitions of n into exactly four even parts and let X^- be the set of partitions of $n - 4$ into exactly four odd parts. Thus X has the partitions occurring in I_L and X^- has those occurring in $\bigwedge^2 V_{L'}$ (Proposition 6.5). For a partition $\lambda \in X$ we define a corresponding partition $\lambda^- \in X^-$ by removing one box from each row in the Young diagram: $(a, b, c, d) \mapsto (a - 1, b - 1, c - 1, d - 1)$. This bijection of sets $X \rightarrow X^-$ has a representation-theoretic characterization:

Lemma 9.5. *Let $\lambda^- \in X^-$ and $\mu \in X$ be given. (Recall from §6.4 that M_λ is the irreducible representation of \mathfrak{S}_L corresponding to partition λ .) Then*

$$\dim \text{Hom}_{\mathfrak{S}_{L'} \times \mathfrak{S}_U}(M_{\lambda^-} \otimes \epsilon, M_\mu) = \begin{cases} 1 & \text{if } \lambda = \mu \\ 0 & \text{otherwise} \end{cases}$$

Here ϵ denotes the sign representation of \mathfrak{S}_U .

Proof. By Frobenius reciprocity, the dimension of the Hom space is equal to the multiplicity of M_μ occurring in the induction $\text{Ind}_{\mathfrak{S}_{L'} \times \mathfrak{S}_U}^{\mathfrak{S}_L}(M_{\lambda^-} \otimes \epsilon)$. This can be computed using the Littlewood-Richardson rule, which is simple in this case because ϵ is just the alternating representation (the Young diagram is a single column of four boxes). The key point is that the only way to add on four boxes to λ and end up with something in X is to put one box at the end of each row. \square

Corollary 9.6. *The map ι carries $(\bigwedge^2 V_{L'})[\lambda^-] \otimes \bigwedge^2 V_U$ into $I_L^{(2)}[\lambda]$ for any $\lambda \in X$, where as usual $V[\lambda]$ is the λ -isotypic part of a representation V of the symmetric group.*

We now can prove Theorem 9.1.

Proof of Theorem 9.1. According to the above corollary, ι induces a map

$$\left(\bigwedge^2 V_{L'}\right)[\lambda^-] \otimes \bigwedge^2 V_U \rightarrow I_L^{(2)}[\lambda]$$

for any $\lambda \in X$. This map is injective (Lemma 9.3), and $(\bigwedge^2 V_{L'})[\lambda^-]$ is non-zero (Proposition 6.5). Thus the image of ι has non-zero projection to each $I_L^{(2)}[\lambda]$. Since $I_L^{(2)}$ is multiplicity free (Proposition 6.5), the image of ι generates $I_L^{(2)}$ as an \mathfrak{S}_L -module. Hence the simple binomial relations span $I_L^{(2)}$. Since every simple binomial relation is a linear combination of simplest binomial relations (Proposition 9.2), we deduce Theorem 9.1. \square

REFERENCES

- [FS] E. Freitag and R. Salvati Manni, *The modular variety of hyperelliptic curves of genus three*, arXiv:0710.5920v1.
- [HMSV1] B. Howard, J. Millson, A. Snowden and R. Vakil, *The equations for the moduli space of n points on the line*, Duke Math. J., **146** No. 2 (2009), 175-226.
- [HMSV2] B. Howard, J. Millson, A. Snowden, and R. Vakil, *The relations among invariants of points on the projective line*, arXiv:0906.2437, C. R. Math. Acad. Sci. Paris, to appear.
- [HMSV3] B. Howard, J. Millson, A. Snowden, and R. Vakil, *The ring of projective invariants of eight points on the line via representation theory*, arXiv:0809.1233.
- [HMSV4] B. Howard, J. Millson, A. Snowden and R. Vakil, *The ideal of relations for the ring of invariants of n points on the line: integral results*, in preparation.
- [HMSV5] B. Howard, J. Millson, A. Snowden and R. Vakil, <http://www.math.princeton.edu/~asnowden/hmsv2/>
- [Ke] A. Kempe, *On regular difference terms*, Proc. London Math. Soc. **25** (1894), 343-350.
- [Ko] K. Koike, *The projective embedding of the configuration space $X(2, 8)$* , Technical Reports of Mathematical Sciences, Chiba University, **16** (2000).
- [Ku] S. Kumar, *A remark on universal connections*, Math. Ann. **260** (1982), 453-462.
- [M] C. Manon, *Presentations of semigroup algebras of weighted trees*, arXiv:0808.1320.
- [SS] D. Speyer and B. Sturmfels, *The tropical Grassmannian*, Adv. Geom. **4**, no. 3, (2004), 389-411.
- [S] B. Sturmfels, *Gröbner Bases and Convex Polytopes*, University Lecture Series **8**, Amer. Math. Soc., Providence, RI, 1996.
- [W] H. Weyl, *The Classical Groups: Their Invariants and Representations*, Princeton U. P., Princeton NJ, 1997.