Canad. J. Math. Vol. **70** (6), 2018 pp. 1416–1435
http://dx.doi.org/10.4153/CJM-2018-006-5
© Canadian Mathematical Society 2018

A Special Case of Completion Invariance for the c_2 Invariant of a Graph

Karen Yeats

Abstract. The c_2 invariant is an arithmetic graph invariant defined by Schnetz. It is useful for understanding Feynman periods. Brown and Schnetz conjectured that the c_2 invariant has a particular symmetry known as completion invariance. This paper will prove completion invariance of the c_2 invariant in the case where we are over the field with 2 elements and the completed graph has an odd number of vertices. The methods involve enumerating certain edge bipartitions of graphs; two different constructions are needed.

1 Introduction

The c_2 invariant is an arithmetic graph invariant defined by Schnetz [24] (see Definition 1.1) that is useful for understanding Feynman periods. Feynman periods are a class of integrals defined from graphs that are simpler than but closely related to Feynman integrals. Based on this connection and on computational evidence, there are certain symmetries that the c_2 invariant is believed to have. A key such symmetry, known as completion invariance and defined below, was first conjectured by Brown and Schnetz in 2010 [8] and has turned out to be quite difficult to prove. The main result of this paper, Theorem 1.2, is the completion invariance of the c_2 invariant in the case where $p = 2$ and the completed graph has an odd number of vertices.

Throughout, K will be a connected 4-regular simple graph. (For unexplained graph theory language or notation, see [16].) The result of removing any vertex of K is called a *decompletion* of K. Different decompletions will typically be non-isomorphic, and K can be uniquely reconstructed from any of its decompletions. We will say that K is the *completion* of any of its decompletions. See Figure 1 for an example.

Suppose G is a decompletion of K. Then we can view G as a Feynman diagram in ϕ^4 theory. For those not familiar with Feynman diagrams and quantum field theory, briefly: The edges of the graphs represent particles; the vertices particle interactions. This particular quantum field theory has only a quartic interaction (that is the 4 in ϕ^4), and so all vertices must have degree 4. The vertices of G that are no longer 4-regular due to decompletion are taken to have additional external edges attached that represent particles entering or exiting the system. Completing G then means attaching all the external edges to a new vertex that one can think of as being "at infinity", hence corresponding to completion in the geometric sense. This connection to geometric completion is easily made precise at the level of the Feynman period; see [23].

Received by the editors July 3, 2017; revised January 25, 2018.
Published electronically April 27, 2018.
The author is supported by an Alexander von Humboldt fellowship and an NSERC Discovery grant.
AMS subject classification: **05C31**, 05C30, 81T18.
Keywords: c_2 invariant, Feynman graph, edge partition, spanning forest, completion conjecture.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

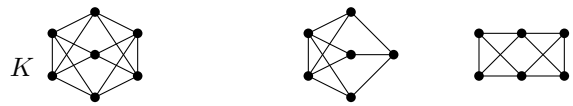

Figure 1: A graph K with two non-isomorphic decompletions

The Feynman integral of a Feynman diagram is what is actually needed for physics, because it computes the contribution of that Feynman diagram to whatever physical process one is interested in. For more on perturbative quantum field theory, see [19].

Given any graph G (but most interestingly for the present purposes G may be some decompletion of K), assign a variable a_e to each edge $e \in E(G)$ and define the (dual) *Kirchhoff polynomial* or first Symanzik polynomial to be

$$\Psi_G = \sum_T \prod_{e \notin T} a_e,$$

where the sum runs over all spanning trees of G. For example, the Kirchhoff polynomial of a 3-cycle with edge variables a_1, a_2, a_3 is $a_1 + a_2 + a_3$, since removing any one edge of a cycle gives a spanning tree of the cycle.

Next define the *Feynman period* to be

$$\int_{a_e \geq 0} \frac{\prod da_e}{\Psi_G^2}\bigg|_{a_1=1}.$$

This is an affine form of the integral; there is also a projective form; see [23]. This integral converges provided G is at least internally 6-edge connected.[1] This corresponds to the Feynman diagram having no subdivergences, because edge cuts with four or fewer edges, other than ones giving isolated vertices, represent sub-processes that yield divergent integrals. There has been a lot of interest in the Feynman period, because it is a sensible algebro-geometric, or even motivic, object [3,4,6,7,21,23], but it is also a key piece of the Feynman integral. It is a sort of coefficient of divergence for the Feynman integral and has the benefit of not depending on the many parameters and choices which the full Feynman integral depends on, so it is mathematically much tidier. Nonetheless, it still captures some of the richness of Feynman integrals. This is best illustrated by the number-theoretic richness of the numbers that Feynman periods give; see [3–5,22,23].

Given these connections to many deep and difficult things, it should not be surprising that Feynman periods are still difficult to understand and compute, and so in [24] Schnetz introduced the c_2 invariant in order to better understand the Feynman period.

Definition 1.1 Let p be a prime, let $\mathbb{F}_p$ be the finite field with p elements, let G be a connected graph with at least 3 vertices, and let $[\Psi_G]_p$ be the number of $\mathbb{F}_p$-rational

[1]That is, any way of removing fewer than 6 edges either leaves the graph connected or breaks the graph into two components one of which is a single vertex.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

points in the affine variety $\{\Psi_G = 0\} \subseteq \mathbb{F}_p^{|E(G)|}$. Then the c_2-*invariant* of G at p is

$$c_2^{(p)}(G) = \frac{[\Psi_G]_p}{p^2} \mod p.$$

That this is well defined is proved in [24]. The same definition can be made for prime powers, but we will stick to primes p. The c_2 invariant has interesting properties [8,10,17,18], yields interesting sequences in p such as coefficient sequences of modular forms [9, 20], and predicts properties of the Feynman period. A simple and striking example of the last of these is that $c_2^{(p)}(G) = 0$ for all p corresponds to when the Feynman period apparently has a drop in transcendental weight relative to the size of the graph.[2] For example, the two decompleted graphs in Figure 1 both have $c_2^{(p)} = 0$ for all p and both have period $36\zeta(3)^2$, which has weight $3 + 3 = 6$, while some other graphs with the same number of edges have weight 7; see [23].

The c_2 invariant should have something to do with the Feynman period, because both counting points and taking period integrals are controlled by the geometry of the variety of Ψ_G. If the Feynman periods are nice, then the point counts as a function of p should be nice and vice versa. More specifically, inspired by known Feynman periods at the time being multiple zeta values,[3] Kontsevich informally conjectured that $[\Psi_G]_p$ should be a polynomial in p. This turned out to be very false [2]. The c_2 invariant is one measure of whether or how badly Kontsevich's conjecture fails for a given graph; if it holds for that graph, then the c_2 invariant is the quadratic coefficient of the polynomial, thus explaining the 2 in c_2.

The c_2 invariant has or is believed to have all of the symmetries of the Feynman period. In particular, it is known that the Feynman periods of two graphs with the same completion are the same; see [23]. The c_2 invariant was also conjectured by Brown and Schnetz in 2010 [8, Conjecture 4] to have this completion symmetry. This is arguably the most interesting open problem on the arithmetic of Feynman periods. Very little progress has been made. We do know that if G has two triangles sharing an edge, then the question of completion invariance for $c_2^{(p)}(G)$ can be reduced to completion invariance of a particular smaller graph. This is known as double-triangle reduction; see [8, Corollary 34].

The main result of this paper puts the first crack into the conjecture itself, proving it in the special case when $p = 2$ and the completed graph has an odd number of vertices.

Theorem 1.2 *Let K be a connected 4-regular graph with an odd number of vertices. Let v and w be vertices of K. Then $c_2^{(2)}(K - v) = c_2^{(2)}(K - w)$.*

The basic approach takes a graph-theoretic perspective on the c_2 invariant. The same approach also underlies the c_2 results for families of graphs from [13, 26], and

[2]Given that proving transcendentality for single zeta values, much less multiple zeta values or other more exotic numbers appearing in Feynman periods, is completely out of reach, one must recast this more formally in order to make a precise statement or prove a precise result along these lines; see [8,24].

[3]This is no longer the case unless there are some outrageous identities in play; as with transcendentality questions for the multiple zeta values themselves, proving an absence of relations is almost impossibly hard even when we have solid theoretical reasons to think there should be none.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

the reader may like to look there for further examples. The strength of the approach is that it is sufficiently different from a more algebraic or geometric approach to possibly make progress where other methods stalled. The weakness of the approach is that extending beyond $p = 2$ will be tricky, perhaps impossible. See Section 6 for further comments. The vertex parity restriction is both more mysterious and perhaps more tractable; again, see Section 6. A different graph-theoretic perspective on the c_2 invariant is given by Crump in [14].

The structure of the paper is as follows. Section 2 defines the different graph polynomials that will be needed and collects together some important lemmas from other sources. Section 3 then uses these lemmas to reduce the problem of proving the main theorem to a problem of determining the parity of the number of certain edge bipartitions. The real work of proving Theorem 1.2 then comes in Sections 4 and 5. The set of these edge bipartitions is divided into five pieces. Two of the pieces are proved to have even size in Section 4 by finding fixed-point free involutions. The other three pieces are proved to have even size in Section 5 by giving a cycle swapping rule to transform edge bipartitions. Finally, Section 6 concludes with some comments about the result, the proof, and the way forward.

2 Background

The first step is to define some additional graph polynomials that will be needed in order to prove the main result.

Let G be a graph. Choose an arbitrary order for the edges and the vertices of G and choose an arbitrary orientation for the edges of G. Let E be the signed incidence matrix of G (with the vertices corresponding to rows and the edges corresponding to columns) with one arbitrary row removed and let Λ be the diagonal matrix with the edge variables of G on the diagonal. Let $M = \left[\begin{smallmatrix} \Lambda & E^t \\ -E & 0 \end{smallmatrix}\right]$. Then $\det M = \Psi_G$. This can be proved by expanding the determinant (see [7, Proposition 21]), or by using the Schur complement and the Cauchy–Binet formula, see [25]. In both cases, it comes down to the fact that the square full rank minors of E are ± 1 for columns corresponding to the edges of a spanning tree of G, which is the essence of the matrix tree theorem. In this and other ways, the matrix M behaves much like the Laplacian matrix of a graph (with variables and one matching row and column removed), but the pieces that make it up are separated out, so call M the *expanded Laplacian* of G.

Assume that we have made a fixed choice of orders and orientation so as to define a fixed M in all that follows. Thus, it will not matter whether we talk of edges or edge indices. In particular, we will use G/e and $G\backslash e$ for the contraction and deletion respectively in the graph G with the same meaning whether e is an edge or an edge index.

If I and J are sets of edge indices, then let $M(I, J)$ be the matrix M with the rows indexed by elements of I removed and the columns indexed by elements of J removed. Brown [7] defined the following Dodgson polynomials.

Definition 2.1 Let I, J, and K be sets of edge indices with $|I| = |J|$. Define

$$\Psi^{I,J}_{G,K} = \det M(I, J)|_{a_e=0 \text{ for } e\in K}.$$

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

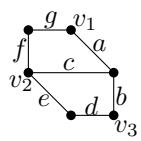

Figure 2: A graph for a spanning forest example.

We can leave out the G subscript when the graph is clear and leave out the K when $K = \varnothing$.

Dodgson polynomials have nice contraction-deletion properties and many interesting relations; see [7]. Making different choices in the construction of M may change the overall sign of a Dodgson polynomial, but since we will be concerned with counting zeros of these polynomials, the overall sign is of no concern.

Dodgson polynomials can also be expressed in terms of spanning forests; this is a consequence of the all minors matrix tree theorem [12]. The following spanning forest polynomials are convenient for this purpose.

Definition 2.2 Let P be a set partition of a subset of the vertices of G. Define

$$\Phi_G^P = \sum_F \prod_{e \notin F} a_e,$$

where the sum runs over spanning forests F of G with a bijection between the trees of F and the parts of P where each vertex in a part lies in its corresponding tree.

Note that trees consisting of isolated vertices are permitted. In most of the argument, we will be working with spanning forest polynomials where P has exactly two parts. The corresponding spanning forests thus have exactly two parts and will be known as 2-*forests*. For example, if G is as in Figure 2, then

$$\Phi_G^{\{v_1,v_2\},\{v_3\}} = (e+d)(ca+cb+ab+fb+gb).$$

The precise relationship between Dodgson polynomials and spanning forest polynomials is given in the following proposition.

Proposition 2.3 ([11, Proposition 12]) *Let I, J, and K be sets of edge indices of G with $|I| = |J|$. Then*

$$\Psi_{G,K}^{I,J} = \sum \pm \Phi_{G\backslash(I\cup J\cup K)}^P,$$

where the sum runs over all set partitions P of the end points of edges of $(I\cup J\cup K)\backslash(I\cap J)$ such that all the forests corresponding to P become spanning trees in both $G\backslash I/(J\cup K)$ and $G\backslash J/(I\cup K)$.

The signs appearing in the proposition can be determined (see [11]); however, they are of no concern for the present, since we will be working modulo 2. Note that if an edge index is in both I and J, then it is deleted in both $G\backslash I$ and in $G\backslash J$ and so cannot then be contracted; in the contraction, simply ignore edge indices whose edges are no longer there.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

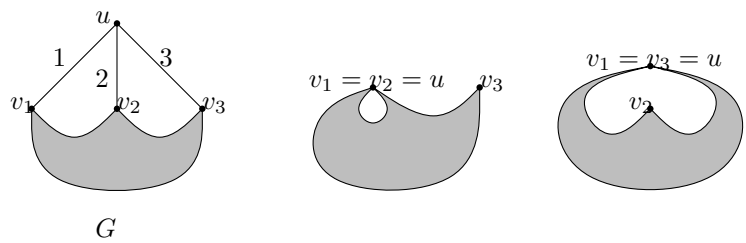

Figure 3: An example graph for some Dodgson polynomial to spanning forest polynomial computations and two of its minors.

The instances of this proposition that will be needed below also serve as a good example of these objects and how to manipulate them.

Example 2.4 Suppose we have a graph G with a 3-valent vertex u and let the edges incident to u be 1, 2, and 3, as in the first graph in Figure 3.

Consider the Dodgson polynomial $\Psi_G^{12,13}$. To find its expression in terms of spanning forest polynomials, we need to look for sets of edges of $G\backslash 123$ that give a spanning tree in $G\backslash 12/3$ and in $G\backslash 13/2$. These smaller graphs are each isomorphic to $G-u$, and so we obtain

$$\Psi_G^{12,13} = \Psi_{G-u} = \Phi_{G\backslash 123}^{\{u\},\{v_1,v_2,v_3\}} = \Phi_{G-u}^{\{v_1,v_2,v_3\}}.$$

Consider the Dodgson polynomial $\Psi_{G,1}^{2,3}$. To find its expression in terms of spanning forest polynomials, we need to look for sets of edges of $G\backslash 123$ that give a spanning tree in $G\backslash 2/13$ and in $G\backslash 3/12$. These two smaller graphs are the second and third graphs in Figure 3. If a set of edges is a forest for both of these graphs, then on $G\backslash 123$ we must have v_1 and v_2 in different trees as well as v_1 and v_3 in different trees; u must also be in a different tree from all the other vertices. Furthermore, to get a tree in the two minors of Figure 3, the spanning forest of $G\backslash 123$ must have exactly three trees, including the tree consisting of the isolated vertex u alone. There is one partition of $\{u, v_1, v_2, v_3\}$ that satisfies these properties, namely $\{u\}, \{v_1\}, \{v_2, v_3\}$, and all forests compatible with this partition give trees in the two minors of Figure 3. Therefore,

$$\Psi_{G,1}^{2,3} = \Phi_{G\backslash 123}^{\{u\},\{v_1\},\{v_2,v_3\}}.$$

Since the vertex u is isolated in $G\backslash 123$, we can rewrite this without u as

$$\Psi_{G,1}^{2,3} = \Phi_{G-u}^{\{v_1\},\{v_2,v_3\}}.$$

We can use the particular Dodgson polynomials from the example to compute the c_2 invariant more easily. Continue for the next two lemmas with the notation $[F]_p$ for the number of $\mathbb{F}_p$-rational points on the affine variety defined by the polynomial F reduced modulo p. The polynomials we will deal with will always come from graphs, and so the dimension of the affine space in question will be the number of edges of the graph.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

Lemma 2.5 ([8, Lemma 24] along with inclusion-exclusion) *Suppose G is a graph with $2+|E(G)| \le 2|V(G)|$. Let i, j, k be distinct edge indices of G and let p be a prime. Then*

$$c_2^{(p)}(G) = -[\Psi_G^{ik,jk}\Psi_{G,k}^{i,j}]_p \mod p.$$

Note that decompletions of 4-regular graphs satisfy the hypotheses of this lemma. This lemma is useful, because we no longer have to divide by p^2, but rather can obtain the c_2 invariant directly as a point count modulo p.

The last lemma we need is a corollary of one of the standard proofs of the Chevalley–Warning theorem; see [1, Section 2].

Lemma 2.6 *Let F be a polynomial of degree N in N variables, $x_1, \dots, x_N$, with integer coefficients. The coefficient of $x_1^{p-1}\cdots x_N^{p-1}$ in F^{p-1} is $[F]_p$ modulo p.*

This last lemma is particularly useful when used in conjunction with the previous lemma. Suppose G is a decompletion of a 4-regular graph and i, j, k are distinct edge indices. Then $\Psi_G^{ik,jk}\Psi_{G,k}^{i,j}$ has both degree and number of variables equal to the number of edges of $G\backslash ijk$. Thus, $c_2^{(p)}(G)$ is equal to the coefficient of $x_1^{p-1}\cdots x_N^{p-1}$ in $(\Psi_G^{ik,jk}\Psi_{G,k}^{i,j})^{p-1}$ modulo p. In view of this, we will no longer need to consider point counts and so that frees up square brackets for the usual algebraic combinatorics notation for the *coefficient-of* operator. Namely, the coefficient of a monomial m in a polynomial F is denoted $[m]F$. Rewriting the previous observation in this notation, we get

$$c_2^{(p)}(G) = [x_1^{p-1}\cdots x_N^{p-1}](\Psi_G^{ik,jk}\Psi_{G,k}^{i,j})^{p-1} \mod p,$$

where N is the number of edges of $G\backslash ijk$. If we further suppose that i, j, and k meet at a 3-valent vertex u and their other ends are v_1, v_2, and v_3 (this case will be sufficient for our purposes), then we can incorporate the computations of Example 2.4 to get

$$c_2^{(p)}(G) = [x_1^{p-1}\cdots x_N^{p-1}](\Phi_{G-u}^{\{v_1\},\{v_2,v_3\}}\Psi_{G-u})^{p-1} \mod p. \tag{2.1}$$

The polynomials do in general depend on the choice of v_1, though the coefficient in question modulo p clearly does not. We will use the freedom of choice of v_1 later.

When $p = 2$, (2.1) has a particularly nice graphical interpretation that can be derived as follows. We are interested in the parity of $[x_1\cdots x_N]\Phi_{G-u}^{\{v_1\},\{v_2,v_3\}}\Psi_{G-u}$, but the coefficient of $x_1\cdots x_N$ in $\Phi_{G-u}^{\{v_1\},\{v_2,v_3\}}\Psi_{G-u}$ is the number of ways to partition the variables into two monomials, one from $\Phi_{G-u}^{\{v_1\},\{v_2,v_3\}}$ and one from Ψ_{G-u}, since both of these polynomials[4] have all monomials appearing with coefficient 1. The variables correspond to edges, so this is equivalent to counting the number of ways to partition the edges of $G-u$ into two parts, one part when removed gives a spanning tree and the other part when removed gives a spanning 2-forest compatible with $\{v_1\}, \{v_2, v_3\}$. Swapping the roles of the two parts, this is equivalent to counting the number of ways to partition the edges of $G-u$ into two parts, one of which is a spanning tree and the

[4]The same idea works for more general products of Dodgson polynomials, but some care must be taken as Dodgson polynomials are in general signed sums of spanning forest polynomials. For some examples computing in this way, see [13, 26].

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

other of which is a spanning 2-forest compatible with $\{v_1\}, \{v_2, v_3\}$. The c_2 invariant at $p = 2$ is simply the parity of this count, so we are determining c_2 by counting assignments of the edges of $G - u$ to particular spanning trees and forests.

For $p > 2$, the same idea works, but we must assign $p-1$ copies of each edge among $p-1$ copies of each polynomial. We can view this as partitioning the edges of the graph we obtain by taking $G - u$ and replacing each edge with $p - 1$ edges in parallel. There are many more possibilities when partitioning all these multiple edges into $2p - 2$ polynomials; the practicalities of working with this approach for $p > 2$ are daunting.

3 Reduction to a Combinatorial Counting Problem

Take K to be a (fixed) connected 4-regular graph with an odd number of vertices and take v, w to be vertices of K. Since K is connected, there is a path between any two vertices of K, and thus it suffices to prove Theorem 1.2 in the case where v and w are joined by an edge. Thus, add as an assumption that v and w are joined by an edge.

If v and w have three common neighbours then $K - v$ and $K - w$ are isomorphic so the result is trivial. Therefore, we can assume that v and w have at most two common neighbours.

The case when v and w have two common neighbours is special for two reasons. First, when v and w have two common neighbours, there is a double triangle (two triangles sharing an edge) involving v, w, and their common neighbours. The arguments of [8] on double triangle invariance of c_2 have not been extended to the case where one of the vertices of the double triangle is the completion vertex, and so these arguments cannot be used here. However, the double triangle arguments should be readily generalizable to this situation and would work for all p, but the required techniques and setup are somewhat different and so this will not be pursued here. Rather, we will simply leave it as a comment that the case where v and w have exactly two common neighbours should be a consequence of the other cases because of the double triangle, and instead will prove this case directly with the present methods.

Second, [5] the case where v and w have two common neighbours is special enough that we can remove the requirement that K have an odd number of vertices. This argument is suggestive of how the parity restriction should hopefully be removable in general and is discussed in Section 6.

In the end we must consider 0, 1, or 2 common neighbours for v and w. These three cases are illustrated in Figure 4; label the vertices of K as in the figure.

If v and w have no common neighbours, let $R = K - \{v, w\}$; that is R is the grey blob on the right-hand side of Figure 4. If v and w have one common neighbour, let $S = K - \{v, w\}$; that is, S is the grey blob in the middle of Figure 4. Finally, if v and w have two common neighbours, let $T = K - \{v, w\}$; that is, T is the grey blob on the left-hand side of Figure 4.

By the fact that K is 4-regular and has an odd number of vertices, the number of edges of $K - \{v, w\}$ is always $-1 \mod 4$, and so for each of R, S, and T, we will use $x_1, x_2, \ldots, x_{4k-1}$ for the edge variables.

[5] Thanks to a referee for this observation.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

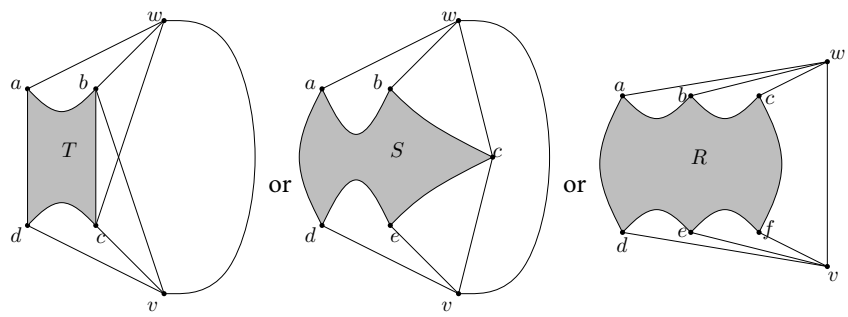

Figure 4: The graph K

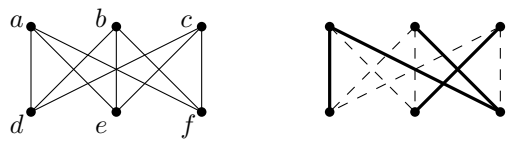

Figure 5: A graph and an edge bipartition of it compatible with $\{a,b\},\{c\}$.

The next step is to use the results of Section 2 to rewrite the $c_2^{(2)}$ in the R, S, and T cases.

Call a set partition with two parts a *bipartition* and use the following notation.

Definition 3.1 Suppose that P is a bipartition of a subset of the vertices of R. Then let $\mathcal{R}_P$ be the set of bipartitions of the edges of R such that one part is a spanning tree of R and the other part is a spanning 2-forest where one tree of the 2-forest contains all vertices of the first part of P and the other contains all vertices of the second part of P. Furthermore, let $r_P = |\mathcal{R}_P|$.

Define $\mathcal{S}_P$ and s_P similarly for a bipartition of a subset of the vertices of S and define $\mathcal{T}_P$ and t_P similarly for T.

For example, if $R = K_{3,3}$ as shown in the first part of Figure 5, then one of the elements of $\mathcal{R}_{\{a,b\},\{c\}}$ is marked by the thick and dotted lines in the second part of the figure. Permuting a and b and permuting d, e, and f, we get 12 elements of $\mathcal{R}_{\{a,b\},\{c\}}$. In this case one other form can occur namely where the isolated vertex c is the tree containing c in the spanning 2-forest. There are 6 such elements, and so $r_{\{a,b\},\{c\}} = 18$ in this case.

Proposition 3.2 *With notation as above, when v and w have no common neighbours,*

$$c_2^{(2)}(K-\{v\}) = r_{\{a,b\}\{c\}} + r_{\{a,c\},\{b\}} + r_{\{b,c\},\{a\}} \mod 2, \tag{3.1}$$

$$c_2^{(2)}(K-\{w\}) = r_{\{d,e\},\{f\}} + r_{\{d,f\},\{e\}} + r_{\{e,f\},\{d\}} \mod 2, \tag{3.2}$$

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

while when v and w have only neighbour c,

$$c_2^{(2)}(K-\{v\}) = s_{\{a,b\},\{c\}} \mod 2, \tag{3.3}$$

$$c_2^{(2)}(K-\{w\}) = s_{\{d,e\},\{c\}} \mod 2, \tag{3.4}$$

and when v and w have neighbours b and c,

$$c_2^{(2)}(K-\{v\}) = t_{\{a\},\{b,c\}} \mod 2, \tag{3.5}$$

$$c_2^{(2)}(K-\{w\}) = t_{\{d\},\{b,c\}} \mod 2. \tag{3.6}$$

Proof By Lemmas 2.5 and 2.6 and Example 2.4, as encapsulated in (2.1), when v and w have no common neighbours,

$$\begin{aligned} c_2^{(2)}(K-\{v\}) &= [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{a,b\},\{c\}}\Psi_R \mod 2 \\ &= [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{a,c\},\{b\}}\Psi_R \mod 2 \\ &= [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{b,c\},\{a\}}\Psi_R \mod 2, \end{aligned}$$

and so by the joys of working modulo 2,

$$\begin{aligned} c_2^{(2)}(K-\{v\}) = {} & [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{a,b\},\{c\}}\Psi_R \\ & + [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{a,c\},\{b\}}\Psi_R \\ & + [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{b,c\},\{a\}}\Psi_R \mod 2. \end{aligned}$$

Similarly,

$$\begin{aligned} c_2^{(2)}(K-\{w\}) = {} & [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{d,e\},\{f\}}\Psi_R \\ & + [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{d,f\},\{e\}}\Psi_R \\ & + [x_1x_2\cdots x_{4k-1}]\Phi_R^{\{e,f\},\{d\}}\Psi_R \mod 2. \end{aligned}$$

Thus, by the edge assignment interpretation discussed at the end of Section 2, $c_2^{(2)}(K-\{v\})$ is equal modulo 2 to the number of ways to partition the edges of R into two parts where one part is a spanning tree and the other part is a spanning 2-forest where one tree of the 2-forest includes two vertices among $\{a,b,c\}$ and the other tree of the 2-forest includes the remaining vertex of $\{a,b,c\}$. Similarly, $c_2^{(2)}(K-\{w\})$ is equal modulo 2 to the number of ways to partition the edges of R into two parts where one part is a spanning tree and the other part is a spanning 2-forest where one tree of the 2-forest includes two vertices among $\{d,e,f\}$ and the other tree of the 2-forest includes the remaining vertex of $\{d,e,f\}$. Restated using our notation, this is the statement of the proposition when v and w have no common neighbours.

When v and w have only common neighbour c, we can calculate similarly. We could again sum over the three possible bipartitions of $\{a,b,c\}$ and similarly for $\{c,d,e\}$. However, the common neighbour c breaks the symmetry, and thus it will be more convenient for the remainder of the argument simply to work with the partitions $\{a,b\},\{c\}$ and $\{c\},\{d,e\}$ giving the simpler result of the statement of the proposition in the case where v and w have common neighbour c.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

Likewise, when v and w have common neighbours b and c, because the symmetry is broken it is most convenient only to take the partitions $\{a\}, \{b, c\}$ and $\{d\}, \{b, c\}$ and otherwise argue as above to obtain the result. ■

In view of the above proposition, to prove Theorem 1.2, it suffices to show that the parity of the number of edge partitions that contribute to the right-hand side of (3.1) is the same as the parity of the number of edges partitions that contribute to the right-hand side of (3.2), and similarly for (3.3) and (3.4) and for (3.5) and (3.6).

Proposition 3.3 *With notation as above, when v and w have no common neighbours,*

$$\begin{aligned} c_2^{(2)}(K-\{v\}) - c_2^{(2)}(K-\{w\}) = {} & r_{\{a,b,d,e,f\},\{c\}} + r_{\{a,c,d,e,f\},\{b\}} \\ & + r_{\{a\},\{b,c,d,e,f\}} + r_{\{a,b\},\{c,d,e,f\}} \\ & + r_{\{a,c\},\{b,d,e,f\}} + r_{\{a,d,e,f\},\{b,c\}} \\ & + r_{\{a,b,c,d,e\},\{f\}} + r_{\{a,b,c,d,f\},\{e\}} \\ & + r_{\{a,b,c,e,f\},\{d\}} + r_{\{a,b,c,f\},\{d,e\}} \\ & + r_{\{a,b,c,e\},\{d,f\}} + r_{\{a,b,c,d\},\{e,f\}} \mod 2 \end{aligned}$$

while when v and w have only neighbour c,

$$\begin{aligned} c_2^{(2)}(K-\{v\}) - c_2^{(2)}(K-\{w\}) = {} & s_{\{a,b,d\},\{c,e\}} + s_{\{a,b,e\},\{c,d\}} \\ & + s_{\{a,b\},\{c,d,e\}} + s_{\{a,d,e\},\{b,c\}} \\ & + s_{\{a,c\},\{b,d,e\}} + s_{\{a,b,c\},\{d,e\}} \mod 2, \end{aligned}$$

and when v and w have neighbour b and c,

$$c_2^{(2)}(K-\{v\}) - c_2^{(2)}(K-\{w\}) = t_{\{a\},\{b,c,d\}} + t_{\{a,b,c\},\{d\}} \mod 2.$$

Proof Consider the right-hand sides of (3.1) and (3.2). Enumerating all possibilities,

$$\begin{aligned} r_{\{a,b\},\{c\}} = {} & r_{\{a,b,d,e,f\},\{c\}} + r_{\{a,b,d,e\},\{c,f\}} + r_{\{a,b,d,f\},\{c,e\}} + r_{\{a,b,e,f\},\{c,d\}} \\ & + r_{\{a,b,d\},\{c,e,f\}} + r_{\{a,b,e\},\{c,d,f\}} + r_{\{a,b,f\},\{c,d,e\}} + r_{\{a,b\},\{c,d,e,f\}} \end{aligned}$$

and similarly for the other terms. Collecting these calculations together, simplifying modulo 2, and performing analogous calculations with regard to the right-hand sides of (3.3), (3.4), (3.5), and (3.6), we get the proposition. ■

It is best, in my view, to keep a graphical viewpoint with the above formulas. For example, one can represent a given term by drawing the graph and marking the partition by using different vertex shapes. Then the second equation of the statement of Proposition 3.3 looks like

$c_2^{(2)}(K-\{v\}) - c_2^{(2)}(K-\{w\}) =$

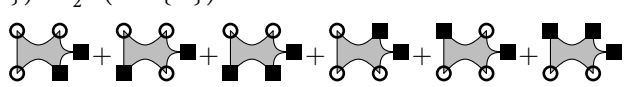

mod 2.

The reader is encouraged to translate all the equations from here on out into this notation in order to better see the intuition behind the argument. Note that this graphical

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

notation is not the same as in [11,13,26], where a graph with a marked partition in this way represented the spanning forest polynomial for that graph and partition.

One's first thought for proceeding from here would likely be to find a fixed-point-free involution of the union of the $\mathcal{R}$, $\mathcal{S}$, or $\mathcal{T}$ sets appearing (via their counts) in Proposition 3.3. It is not clear how to do this for all the $\mathcal{R}$, $\mathcal{S}$, and $\mathcal{T}$ sets, and so the method will be more complicated. First we will deal with some of the $\mathcal{R}$ and $\mathcal{S}$ sets via involutions. For the sets that remain we will build auxiliary graphs that use the properties of certain cycles along with the parity hypotheses on K to show these remaining sets all have an even contribution.

4 Some Involutions from Swapping around Particular Vertices

The involution is simplest in the case where v and w have only common neighbour c and the partition breaks a, b, c, d, e into $\{a, b, c\}, \{d, e\}$ or $\{a, b\}, \{c, d, e\}$, so we will begin by discussing that case.

Lemma 4.1 *Let σ be a bipartition of the edges of S so that each part is a spanning forest (either or both of which may potentially be a spanning forest with one tree, that is a spanning tree), where in each forest every tree contains at least one of a, b, d, or e. Then of the two edges incident to c, exactly one is in each part of σ and swapping which part these two edges are in yields a new partition σ' with the above listed properties, but not necessarily partitioning the vertices among the trees in the same way.*

Proof Note that c is 2-valent in S. Let y and z be the two neighbours of c.

By hypothesis, in each part of σ, the vertex c is connected to some other vertices of S since every tree contains at least one of a, b, d, or e, and so at least one edge incident to c is assigned to this part. Since c is 2-valent, this means that exactly one incident edge to c is in each part of σ. Consider the spanning forest F corresponding to one part of σ. In F, c is a leaf and so removing the edge incident to c isolates c and does not otherwise change the connectivity of the trees in F. Without loss of generality say that y was the other end of this edge. Vertex z is in some tree of F and adding the edge between z and c reconnects c to one of the trees of F while maintaining a forest structure. The same holds for the other part of σ. ∎

Lemma 4.2 *There is a fixed-point free involution on $\mathcal{S}_{\{a,b,c\},\{d,e\}} \cup \mathcal{S}_{\{a,b\},\{c,d,e\}}$ and thus $s_{\{a,b,c\},\{d,e\}} + s_{\{a,b\},\{c,d,e\}} = 0 \mod 2$.*

Proof All edge bipartitions in $\mathcal{S}_{\{a,b,c\},\{d,e\}} \cup \mathcal{S}_{\{a,b\},\{c,d,e\}}$ satisfy the hypotheses of Lemma 4.1. To each such an edge partition swap the parts to which the edges incident to c belong. Consider an edge partition coming from $\mathcal{S}_{\{a,b,c\},\{d,e\}}$. Removing the edges incident to c disconnects c in both the 2-forest and the tree. Putting the edges back in, but in opposite parts of the bipartition, reconnects c to one of the trees of the 2-forest as well as to the single tree. If c is reconnected to the tree of the 2-forest including $\{a, b\}$, then we now have another edge partition in $\mathcal{S}_{\{a,b,c\},\{d,e\}}$. If c is now connected to the tree of the 2-forest involving $\{d, e\}$, then we now have an edge partition in $\mathcal{S}_{\{a,b\},\{c,d,e\}}$. The edge partition must be distinct from the initial partition

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

since which edge incident to c corresponds to the tree has changed. This operation is clearly an involution. Thus we get a fixed-point free involution on $\mathcal{S}_{\{a,b,c\},\{d,e\}} \cup \mathcal{S}_{\{a,b\},\{c,d,e\}}$. Consequently, the size of this set is even. ∎

We can also apply the swapping map described in the previous proof to the other $\mathcal{S}$ sets. However, an edge partition in $\mathcal{S}_{\{a,b,d\},\{c,e\}}$ can either be mapped to another partition from the same set or it can be mapped to $\mathcal{S}_{\{a,b,c,d\},\{e\}}$, which is not an $\mathcal{S}$ set appearing in Proposition 3.3. An analogous situation occurs for the other $\mathcal{S}$ sets not dealt with in the previous lemma. For these remaining $\mathcal{S}$ sets, use the methods of Section 5 instead.

Next we look at some of the $\mathcal{R}$ sets that can be tackled with a generalized version of the above argument.

Lemma 4.3 *Let $p_1 \cup p_2$ be a partition of $\{a, b, c, d, e, f\}$, where p_1 consists of either all of $\{a, b, c\}$ and exactly one of $\{d, e, f\}$ or all of $\{d, e, f\}$ and exactly one of $\{a, b, c\}$. Let x be the element of p_1 that is alone from its trio. Consider any edge partition τ in $\mathcal{R}_{p_1,p_2}$. Let t be the tree of the 2-forest of τ associated with p_1.*

There is a unique vertex y with the following properties:

(i) *Either $y \in p_1$ and is 2-valent in t or $y \notin p_1$ and y is 3-valent in t.*
(ii) *Removing y from t gives a component containing exactly two vertices of p_1 and a component containing exactly one vertex of p_1 (hence the third component, if it exists also contains exactly one vertex of p_1).*
(iii) *Either $x = y$ or x is in one of the components that, after removing y, contains exactly one vertex of p_1.*

We will call the vertex y defined in the above lemma the *control vertex* of τ, and we call the vertex x the *outsider vertex*.

Proof Since the spanning tree of τ spans, every vertex of t is at most 3-valent and the vertices of p_1 are at most 2-valent in t. The union of the paths in t between the vertices of p_1 gives a subtree of t where every leaf is in p_1. There are only finitely many configurations for this subtree; these are illustrated in Figure 6 where the edges in the figure represent paths in t. For each configuration, the three properties and uniqueness can be checked directly, remembering that the degree of y remains the same in t as in the subtree. ∎

Note that all the different configurations in Figure 6 can be viewed as special cases of the top left configuration where some of the paths have been contracted.

Lemma 4.4 *There is a fixed-point free involution on*

$$\mathcal{R}_{\{a,b\},\{c,d,e,f\}} \cup \mathcal{R}_{\{a,c\},\{b,d,e,f\}} \cup \mathcal{R}_{\{a,d,e,f\},\{b,c\}} \cup \mathcal{R}_{\{a,b,c,d\},\{e,f\}} \cup \mathcal{R}_{\{a,b,c,e\},\{d,f\}} \cup \mathcal{R}_{\{a,b,c,f\},\{d,e\}},$$

and thus

$$r_{\{a,b\},\{c,d,e,f\}} + r_{\{a,c\},\{b,d,e,f\}} + r_{\{a,d,e,f\},\{b,c\}} + r_{\{a,b,c,d\},\{e,f\}} + r_{\{a,b,c,e\},\{d,f\}} + r_{\{a,b,c,f\},\{d,e\}} = 0 \mod 2.$$

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

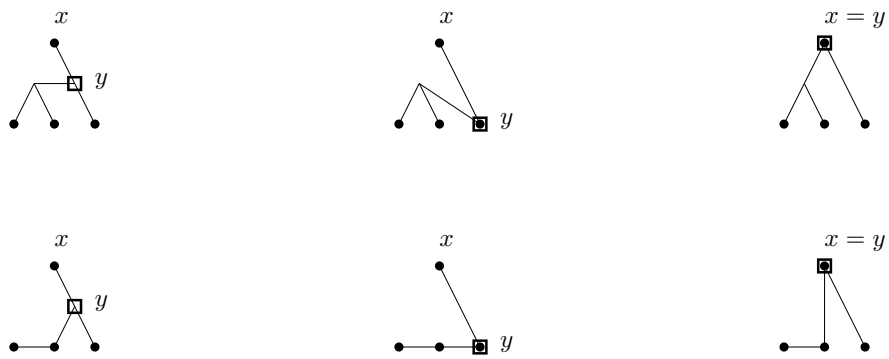

Figure 6: Schematic determining the control vertex y. The lower three vertices in each schematic are $p_1 - \{x\}$ in any order.

Proof Let τ be an edge bipartition in the union of $\mathcal{R}$ sets in the statement of the lemma. We need to set up some notation to build the involution.

- All of these $\mathcal{R}$ sets satisfy the hypotheses of Lemma 4.3, and so let y be the control vertex of τ.
- There is exactly one edge incident to y in the spanning tree of τ. Let ϵ be that edge.
- Let z be the end of ϵ that is not y.
- Let t_1 be the tree of the 2-forest of τ that contains y and let t_2 be the other tree of the 2-forest.

Now build τ' from τ as follows.

- If $z \in t_2$, then let η be the edge incident to y that leads to the component of $t_1 - \{y\}$ with two vertices from $\{a, b, c, d, e, f\}$. Swap which part of the bipartition τ contains ϵ and which contains η to obtain τ'.
- If $z \in t_1$, then let η be the edge incident to y which leads to the component of $t_1 - \{y\}$ that contains z. Swap which part of the bipartition τ contains ϵ and which contains η to obtain τ'.

First let us check that τ' is in the union of $\mathcal{R}$ sets in the statement. Similarly to the proof of Lemma 4.1, y is a leaf in the spanning tree of τ so removing ϵ disconnects y from the spanning tree and adding η reconnects y maintaining a spanning tree structure. Removing η further disconnects the 2-forest of τ into three components. If we constructed τ' by the second case, then adding ϵ to the 2-forest reconnects the same components that were disconnected by the removal of η. If we constructed τ' by the first case, then removing η cuts off the component of $t_1 - \{y\}$ containing exactly two vertices from $\{a, b, c, d, e, f\}$ from the rest of t_1; adding ϵ reconnected t_2 instead. Furthermore, the outsider vertex is in the part of t_1 that gets connected with t_2. The result is an edge partition that corresponds to a partition of $\{a, b, c, d, e, f\}$ satisfying the hypotheses of Lemma 4.3, and so is in one of the $\mathcal{R}$ sets in the statement.

Next note that the map $\tau \mapsto \tau'$ is fixed-point free, since which edge of the control vertex is in the spanning tree changes.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

Finally, the control vertex of τ and τ' are the same and so applying the map twice is the identity. Thus, we get a fixed-point free involution on the union of $\mathcal{R}$ sets in the statement, and so the size of this set is even. ■

Note that in the case where v and w had a common neighbour c, c always plays the role of the control vertex, but compared to the configurations in Figure 6, more of the paths have been contracted away. In this way, the case with no common neighbours generalized the simpler argument in the common neighbour case.

5 Compatible Cycles

This section defines certain special cycles and investigates their properties. These cycles will let us determine the parity of the remaining r_P and s_P in Proposition 3.3 as well as the t_P.

Definition 5.1

(i) Call a bipartition of the edges of any graph (for our purposes, either R, S, or T) such that one part gives a spanning tree and the other part gives a spanning 2-forest a *valid* edge partition.

(ii) Suppose we have a valid edge partition. This gives a bipartition of all of the vertices of the graph according to which tree of the 2-forest they are in. Call a cycle C *compatible* with the edge partition if all vertices of C are in the same part of the vertex partition and exactly one edge of C is in the part of the edge partition corresponding to the spanning tree.

Lemma 5.2 *Suppose we have a valid edge partition. The number of compatible cycles is the same as the number of edges of the graph that are in the part of the edge partition corresponding to the spanning tree but where both ends of the edge are in the same tree of the* 2*-forest.*

Proof Adding an edge joining two vertices of a tree gives a graph with a unique cycle. When this fact is applied to one of the trees of the 2-forest, this cycle is compatible and every compatible cycle has this form. ■

Lemma 5.3 *Suppose we have a valid edge partition and let V_1, V_2 be the associated vertex partition. Suppose that $\sum_{v\in V_i} \deg v$ is odd for $i = 1, 2$, and the total number of vertices of the graph is odd. Then the number of compatible cycles is odd.*

Proof Let ℓ be the number of edges crossing the vertex partition and let e_i be the number of edges not in the 2-forest but with both ends in the tree of V_i for $i = 1, 2$. The number of edges leaving the tree of the 2-forest associated with V_i is

$$\sum_{v\in V_i} \deg v - 2(|V_i| - 1) - 2e_i = \ell$$

for $i = 1, 2$. Therefore, by the degree hypothesis, ℓ is odd. The number of edges of the spanning tree of the valid edge partition is

$$|V_1| + |V_2| - 1 = e_1 + e_2 + \ell.$$

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

However, $|V_1| + |V_2|$ is the total number of vertices of the graph, and so by hypothesis is odd. Therefore, $e_1 + e_2$ is also odd. By Lemma 5.2, the number of compatible cycles is $e_1 + e_2$, and hence is odd, as desired. ∎

Now we want to apply this lemma to the $\mathcal{R}$ and $\mathcal{S}$ of Proposition 3.3 sets that were not dealt with in the previous section, as well as the $\mathcal{T}$ sets of Proposition 3.3.

Lemma 5.4 *Every edge partition of*

$$\begin{aligned}\mathcal{R}_{\{a,b,d,e,f\},\{c\}} &\cup \mathcal{R}_{\{a,c,d,e,f\},\{b\}} \cup \mathcal{R}_{\{a\},\{b,c,d,e,f\}} \cup \mathcal{R}_{\{a,b,c,d,e\},\{f\}}\\ &\cup \mathcal{R}_{\{a,b,c,d,f\},\{e\}} \cup \mathcal{R}_{\{a,b,c,e,f\},\{d\}}\\ &\cup \mathcal{S}_{\{a,b,d\},\{c,e\}} \cup \mathcal{S}_{\{a,b,e\},\{c,d\}} \cup \mathcal{S}_{\{a,c\},\{b,d,e\}} \cup \mathcal{S}_{\{b,c\},\{a,d,e\}}\\ &\cup \mathcal{T}_{\{a\},\{b,c,d\}} \cup \mathcal{T}_{\{a,b,c\},\{d\}}\end{aligned}$$

has an odd number of compatible cycles.

Proof It suffices to check the hypotheses of Lemma 5.3.

The number of vertices of each of R, S, and T is odd by our running assumptions on K.

Take any edge partition in the union above. All the vertices other than $\{a, b, c, d, e, f\}$ are degree 4, so it suffices to check that the sums of the degrees of the vertices in each part of the defining partition of $\{a, b, c, d\}$, or $\{a, b, c, d, e\}$, or $\{a, b, c, d, e, f\}$ are odd. For the $\mathcal{R}$ sets, one part has degree sum $5 \cdot 3$, and the other has degree sum 3; for the $\mathcal{S}$ sets, one part has degree sum $3 \cdot 3$, and the other has degree sum $3 + 2$, and for the $\mathcal{T}$ sets, one part has degree sum 3, and the other has degree sum $3 + 2 + 2$. All of these are odd. ∎

Lemma 5.5 *Suppose we have a valid edge partition and a compatible cycle C. Let f be the one edge of C not in the 2-forest. If we were to remove f, the spanning tree would split into exactly two trees, call them t_1 and t_2. There are a non-zero even number of edges of C with one end in t_1 and the other end in t_2 (including f as one of the possibilities). If we take any such edge other than f and swap it with f in the edge partition, then we obtain a valid edge partition corresponding to the same vertex partition.*

Proof By construction, f has one end in t_1 and one end in t_2. The vertices of C can be bipartitioned based on whether they are in t_1 or in t_2. Running around C, we must change which part of the bipartition we are in an even number of times in order to return to where we started, giving a non-zero even number of edges of the type described in the statement.

Let $f' \neq f$ be another edge of C where we change from t_1 to t_2. Removing f from the spanning tree disconnects it into t_1 and t_2 while adding f' reconnects t_1 and t_2 to obtain a spanning tree again. Adding f to the 2-forest creates one cycle, specifically C. Removing any edge of C, in particular f', returns us to a 2-forest with the same vertex partition. ∎

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

Lemma 5.6

(i) $$r_{\{a,b,d,e,f\},\{c\}} + r_{\{a,c,d,e,f\},\{b\}} + r_{\{a\},\{b,c,d,e,f\}} + r_{\{a,b,c,d,e\},\{f\}} + r_{\{a,b,c,d,f\},\{e\}} + r_{\{a,b,c,e,f\},\{d\}} = 0 \mod 2.$$

(ii) $$s_{\{a,b,d\},\{c,e\}} + s_{\{a,b,e\},\{c,d\}} + s_{\{a,c\},\{b,d,e\}} + s_{\{b,c\},\{a,d,e\}} = 0 \mod 2.$$

(iii) $$t_{\{a\},\{b,c,d\}} + t_{\{a,b,c\},\{d\}} = 0 \mod 2.$$

Proof The construction is the same for all three cases. It will be described explicitly in the $\mathcal{R}$ case. Construct a graph X_R as follows. The vertices of X_R are the edge partitions in

$$\mathcal{R}_{\{a,b,d,e,f\},\{c\}} \cup \mathcal{R}_{\{a,c,d,e,f\},\{b\}} \cup \mathcal{R}_{\{a\},\{b,c,d,e,f\}} \cup \mathcal{R}_{\{a,b,c,d,e\},\{f\}} \cup \mathcal{R}_{\{a,b,c,d,f\},\{e\}} \cup \mathcal{R}_{\{a,b,c,e,f\},\{d\}}.$$

Two vertices of X_R are adjacent if they are related by a swap as given in Lemma 5.5. Note that if a given edge assignment goes to another via such a swap, then the second also goes to the first by such a swap, since the cycle is compatible for either edge partition and the t_1, t_2 partition (in the notation of the proof of Lemma 5.5) is also the same for both edge partitions.

By Lemma 5.4, for any vertex x in X_R, there are an odd number of cycles that can yield swaps corresponding to edges incident to x. Distinct cycles must give distinct swaps, hence distinct edges. By Lemma 5.5, each one of these cycles gives an odd number of edges incident to x (one for each edge of the type described in the statement of Lemma 5.5 other than f itself). All edges incident to x are obtained in this way, so x has odd degree. This is true for all vertices of X_R, but by basic counting, any graph has an even number of vertices of odd degree, so X_R has an even number of vertices.

Therefore, the union of $\mathcal{R}$ sets defining X_R has even size, which is the first statement of the lemma.

The argument for X_S and X_T is analogous, using the edge partitions in

$$\mathcal{S}_{\{a,b,d\},\{c,e\}} \cup \mathcal{S}_{\{a,b,e\},\{c,d\}} \cup \mathcal{S}_{\{a,c\},\{b,d,e\}} \cup \mathcal{S}_{\{b,c\},\{a,d,e\}}$$

and

$$\mathcal{T}_{\{a\},\{b,c,d\}} \cup \mathcal{T}_{\{a,b,c\},\{d\}},$$

respectively. ■

Note that the construction of X_R, X_S, and X_T is closely related to the spanning tree graph (often just called the *tree graph*, see [15]) construction. The vertices of the tree graph of a graph G are the spanning trees of G, and two vertices of the tree graph are joined by an edge if the two spanning trees differ by removing one edge and replacing it with another.

This is all we need to prove the main theorem.

Proof of Theorem 1.2 As discussed at the beginning of Section 3 it suffices to consider v and w joined by an edge and with zero, one, or two common neighbours.

By Proposition 3.3, we need only check that the parity of a certain sum of r_{p_1,p_2} is even in the case where v and w have no common neighbours, or a certain sum of s_{p_1,p_2} is even in the case where v and w have one common neighbour, or a certain sum of t_{p_1,p_2} is even in the case where v and w have two common neighbours.

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

In the case where v and w have two common neighbours, the third part of Lemma 5.6 gives that the required sum is even.

In the case where v and w have one common neighbour, Lemma 4.2 and the second part of Lemma 5.6 give that the required sum is even.

In the case where v and w have no common neighbours, Lemma 4.4 and the first part of Lemma 5.6 give that the required sum is even. ■

6 Discussion

It should be the case that Theorem 1.2 is true for all p and without the restriction that K has an odd number of vertices; see [8, conjecture 4]. This would be an excellent conjecture to prove, because it would support the deep connection between the c_2 invariant and the Feynman period. It is a surprisingly difficult and very interesting conjecture.

The restriction to $p = 2$ came about because Lemma 2.6 is much simplified in the $p = 2$ case. For higher values of p, there is a still an edge assignment interpretation as discussed at the end of Section 2, but each edge must be assigned $p-1$ times to build $p-1$ spanning trees and $p-1$ spanning 2-forests. This greatly increases the complexity. However, in principle these ideas may be extendable to $p > 2$. The practicalities would certainly be very difficult. The question is whether or not the practicalities would be so difficult as to render the approach unworkable.

The vertex parity condition is more mysterious. Note that it is only needed for the arguments of Section 5 and not for the arguments of Section 4. If the number of vertices of K were even, then the required degree sum parity in Lemma 5.3 would also need to change to preserve the total number of compatible cycles being odd. This translates into the methods of Section 5 applying to the other $\mathcal{R}$, $\mathcal{S}$, and $\mathcal{T}$ sets, namely the ones we already know how to tackle by Section 4. So if K had an even number of vertices then the terms that in the odd case matched by swapping around the control vertex would become the terms that are even by compatible cycles, while the terms, which used be even by compatible cycles would need a new argument. The first place to look for this new argument would be as a generalized control vertex argument; however, it is not clear how to do it. There is no obvious obstruction; rather the required construction is simply not apparent, and so some further cleverness is required to progress.

In the T case, that is the double triangle case where v and w have two common neighbours, things are sufficiently special that we can extend the argument to remove the parity condition. This was graciously pointed out by a referee. This argument manages to succeed by running through additional $\mathcal{T}$ sets that do not appear directly in the expression for c_2 but occur while swapping. I had tried similar things in other cases without finding a path to the result, but the approach remains promising, as this argument shows. Consider $\tau \in \mathcal{T}_{\{a\},\{b,c,d\}} \cup \mathcal{T}_{\{a,b\},\{c,d\}}$. The vertex b is 2-valent and one of its incident edges is in each of the tree and the 2-forest of τ since in both cases b must connect to other vertices. Thus, by the same argument as Lemma 4.1, we can swap the edges around b and obtain a fixed-point free involution on $\mathcal{T}_{\{a\},\{b,c,d\}} \cup \mathcal{T}_{\{a,b\},\{c,d\}}$. Therefore,

$$t_{\{a\},\{b,c,d\}} + t_{\{a,b\},\{c,d\}} = 0 \mod 2.$$

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

Symmetrically, by swapping around vertex c in $\mathcal{T}_{\{a,b,c\},\{d\}} \cup \mathcal{T}_{\{a,b\},\{c,d\}}$, we get

$$t_{\{a,b,c\},\{d\}} + t_{\{a,b\},\{c,d\}} = 0 \mod 2.$$

Adding these two equations and applying Proposition 3.3 we get

$$\begin{aligned} c_2^{(2)}(K - \{v\}) - c_2^{(2)}(K - \{w\}) &= t_{\{a\},\{b,c,d\}} + t_{\{a,b,c\},\{d\}} \mod 2 \\ &= 0 \mod 2. \end{aligned}$$

The statement and proof of Theorem 1.2 arose out of discussions with Dmitry Doryn about looking for special cases where we could hope to progress on understanding the c_2 invariant. The first version of the result had many additional hypotheses including planarity and being restricted to the S-case. With some work, most of the extra hypotheses dropped away, and the current Theorem 1.2 remained. It is not an entirely satisfactory theorem as it stands, but it introduces some new ideas to the game and makes nontrivial progress on the completion conjecture for the c_2 invariant for almost the first time. Hopefully, as the first crack in that conjecture, it will lead, with the help of others, to a proof of the full conjecture.

Acknowledgments The author thanks Dmitry Doryn for discussions that inspired this work and the IBS Centre for Geometry and Physics in Korea for their hospitality in hosting those discussions as well as their financial support from IBS-R003-S1 for the visit. The author thanks Iain Crump for many discussions about graphical interpretations of the c_2 invariant over the last few years. The author also thanks the referees.

References

[1] J. Ax, *Zeroes of polynomials over finite fields.* Amer. J. Math. **86**(1964), 255–261. http://dx.doi.org/10.2307/2373163

[2] P. Belkale and P. Brosnan, *Matroids, motives, and a conjecture of Kontsevich.* Duke Math. J. **116**(2003), 147–188. http://dx.doi.org/10.1215/S0012-7094-03-11615-4

[3] S. Bloch, H. Esnault, and D. Kreimer, *On motives associated to graph polynomials.* Comm. Math. Phys. **267**(2006), 181–225. http://dx.doi.org/10.1007/s00220-006-0040-2

[4] D. Broadhurst and O. Schnetz, *Algebraic geometry informs perturbative quantum field theory.* In PoS, volume LL2014, page 078, 2014. arxiv:1409.5570

[5] D. J. Broadhurst and D. Kreimer, *Knots and numbers in ϕ^4 theory to 7 loops and beyond.* Internat. J. Modern Phys. C **6**(1995), 519–524. http://dx.doi.org/10.1142/S012918319500037X

[6] F. Brown, *Feynman amplitudes, coaction principle, and cosmic Galois group.* Commun. Number Theory Phys. **11**(2017), no. 3, 453–556. http://dx.doi.org/10.4310/CNTP.2017.v11.n3.a1

[7] ———, *On the periods of some Feynman integrals.* arxiv:0910.0114

[8] F. Brown and O. Schnetz, *A K3 in ϕ^4.* Duke Math J. **161**(2012), 1817–1862. http://dx.doi.org/10.1215/00127094-1644201

[9] ———, *Modular forms in quantum field theory.* Commun. Number Theory Phys. **7**(2013), 293–325. http://dx.doi.org/10.4310/CNTP.2013.v7.n2.a3

[10] F. Brown, O. Schnetz, and K. Yeats, *Properties of c_2 invariants of Feynman graphs.* Adv. Theor. Math. Phys. **18**(2014), 323–362. http://dx.doi.org/10.4310/ATMP.2014.v18.n2.a2

[11] F. Brown and K. Yeats, *Spanning forest polynomials and the transcendental weight of Feynman graphs.* Comm. Math. Phys. **301**(2011), 357–382. http://dx.doi.org/10.1007/s00220-010-1145-1

[12] S. Chaiken, *A combinatorial proof of the all minors matrix tree theorem.* SIAM J. Alg. Disc. Meth. **3**(1982), 319–329. http://dx.doi.org/10.1137/0603033

[13] W. Chorney and K. Yeats, *c_2 invariants of recursive families of graphs.* arxiv:1701.01208

[14] I. Crump, *Graph Invariants with Connections to the Feynman Period in 4 Theory.* PhD thesis, Simon Fraser University, 2017. arxiv:1704.06350

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press

[15] R. L. Cummins, *Hamilton circuits in tree graphs.* IEEE Trans. Circuit Theory **CT-13**(1966), 82–90. http://dx.doi.org/10.1109/TCT.1966.1082546
[16] R. Diestel, *Graph theory.* Graduate Texts in Mathematics, 173, Springer-Verlag, New York, 1997.
[17] D. Doryn, *The c_2 invariant is invariant.* arxiv:1312.7271
[18] ———, *Dual graph polynomials and a 4-face formula.* arxiv:1508.03484
[19] C. Itzykson and J.-B. Zuber, *Quantum field theory.* International Series in Pure and Applied Physics, McGraw-Hill, 1980. Dover edition 2005.
[20] A. Logan, *New realizations of modular forms in Calabi-Yau threefolds arising from ϕ^4 theory.* arxiv:1604.04918
[21] M. Marcolli, *Feynman motives.* World Scientific, Hackensack, NJ, 2010.
[22] O. Schnetz, *Numbers and functions in quantum field theory.* arxiv:1606.08598
[23] ———, *Quantum periods: A census of ϕ^4-transcendentals.* Commun. Number Theory Phys. **4**(2010), 1–47. http://dx.doi.org/10.4310/CNTP.2010.v4.n1.a1
[24] ———, *Quantum field theory over $\mathbb{F}_q$.* Electron. J. Combin. **18**(2011), no. 1, Paper 102.
[25] A. Vlasev and K. Yeats, *A four-vertex, quadratic, spanning forest polynomial identity.* Electron. J. Linear Algebra **23**(2012), 923–941. http://dx.doi.org/10.13001/1081-3810.1566
[26] K. Yeats, *A few c_2 invariants of circulant graphs.* Commun. Number Theory Phys. **10**(2016), 63–86. http://dx.doi.org/10.4310/CNTP.2016.v10.n1.a3

Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON N2L 3G1
e-mail: kayeats@uwaterloo.ca

https://doi.org/10.4153/CJM-2018-006-5 Published online by Cambridge University Press