

Differential graded algebras for trivalent plane graphs and their representations

Kevin Sackel

Abstract. To any trivalent plane graph embedded in the sphere Casals and Murphy associate a differential graded algebra (dg-algebra), in which the underlying graded algebra is free associative over a commutative ring. Our first result is the construction of a generalization of the Casals–Murphy dg-algebra to non-commutative coefficients, for which we prove various functoriality properties not previously verified in the commutative setting. Our second result is to prove that rank r representations of this dg-algebra, over a field $\mathbb{F}$, correspond to colorings of the faces of the graph by elements of the Grassmannian $\text{Gr}(r, 2r; \mathbb{F})$ so that bordering faces are transverse, up to the natural action of $\text{PGL}_{2r}(\mathbb{F})$. Underlying the combinatorics, the dg-algebra is a computation of the fully non-commutative Legendrian contact dg-algebra for Legendrian satellites of Legendrian 2-weaves, though we do not prove as such in this paper. The graph coloring problem verifies that, for Legendrian 2-weaves, rank r representations of the Legendrian contact dg-algebra correspond to constructible sheaves of microlocal rank r . This is the first explicit such computation of the bijection between the moduli spaces of representations and sheaves for an infinite family of Legendrian surfaces.

Contents

1. Introduction	2
2. Preliminaries on stable regenerative tame dg-homomorphisms	21
3. The coefficient ring	24
4. The enlarged non-commutative Casals–Murphy dg-algebra	42
5. Functoriality for the enlarged dg-algebra	54
6. Removing the enlargements	80
7. Representations and rank r face colorings	87
A. Computational example	104
References	106

Mathematics Subject Classification 2020: 53D10 (primary); 05C15, 53D42, 57R17 (secondary).

Keywords: trivalent graph, differential algebra, Legendrian contact homology, Legendrian weaves, augmentations are sheaves.

1. Introduction

In this paper, a *trivalent plane graph* G will always mean a non-empty connected trivalent planar graph together with an embedding into S^2 , up to isotopy of embeddings. The adjective *trivalent* indicates that each vertex has exactly three incident edges, whereas the adjective *plane* indicates that such an isotopy class of embedding has been fixed.¹ Once we have fixed an embedding, we obtain a decomposition of S^2 into faces by cutting along the edges of the graph. Throughout, we will denote by V , E , and F the set of vertices, edges, and faces of the graph G .

Casals and Murphy [3] have constructed for each trivalent plane graph G , upon making a number of auxiliary choices, a certain algebraic invariant in the form of a differential graded algebra (dg-algebra). Our first goal in this paper is to build a souped-up version of their dg-algebra. Whereas their underlying algebra is free associative over the group ring $\mathbb{Z}[H_1(\Lambda_G)]$, where Λ_G is topologically a branched double cover of S^2 branched over V , our underlying algebra is instead free associative over the non-commutative group ring $\mathbb{Z}[\pi_1(\Lambda_G)]$, where the coefficients no longer commute with the generators, and where we must keep track of a choice of base point. Upon abelianization of the coefficient ring and commuting coefficients with generators, our dg-algebra recovers the Casals–Murphy version. It was shown in the appendix to the paper of Casals and Murphy, written by the author of the present article, that their dg-algebra encodes certain graph coloring data. As a second goal, and indeed the initial impetus for this work, we extend this graph coloring result to the fully non-commutative setting.

Although this article is mostly presented in the guise of algebraic combinatorics with a few details from algebraic topology, on a deeper layer lies contact geometry, which provides context for the present article, as we now describe.

Associated to a trivalent plane graph is a Legendrian surface, the surface Λ_G of the previous paragraph, as was defined by Treumann and Zaslow [28]. Associated to any Legendrian (in a nice enough contact space) is a Legendrian contact dg-algebra, as originally discussed by Eliashberg [13] and simultaneously computed for Legendrian knots by Chekanov [7], and later formally shown to be well-defined by Ekholm, Etnyre, and Sullivan [9, 11]. The dg-algebra constructed by Casals and Murphy, though presented purely combinatorially, is a computation of a certain version of the Legendrian contact dg-algebra attached to Λ_G . The computation utilizes gradient flow trees, a technique pioneered by Ekholm [8]. Our fully non-commutative dg-algebra simply keeps track of a little more geometric detail: the homotopy classes

¹We use the adjective *plane* instead of *planar* intentionally. A planar graph is abstractly embeddable in S^2 . Meanwhile, the choice of (isotopy class of) embedding is important for our purposes and will affect all constructions which follow.

(as opposed to the homology classes) of the curves which comprise the boundary of the involved gradient flow trees.

Our graph coloring result also admits a contact-geometric interpretation. With underpinnings arising in the work of Nadler and Zaslow [22] and Nadler [21] providing a quasi-equivalence between the Fukaya category of a cotangent bundle with a category of constructible sheaves on the base, it was conjectured that representations of the Legendrian contact dg-algebra correspond to certain subclasses of constructible sheaves (see the work of Chantraine, Dimitroglou Rizell, Ghiggini, and Golovko [5] for the higher-rank version). Since this article originally appeared on the arXiv, “representations are sheaves” was proved in full generality by Asplund and Ekholm [2]. Indeed, they proved a conjecture of Ekholm and Lekili [12] that the Legendrian contact dg-algebra with coefficients given by chains on the based loop space is quasi-equivalent to a certain partially wrapped Floer cohomology A_∞ -algebra, modules over which were proved by Ganatra, Pardon, and Shende [17] to correspond to constructible sheaves.

This article is more hands-on, continuing a history of explicit computations of this correspondence that are interesting in their own right and difficult in general to produce from the abstract theory alone. For Legendrian knots in $\mathbb{R}^3$, and rank 1 representations, called augmentations, the correspondence was verified as an A_∞ -categorical equivalence by Ng, Rutherford, Shende, Sivek, and Zaslow [23]. Higher rank representations on Legendrian knots have also been understood for torus knots by Chantraine, Ng, and Sivek [6]. Increasing to contact dimension 5 (and Legendrians of dimension 2), there have been explicit verifications for the rank 1 version of the correspondence for knot and link conormals by Gao [18, 19] and for cellular Legendrian surfaces by Rutherford and Sullivan [26]. The graph coloring result in this paper provides a bijection between representations and sheaves for all Treumann–Zaslow Legendrian surfaces Λ_G , marking the first explicit non-trivial computation of the higher-rank correspondence for an infinite collection of Legendrian surfaces.

We will keep the contact geometry mostly in the background, with two notable exceptions. First, throughout the introduction, we will expand, when appropriate, upon the short description above, primarily for the purposes of providing motivation for the results and constructions we discuss in this article. This includes a contact-geometric description of much of the combinatorial package in Section 1.4 so that contact geometers should feel satisfied that our dg-algebra is indeed the Legendrian contact dg-algebra, even if it is not fully proved. Second, we fundamentally use the fact that Λ_G is a double cover of S^2 branched over the vertices V of G to find a family of combinatorial models for the coefficient ring $\mathbb{Z}[\pi_1(\Lambda_G)]$ with respect to different base points and, further, to describe how to transition between the various members of this combinatorial family. Our computation of the Legendrian contact dg-algebra, therefore, is not just a single dg-algebra, but a collection of dg-algebras, together with

intertwining dg-isomorphisms which compose in a functorial manner depending upon a number of combinatorial choices. In practice, the reader interested in the algebro-combinatorial content of the paper should be able to read what follows with only this minimal topological input (explicitly described in the statement and proof of Theorem 3.10 of Section 3.3).

1.1. Generalizing the Casals–Murphy dg-algebra

Associated to any non-empty trivalent plane graph G is a constant $g \geq 0$, which we call the *genus*, such that G has $|V| = 2g + 2$ vertices, $|E| = 3g + 3$ edges, and $|F| = g + 3$ faces. It is called the genus because the double cover of the sphere branched over the vertices is a surface of genus g . This is more than just a convenience in our situation; there exists a Legendrian surface Λ_G in the unit cotangent sphere bundle $S^*\mathbb{R}^3$ with its standard contact structure, as constructed by Treumann and Zaslow [28], which is a double cover of the sphere branched over the vertices. In the more modern terminology of N -graphs and Legendrian N -weaves initially studied by Casals and Zaslow [4] and greatly generalizing the Treumann–Zaslow examples, G is a 2-graph on the sphere, and Λ_G is the associated Legendrian 2-weave. For the rest of the paper, we assume that G is connected.

From our graph, one may single out a face at infinity $f_\infty \in F$ so that we may draw G as embedded in $\mathbb{R}^2$ via stereographic projection away from a point in the chosen face at infinity. We write G_{f_∞} for this trivalent plane graph embedded in $\mathbb{R}^2$, up to isotopy of embeddings (in $\mathbb{R}^2$). In the contact-geometric picture, a choice of face at infinity yields a relatively explicit model for the Legendrian satellite of Λ_G , which lives in a standard contact Darboux ball $\mathbb{R}_{\text{std}}^5$, as described, e.g., in Casals–Zaslow for the case of N -weaves which includes our case when $N = 2$ [4]. By the Legendrian satellite operation, we mean the following. We have that a Legendrian admits a contact germ depending only upon the smooth topology of the Legendrian itself. Since Λ_G is constructed in an arbitrarily small neighborhood of a Legendrian S^2 , one may satellite it along the standard Legendrian unknot $S^2 \subset \mathbb{R}_{\text{std}}^5$. We note that $H_1(\Lambda_G)$ and $\pi_1(\Lambda_G)$ still have meaning when we use the satellite, since the original surface Λ_G is diffeomorphic to its satellite, with diffeomorphism determined up to isotopy.

The algebraic invariant of Casals and Murphy [3], from this perspective, is a computation of the Legendrian contact dg-algebra of this Legendrian satellite using Ekholm’s gradient flow trees [8]. Though a proof that the combinatorics matches the geometry does not currently appear in the literature, we will provide some details of the correspondence in Section 1.4, assigning contact-geometric meaning to the combinatorial package below, providing a sketch of why the object we construct is indeed the Legendrian contact dg-algebra.

Suppose that G_{f_∞} is a trivalent plane graph with a specified face at infinity. Let $F_{\text{fin}} = F - \{f_\infty\}$ be the set of all finite faces, where we will always assume that the choice of f_∞ has been made clear from the context. Over the coefficient ring $R = \mathbb{Z}[H_1(\Lambda_G)]$, we form the free associative algebra

$$\mathcal{A}_{G, f_\infty} = R\langle F_{\text{fin}}, x, y, z \rangle \cong R\langle f_1, f_2, \dots, f_{g+2}, x, y, z \rangle,$$

where, in the last isomorphism, $f_1, \dots, f_{g+2}$ correspond to some labeling of the $g + 2$ finite faces. We turn it into a graded algebra by specifying degrees: elements of R have degree zero, elements of F_{fin} have degree 1, and $|x| = |y| = |z| = 2$. It is this graded algebra that Casals and Murphy upgrade into a dg-algebra.

Definition 1.1. A *differential graded algebra* (dg-algebra) is a pair (A, ∂) , where $A = \bigoplus_{i \in \mathbb{Z}} A_i$ is a graded algebra, and $\partial: A \rightarrow A$ is a linear map of degree -1 (so $\partial(A_i) \subset A_{i-1}$) satisfying

- the *Leibniz rule*: $\partial(xy) = (\partial x)y + (-1)^{|k|}x(\partial y)$ for $x \in A_k$ and $y \in A$,
- the *differential condition*: $\partial^2 = 0$.

A *dg-homomorphism* of dg-algebras is a morphism of graded algebras intertwining the differentials; i.e., $\phi: (A, \partial^A) \rightarrow (B, \partial^B)$ such that $\partial^B \circ \phi = \phi \circ \partial^A$. A *dg-isomorphism* is a dg-homomorphism which is an isomorphism of graded algebras.

A dg-algebra has within it the structure of a chain complex, and dg-homomorphisms are a fortiori homomorphisms of chain complexes. We may therefore recall the following notions from the theory of chain complexes.

Definition 1.2. A dg-homomorphism $\phi: (A, \partial^A) \rightarrow (B, \partial^B)$ is called a *quasi-isomorphism* if the induced morphism on homology $\phi_*: H_*(A, \partial^A) \rightarrow H_*(B, \partial^B)$ is an isomorphism. Two dg-homomorphisms $\phi, \psi: (A, \partial^A) \rightarrow (B, \partial^B)$ are said to be *(linearly/chain) homotopic* if there exists a linear map $H: (A, \partial^A) \rightarrow (B, \partial^B)$ of degree $+1$, called a *(chain) homotopy*, such that $\phi - \psi = H\partial^A + \partial^B H$. In such a case, ϕ and ψ induce the same map on homology.

On a trivalent plane graph G , there are a number of auxiliary combinatorial choices which one can make, which we package into the notion of a garden Γ ; this is Definition 4.1 below. We refer to Figure 6 for an intuitive picture: we choose an orientation of the graph, a point called a “center” in each face, arcs between all adjacent centers and vertices called “threads”, and a collection of non-intersecting coherently oriented paths which each pass through a single center, one for each face, called “tines”, and which start and end at some specified point at infinity, called a “seed”. We will come back to the contact-geometric meaning of the garden in Section 1.4. Casals and Murphy use a slightly more restrictive definition, which we call a finite-type garden. Associated to each garden (regardless of whether it is of finite type or not) is

a specified face $f_\Gamma \in F$ (the one containing the seed), which in the finite-type case we take to be the face at infinity. Casals and Murphy prove that, associated to any finite-type garden, there exists a differential which makes the resulting dg-algebra with respect to the corresponding face at infinity an invariant.

Theorem 1.3 (Casals–Murphy [3]). *For any choice of finite-type garden Γ on G , there exists a differential $\partial_{G,\Gamma}^{\mathcal{A}}$ yielding a dg-algebra $(\mathcal{A}_{G,f_\Gamma}, \partial_{G,\Gamma}^{\mathcal{A}})$, the dg-isomorphism type of which is independent of any two gardens with the same associated face at infinity.*

There are many reasonable names for this dg-algebra:

- the *Casals–Murphy dg-algebra*, in reference to the mathematicians who computed it,
- the invariant dg-algebra of binary sequences, as it was called in the paper of Casals and Murphy [3],
- the Legendrian contact dg-algebra (of the Legendrian satellite), in recognition of the contact-geometric picture,
- the Chekanov–Eliashberg dg-algebra, which is synonymous with the Legendrian contact dg-algebra.

In this paper, we will use the first of this list because the second is rather cumbersome, and the third and fourth invoke the underlying contact geometry which we will largely ignore.

There are a few versions of the Legendrian contact dg-algebra one could reasonably define, depending upon how much structure one wishes to encode. The Casals–Murphy dg-algebra corresponds to one such choice; instead, we use a fully non-commutative version, which on the level of the graded algebra is realized by removing commutativity using the following two-step procedure:

- First, we do not need that the elements of the coefficient ring commute with the other generators. We may instead consider $\mathbb{Z}$ -linear combinations of words of the form $n_1 w_1 n_2 w_2 \cdots n_k w_k$, where $n_i \in R$ and w_i is a word in the other variables.
- We may change our coefficient ring from $R = \mathbb{Z}[H_1(\Lambda_G)]$ to $R^{nc} = \mathbb{Z}[\pi_1(\Lambda_G)]$ instead, with respect to some base point. The marker nc stands for *non-commutative*.

Because of the definition of the differential of a Legendrian contact dg-algebra, in order to apply the second step, we must have also applied the first, so our use of non-commutative coefficients also necessitates removing commutativity between the coefficients and the generators. The reader interested in a contact-geometric explanation of this point may consult, e.g., Appendix A of an article of Chantraine, Dimitroglou Rizell, Ghigini, and Golovko [5].

Let us be precise about our base points. If we fix a vertex $v \in V$, then, since Λ_G is branched at v , there is a unique point $\tilde{v} \in \Lambda_G$ lying over v , which we use as a base point for our computations. Let $\mathcal{B}_{G,f_\infty,v}$ be the resulting graded algebra; we may write it as a free product of algebras

$$\mathcal{B}_{G,f_\infty,v} = \mathbb{Z}[\pi_1(\Lambda_G, \tilde{v})] * \mathbb{Z}\langle F_{\text{fin}}, x, y, z \rangle,$$

where again the elements of $\pi_1(\Lambda_G, \tilde{v})$ are in degree zero, the elements of F_{fin} are in degree 1, and $|x| = |y| = |z| = 2$.

In order to obtain nice combinatorial models for $\pi_1(\Lambda_G, \tilde{v})$ for $v \in V$, it is convenient to specify a set of generators and relations. This can be achieved by picking a tree $T \subset G$ spanning all of the vertices in V except v (recalling our assumption that G is connected so that such T exists). In fact, a choice of such a tree is already performed by Casals and Murphy for essentially the same reason; that is, the Poincaré duals to the cycles in Λ_G lying over T yield a convenient set of generators of $H_1(\Lambda_G)$. We denote the vertex missed by T as $v_T \in V$, and for the lift $\tilde{v}_T$, we use the slightly more convenient notation $*_T \in \Lambda_G$. Accordingly, we build a collection of groups Π_T with explicit generators and relations which come with canonical isomorphisms

$$\Pi_T \cong \pi_1(\Lambda_G, *_T).$$

As these models are purely combinatorial, we may therefore instead use the graded algebras

$$\mathcal{B}_{G,f_\infty,T} = \mathbb{Z}[\Pi_T] * \mathbb{Z}\langle F_{\text{fin}}, x, y, z \rangle$$

with the understanding that if T and T' miss the same vertex $v_T = v_{T'}$, then there are canonical isomorphisms

$$\mathcal{B}_{G,f_\infty,T} \cong \mathcal{B}_{G,f_\infty,v_T} = \mathcal{B}_{G,f_\infty,v_{T'}} \cong \mathcal{B}_{G,f_\infty,T'}.$$

There is good reason for having these graded algebras depend upon T (and not just $*_T$) aside from just that they are combinatorially convenient. Unlike in the version with homology coefficients, in order to obtain the Legendrian contact differential, each gradient flow tree counted in the differential must be connected to the base point $*$ in order to obtain an element in $\pi_1(\Lambda_G, *)$. In contact-geometric parlance, such choices are referred to as capping paths. Conveniently, our trees T spanning the vertices $V - \{v_T\}$ yield natural choices of capping paths and hence fix the differentials. However, even if $v_T = v_{T'}$, the capping paths will not be the same, and hence, a dg-isomorphism between the models will not in general be realized by the natural canonical isomorphisms $\mathcal{B}_{G,f_\infty,T} \cong \mathcal{B}_{G,f_\infty,T'}$ of the previous paragraph, since this will not in general intertwine the differentials. Nonetheless, we have the following, the first main theorem of this paper.

Theorem 1.4. *Let G be a trivalent plane graph. For any choice of finite-type garden Γ on G and tree $T \subset G$ spanning all but one vertex, there exists a differential $\partial_{G,\Gamma,T}^{\mathcal{B}}$ on $\mathcal{B}_{G,f_\Gamma,T}$, yielding a dg-algebra*

$$(\mathcal{B}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T}^{\mathcal{B}}),$$

which, under the map given by abelianizing coefficients and commuting ring elements with generators, induces a dg-homomorphism

$$\text{Ab}: (\mathcal{B}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T}^{\mathcal{B}}) \rightarrow (\mathcal{A}_{G,f_\Gamma}, \partial_{G,\Gamma}^{\mathcal{A}}).$$

If Γ and Γ' are two finite-type gardens and T and T' are trees in G spanning all but one vertex, then the resulting dg-algebras for the pairs (Γ, T) and (Γ', T') are dg-isomorphic.

Theorem 1.3 follows as an immediate corollary from Theorem 1.4.

1.2. Functoriality for the non-commutative Casals–Murphy dg-algebra

We may further refine the dg-isomorphisms of Theorem 1.4, as in Theorem 1.6 below. The refined version provides an understanding of two further questions:

- (1) What properties do the dg-isomorphisms between the dg-algebras associated to (Γ, T) and (Γ', T') satisfy?
- (2) What combinatorial data is required to specify a *canonical* dg-isomorphism for different pairs (Γ, T) and (Γ', T') ?

Let us begin with the first of these questions. In the setting of commutative coefficients, it is well known that the Legendrian contact dg-algebra is invariant up to a notion known as stable tame (dg-)isomorphism; cf. [9]. In our setting, one can easily deduce that this is the case for the commutative coefficient setting by carefully studying the arguments of Casals and Murphy [3]. It was noted previously that there should be a notion of stable tame (dg-)isomorphism in the non-commutative setting [5], though formal definitions did not appear. We make those definitions under the assumption of non-negative grading, which will be the case throughout this paper (all of our generators will have degree 1 or 2).

Definition 1.5. Let R be any unital ring, S any set, and $\deg: S \rightarrow \mathbb{Z}_{\geq 0}$ a collection of (non-negative) degrees for elements of S . Let

$$A_S = R * \mathbb{Z}\langle S \rangle$$

be the fully non-commutative associative graded algebra freely generated by R and elements of S , where the grading of elements of R is 0 and the grading of elements

of S matches the degree. The algebra A_S together with its explicit set of generators S is known as a *semi-free graded algebra*. If A_S comes with a differential ∂ (of degree -1), then we call (A_S, ∂) a *semi-free dg-algebra*. Note that $\partial r = 0$ for all $r \in R$ for degree reasons.

An *elementary automorphism* of a semi-free graded algebra A_S is an automorphism $\Phi: A_S \rightarrow A_S$, which is the identity on elements of R and $S \setminus \{s\}$ for some $s \in S$, and satisfies

$$\Phi(s) = s + \alpha$$

for some $\alpha \in A_{S \setminus \{s\}}$. A *tame automorphism* of A_S is a composition of elementary automorphisms.

A *regeneration* of A_S is a graded algebra automorphism of the form $\Phi(s) = u_s \cdot \sigma(s) \cdot v_s$ on all generators $s \in S$, where $u_s, v_s \in R$ are invertible elements and σ is a permutation on S .

If (A_S, ∂) is a semi-free dg-algebra, then any graded algebra automorphism, say, $\Phi: A_S \rightarrow A_S$, induces a unique differential ∂' on A_S so that $\Phi: (A_S, \partial) \rightarrow (A_S, \partial')$ is a dg-isomorphism. If ∂ is given, we will often conflate Φ as a graded algebra automorphism with Φ as a dg-isomorphism. (For example, we may call a tame automorphism a tame dg-isomorphism.)

The *stabilization* of a semi-free dg-algebra (A_S, ∂) is the semi-free dg-algebra (A_{S^+}, ∂^+) , where $S^+ = S \sqcup \{x, y\}$ with $\deg(y) = \deg(x) + 1$, ∂^+ matching ∂ on S , and

$$\partial^+ y = x, \quad \partial^+ x = 0.$$

The natural inclusion $A_S \hookrightarrow A_{S^+}$ is a dg-homomorphism, called the *stabilization (dg-homomorphism)*. Conversely, A_S is called the *destabilization* of A_{S^+} (with respect to $x, y \in S$), and the dg-homomorphism $A_{S^+} \rightarrow A_S$ which kills x and y but preserves the other generators and elements of R is called the *destabilization (dg-homomorphism)*.

A *stable regenerative tame dg-homomorphism* is a composition of stabilizations, destabilizations, regenerations, and elementary dg-automorphisms. Any of the adjectives stable, regenerative, or tame may be removed one at a time if there are no stabilizations or destabilizations, no regenerations, or no tame dg-homomorphisms. We may similarly replace the word dg-homomorphism with dg-automorphism, where appropriate.

It is a fact that every stable regenerative tame dg-homomorphism is a quasi-isomorphism, as is proved in Proposition 2.4.

Because we allow for stabilizations and destabilizations, we work in a slightly more symmetric setup. Whereas the finite-type gardens considered by Casals and Murphy avoid considering the face at infinity, we may stabilize, including the face

at infinity as an extra degree 1 generator at the cost of also introducing a degree 2 generator w which kills it (up to an elementary dg-homomorphism). We will use the superscript $+$ for this enlargement. For any garden, not just those of finite type, we form the graded algebra

$$\mathcal{B}_{G,T}^+ := \mathbb{Z}[\Pi_T] * \mathbb{Z}\langle F, x, y, z, w \rangle.$$

This allows us to treat all of the faces together in a slightly cleaner manner. As per Proposition 2.4, when we put the proper differential on this graded algebra, we obtain a model quasi-isomorphic to the version without the $+$ enlargement. Contact geometrically, these two extra generators can be produced by a specific Legendrian isotopy of the satellite of Λ_G , and so, it is no surprise that we obtain a quasi-isomorphic model.

The second question is partly motivated by the fact that the constructed dg-isomorphisms seem to require a number of choices, even in the commutative coefficient setting of Casals and Murphy [3] (though they leave this point implicit in their constructions). We work to explain the extra underlying algebraic structure.

One piece of the story is relatively simple. A garden comes with a choice of orientation of the edges of G . If E are the edges of G , then there is always a natural $(\mathbb{Z}_2)^E$ action on gardens given by flipping edge orientations. For two gardens which are the same except for the orientation of edges on G , the underlying graded algebras are the same, and we will see that the dg-isomorphisms are induced by ring automorphisms of $\mathbb{Z}[\Pi_T]$. Other changes of the garden are more complicated.

Aside from changing the garden, we may also change the tree. We have already pointed out that even if $*_T = *_{T'}$, then, for fixed Γ , the canonical identification

$$\mathcal{B}_{G,f_\Gamma,T} \cong \mathcal{B}_{G,f_\Gamma,T'}$$

given by the association

$$\Pi_T \cong \pi_1(\Lambda_G, *_T) = \pi_1(\Lambda_G, *_{T'}) \cong \Pi_{T'}$$

will not induce a dg-isomorphism because the induced capping paths are not the same. More generally, if

$$*_T \neq *_{T'},$$

then to even compare $\pi_1(\Lambda_G, *_T)$ and $\pi_1(\Lambda_G, *_{T'})$ in the first place we must choose a non-canonical homotopy class of path γ from $*_T$ to $*_{T'}$, yielding the isomorphism $\mathcal{C}(\gamma): \pi_1(\Lambda_G, *_T) \rightarrow \pi_1(\Lambda_G, *_{T'})$ defined by

$$\mathcal{C}(\gamma)(\eta) := \gamma^{-1} * \eta * \gamma.$$

We recall here that the composition $*$ means the following: if η_1 and η_2 are two homotopy classes of paths from p to q and q to r , respectively, then $\eta_1 * \eta_2$ is the homotopy

class of paths from p to r with representative given by first following a representative of η_1 and then following a representative of η_2 . Finally, let $\mathcal{C}_T^{T'}(\gamma): \Pi_T \rightarrow \Pi_{T'}$ be the induced isomorphism fitting into the commutative diagram

$$\begin{array}{ccc} \Pi_T & \xrightarrow{\mathcal{C}_T^{T'}(\gamma)} & \Pi_{T'} \\ \wr \downarrow & & \downarrow \wr \\ \pi_1(\Lambda_G, *T) & \xrightarrow{\mathcal{C}(\gamma)} & \pi_1(\Lambda_G, *T') \end{array}$$

which in turn induces graded algebra isomorphisms $\mathcal{C}_T^{T'}(\gamma): \mathcal{B}_{G,T}^+ \rightarrow \mathcal{B}_{G,T'}^+$.

Theorem 1.6. *Let G be a non-empty connected trivalent plane graph of genus g . For any garden Γ and tree T , there exists a canonical differential $\partial_{G,\Gamma,T}^{\mathcal{B}^+}$ on $\mathcal{B}_{G,T}^+$ yielding a dg-algebra. If two gardens Γ and Γ' on G are homotopic and $T = T'$, then their associated dg-algebras are canonically identified via the identity on the underlying graded algebras.*

The group

$$\mathcal{H}(G) := (\mathbb{Z}_2)^E \times F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$$

naturally acts transitively on the set of homotopy classes of gardens on G , where E is the set of edges of G , and where F_{g+2} is the free group on $g + 2$ generators. Suppose that $\zeta \in \mathcal{H}(G)$ with

$$\Gamma' = \zeta \cdot \Gamma,$$

and suppose that γ is a homotopy class of paths from $*T$ to $*T'$ in Λ_G . Then, there exist canonical dg-isomorphisms

$$\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma): (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}) \rightarrow (\mathcal{B}_{G,T'}^+, \partial_{G,\Gamma',T'}^{\mathcal{B}^+})$$

which satisfy the following functoriality properties:

- *Composition property: the dg-isomorphisms compose naturally in the sense that*

$$\Phi_{\Gamma',T'}^{\Gamma'',T''}(\zeta, \gamma_2) \circ \Phi_{\Gamma,T}^{\Gamma',T'}(\theta, \gamma_1) = \Phi_{\Gamma,T}^{\Gamma'',T''}(\zeta\theta, \gamma_1 * \gamma_2).$$

- *Orientation changes: the action of the factor $(\mathbb{Z}_2)^E$ (with $\gamma = \mathbb{1}$) is induced by ring automorphisms of $\mathbb{Z}[\Pi_T]$.*
- *Fixing the trees and capping paths; modifying the garden: if $T = T'$ and $\gamma = \mathbb{1}$ is the constant path at $*T$, then, for any $\zeta \in F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$, the map $\Phi_{\Gamma,T}^{\Gamma',T}(\zeta, \mathbb{1})$ is a regenerative tame dg-isomorphism. Furthermore, if ζ stabilizes Γ , i.e., $\zeta \cdot \Gamma = \Gamma$, then $\Phi_{\Gamma,T}^{\Gamma,T}(\zeta, \mathbb{1})$ is homotopic to the identity and hence acts by the identity in homology.*

- *Fixing the garden; modifying the trees and capping paths:* for any fixed garden Γ , any two trees T and T' , and any homotopy class of path γ from $*_T$ to $*_{T'}$, we have

$$\Phi_{\Gamma,T}^{\Gamma,T'}(1, \gamma) = \mathcal{R}_T^{T'}(\gamma) \circ \mathcal{C}_T^{T'}(\gamma),$$

where $\mathcal{C}_T^{T'}(\gamma)$ acts on the coefficient ring (as discussed just before this theorem statement) and $\mathcal{R}_T^{T'}(\gamma)$ is a regeneration independent of the garden acting on each generator by conjugation by an element of $\Pi_{T'}$ (not necessarily the same for each generator). If $T = T'$, then $\mathcal{R}_T^T(\mathbb{1}) = \mathbb{1}$.

- *Finite-type gardens:* if Γ is a finite-type garden, then the dg-algebra of Theorem 1.4, $(\mathcal{B}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T}^{\mathcal{B}})$, is obtained from $(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+)$ via a canonical composition of an elementary automorphism of $\mathcal{B}_{G,T}^+$ followed by a destabilization.

Notice that we obtain significantly more functorial data than was originally obtained in the commutative coefficient setting by Casals and Murphy [3]. By functoriality, we mean that we may build for each connected non-empty trivalent graph G a category whose objects are pairs (Γ, T) and whose morphisms are pairs (ζ, γ) ; Theorem 1.6 yields a functor from this category to the category of dg-algebras and dg-isomorphisms.

Any specified faces $f, f' \in F$ may arise as $f_\Gamma, f_{\Gamma'}$ for some finite-type gardens Γ, Γ' , and accordingly, for any $\zeta \in \mathcal{H}(G)$ with $\zeta \cdot \Gamma = \Gamma'$ and homotopy class of path γ from $*_T$ to $*_{T'}$, we obtain a canonical quasi-isomorphism from $(\mathcal{B}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T})$ to $(\mathcal{B}_{G,f_{\Gamma'},T'}, \partial_{G,\Gamma',T'})$ given by the composition

$$\mathcal{B}_{G,f_\Gamma,T} \xrightarrow{\text{stab}} \mathcal{B}_{G,T}^+ \xrightarrow{\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta,\gamma)} \mathcal{B}_{G,T'}^+ \xrightarrow{\text{destab}} \mathcal{B}_{G,f_{\Gamma'},T'},$$

where the stabilization and destabilization maps include an extra elementary automorphism of $\mathcal{B}_{G,T}^+$ and $\mathcal{B}_{G,T'}^+$, respectively. Theorem 1.4 follows from Theorem 1.6 so long as one proves that the above maps are dg-isomorphisms. Partial abelianization then yields all functoriality properties satisfied by the original Casals–Murphy dg-algebras $(\mathcal{A}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T}^{\mathcal{A}})$.

Our construction of the differentials occurring in Theorem 1.6 borrows from an idea of Casals and Murphy [3]. Instead of working with the coefficient ring $\mathbb{Z}[\Pi_T]$, we may work with instead a coefficient ring $\mathbb{Z}[\Pi]$, where Π is a group coming with projections $\pi_T: \Pi \rightarrow \Pi_T$ and inclusions $i_T: \Pi_T \hookrightarrow \Pi$ with

$$\pi_T \circ i_T = \mathbb{1}.$$

In other words, we separate the dependence on the tree T , forming the dg-algebra

$$\tilde{\mathcal{B}}_G^+ := \mathbb{Z}[\Pi] * \mathbb{Z}\langle F, x, y, z, w \rangle$$

and accordingly finding differentials $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ for each garden Γ . (One may of course work without the $+$ enlargement at the cost of restricting to finite-type gardens.) The differentials over this enlarged ring then induce the differentials $\partial_{G,\Gamma,T}^{\mathcal{B}^+}$ of Theorem 1.6, and the functoriality properties follow from a careful consideration of the algebraic structure of Π with its various projections π_T . The functoriality properties for the dg-algebras $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})$ are given by Theorem 5.1.

1.3. Representations and colors

If R is a ring, which we consider as a graded $\mathbb{Z}$ -algebra concentrated in degree zero, then we obtain a dg-algebra $(R, 0)$.

Definition 1.7. An R -augmentation of a dg-algebra (A, ∂) is a dg-homomorphism

$$\varepsilon: (A, \partial) \rightarrow (R, 0).$$

In the case that $R = \text{Mat}_r(\mathbb{F})$ for a field $\mathbb{F}$, we call the corresponding R -augmentation a *rank r representation over $\mathbb{F}$* . A rank 1 representation is sometimes just called an *augmentation*. We denote the set of representations of rank r by

$$\text{Rep}_r(A, \partial; \mathbb{F}) := \{\text{rank } r \text{ representations of } (A, \partial) \text{ over } \mathbb{F}\}.$$

Notice that there is a conjugation action of $\text{GL}_r(\mathbb{F})$ on $\text{Rep}_r(A, \partial; \mathbb{F})$: if ε is a rank r representation and $g \in \text{GL}_r(\mathbb{F})$, then

$$(g \cdot \varepsilon)(x) := g \cdot \varepsilon(x) \cdot g^{-1}.$$

We will prefer to consider two rank r representations to be equivalent if they are related by the conjugation action.

Definition 1.8. The *moduli space of rank r representations over $\mathbb{F}$* of a dg-algebra (A, ∂) is just the set of representations modulo equivalence by this conjugation action, denoted by

$$\mathcal{M}_r^{\text{Rep}}(A, \partial; \mathbb{F}) := \text{Rep}_r(A, \partial; \mathbb{F}) / \text{GL}_r(\mathbb{F}).$$

The conjugation $\text{GL}_r(\mathbb{F})$ -action descends to a $\text{PGL}_r(\mathbb{F})$ -action, so we could alternatively have written the quotient with respect to $\text{PGL}_r(\mathbb{F})$ instead. In the case $r = 1$, since $\text{PGL}_1(\mathbb{F})$ is trivial, we therefore have

$$\mathcal{M}_1^{\text{Rep}}(A, \partial; \mathbb{F}) = \text{Rep}_1(A, \partial; \mathbb{F}).$$

For higher r , there is in general no such equality except for highly degenerate situations.

There is a general correspondence between the moduli space of rank r representations of a Legendrian contact dg-algebra and a certain moduli space of constructible sheaves of microlocal rank r depending upon the Legendrian.² As such, since we have stated throughout this introduction that the dg-algebras we are constructing are combinatorial models for the Legendrian contact dg-algebra associated to Λ_G , in order to verify this correspondence, we need also to understand the corresponding moduli space of constructible sheaves. See the work of Shende, Treumann, Williams, and Zaslow [27] for a precise definition of the desired constructible sheaves of microlocal rank r and Treumann and Zaslow [28] for a description in the case $r = 1$ for the case we are considering. From these references, one may easily deduce a clean combinatorial description of this moduli space of sheaves purely in terms of the trivalent plane graph G , which we now describe.

Suppose that we are given a trivalent plane graph G , and let F be the faces of G . Consider the Grassmannian $\text{Gr}(r, 2r; \mathbb{F})$ consisting of subspaces of $\mathbb{F}^{2r}$ of dimension r .

Definition 1.9. A rank r face coloring of G over $\mathbb{F}$ is an assignment

$$\Phi: F \rightarrow \text{Gr}(r, 2r; \mathbb{F})$$

such that if f and g share an edge, then $\Phi(f)$ and $\Phi(g)$ are transverse. We denote the set of such rank r face colorings by

$$\text{Col}_r(G; \mathbb{F}) := \{\text{rank } r \text{ face colorings of } G \text{ over } \mathbb{F}\}.$$

In the case that $r = 1$, this recovers the usual notion of a face coloring of the graph G (or equivalently a vertex coloring of the dual graph G^*) by elements of the projective line $\mathbb{F}\mathbb{P}^1$. We note that in the case that

$$\mathbb{F} = \mathbb{F}_q$$

is a finite field of order q , then the number of face colorings of rank 1 is therefore

$$|\text{Col}_1(G; \mathbb{F}_q)| = \chi_{G^*}(q + 1),$$

where χ_{G^*} is the famous chromatic polynomial associated to a trivalent plane graph. Here, we use the fact that $|\mathbb{F}_q\mathbb{P}^1| = q + 1$.

Notice that $\text{PGL}_{2r}(\mathbb{F})$ acts on $\text{Gr}(r, 2r; \mathbb{F})$, and if $\Phi \in \text{PGL}_{2r}(\mathbb{F})$ and we have $X, Y \in \text{Gr}(r, 2r; \mathbb{F})$ being transverse, then $\Phi(X)$ and $\Phi(Y)$ are also transverse. Therefore, $\text{Col}_r(G; \mathbb{F})$ is an invariant set for the $\text{PGL}_{2r}(\mathbb{F})$ -action.

²In our case, representations are homotopic if and only if they are equal for degree reasons; the more general conjecture requires taking representations modulo homotopies.

Definition 1.10. The *moduli space of rank r colorings over $\mathbb{F}$* is the set given by

$$\mathcal{M}_r^{\text{Col}}(G; \mathbb{F}) := \text{Col}_r(G; \mathbb{F}) / \text{PGL}_{2r}(\mathbb{F}).$$

In the case that $r = 1$, so long as the trivalent plane graph G is non-empty, the action of $\text{PGL}_2(\mathbb{F})$ is free. In particular, in the case that $\mathbb{F} = \mathbb{F}_q$ has order q , we have

$$|\mathcal{M}_1^{\text{Col}}(G; \mathbb{F}_q)| = \frac{\chi_{G^*}(q+1)}{q^3 - q}.$$

This description of $\mathcal{M}_1^{\text{Col}}(G; \mathbb{F}_q)$ appeared in the work of Treumann and Zaslow [28]. In the appendix to the paper of Casals and Murphy [3], written by the present author, the “augmentations are sheaves” result was explicitly verified.

Theorem 1.11 (See [3, Appendix A]). *For any field $\mathbb{F}$, there is a natural bijection*

$$\mathcal{M}_1^{\text{Rep}}(\mathcal{A}_{G, f_{\Gamma}, T}, \partial_{G, \Gamma, T}^{\mathcal{A}}; \mathbb{F}) \cong \mathcal{M}_1^{\text{Col}}(G; \mathbb{F}).$$

It is worth noting that in the case that $\mathbb{F} = \mathbb{F}_3$ is the field of order 3, the four-color theorem, famous for its eventual solution by Appel and Haken [1] in the 1970s using computers, as well as the current non-existence of a non-computer-based proof, is the assertion that $\mathcal{M}_1^{\text{Col}}(G; \mathbb{F}_3)$ is non-empty whenever G is bridgeless, meaning there is no face which borders itself. Therefore, we obtain the equivalent result that, for bridgeless G , the Legendrian Λ_G admits an $\mathbb{F}_3$ -augmentation.

Because $\text{GL}_1(\mathbb{F}) = \mathbb{F}^*$ is commutative, we notice that there is a canonical identification

$$\mathcal{M}_1^{\text{Rep}}(\mathcal{A}_{G, f_{\Gamma}, T}, \partial_{G, \Gamma, T}^{\mathcal{A}}; \mathbb{F}) \cong \mathcal{M}_1^{\text{Rep}}(\mathcal{B}_{G, f_{\Gamma}, T}, \partial_{G, \Gamma, T}^{\mathcal{B}}; \mathbb{F}).$$

In particular, the previous theorem is a special case of “representations are sheaves” for the non-commutative setting, which we verify along with a number of functoriality properties.

With regards to functoriality, notice that for a dg-homomorphism $\Phi: (A, \partial) \rightarrow (A', \partial')$, the pullback maps

$$\Phi^*: \text{Rep}_r(A', \partial'; \mathbb{F}) \rightarrow \text{Rep}_r(A, \partial; \mathbb{F})$$

intertwine the conjugation action. Hence, Φ^* descends to the moduli spaces

$$\Phi^*: \mathcal{M}_r^{\text{Rep}}(A', \partial'; \mathbb{F}) \rightarrow \mathcal{M}_r^{\text{Rep}}(A, \partial; \mathbb{F}).$$

In our situation, $\Phi_{\Gamma, T}^{\Gamma', T'}(\zeta, \gamma)$ are all dg-isomorphisms yielding pullback isomorphisms

$$(\Phi_{\Gamma, T}^{\Gamma', T'}(\zeta, \gamma))^*: \text{Rep}_r(\mathcal{B}_{G, T'}^+, \partial_{G, \Gamma', T'}^{\mathcal{B}^+}; \mathbb{F}) \rightarrow \text{Rep}_r(\mathcal{B}_{G, T}^+, \partial_{G, \Gamma, T}^{\mathcal{B}^+}; \mathbb{F})$$

which intertwine the conjugation action of $\text{GL}_r(\mathbb{F})$ and hence also induce isomorphisms at the level of the moduli space of rank r representations.

Theorem 1.12. *Let G be a trivalent plane graph, $\mathbb{F}$ a field, and r a positive integer. For each garden Γ and subtree $T \subset G$ spanning all but one vertex, there is a canonical map*

$$\Psi_{G,\Gamma,T}: \text{Rep}_r(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}; \mathbb{F}) \rightarrow \mathcal{M}_r^{\text{Col}}(G; \mathbb{F})$$

with the following properties:

- *Functoriality: we have*

$$\Psi_{G,\Gamma,T} \circ (\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma))^* = \Psi_{G,\Gamma',T'}.$$

- *Representations are sheaves: $\Psi_{G,\Gamma,T}$ is invariant under the conjugation action and induces a bijection*

$$\mathcal{M}_r^{\text{Rep}}(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}; \mathbb{F}) \cong \mathcal{M}_r^{\text{Col}}(G; \mathbb{F}).$$

If one wishes to work with finite-type gardens, as was done in the original commutative setting, one simply notices that, so long as all generators are in positive degree, the pullbacks of stabilization and destabilization maps are inverses on the sets of representations; see Proposition 2.6 below. It follows immediately that

$$\text{Rep}_r(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}; \mathbb{F}) \cong \text{Rep}_r(\mathcal{B}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T}^{\mathcal{B}}; \mathbb{F})$$

since the dg-algebras are related by a composition of (elementary) dg-isomorphisms together with stabilization/destabilization, where all generators involved are of positive degree. Furthermore, since pullbacks of dg-homomorphisms intertwine the conjugation action, we furthermore have

$$\mathcal{M}_r^{\text{Rep}}(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}; \mathbb{F}) \cong \mathcal{M}_r^{\text{Rep}}(\mathcal{B}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T}^{\mathcal{B}}; \mathbb{F}).$$

In the case $r = 1$, one recovers Theorem 1.11 from combining this observation with Theorem 1.12.

1.4. Contact geometry

We now give a few details as to why one expects that the dg-algebras $(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+})$ are indeed models for a Legendrian contact dg-algebra. To begin, we will explain why $\mathcal{B}_{G,T}^+$ is the correct graded algebra, which requires identifying the generators as the Reeb chords of the Legendrian. We will then partially discuss the differential, including a description of the types of objects which are counted. Upon reading Section 4.2, the contact-geometric reader may wish to come back to this section to check their understanding of how each binary sequence (Definition 4.3) corresponds to a term in the differential, which essentially finishes the translation of the rest of the article back into the language of contact geometry. A full proof of this correspondence




Figure 1. We have a sketch of the front of the Legendrian satellite of Λ_G . We have not drawn the details of the flat portion for simplicity, though we sketch in green one of the Reeb chords corresponding to a center of the garden inside one of the faces of the graph G . We have also drawn the four Reeb chords corresponding to the generators w, x, y, z .

would require also a description of the signs, which we do not discuss at all, as well as an understanding of gradient flow trees near the trivalent vertices, which requires a perturbation of certain singularities in the front.

Recall that $\mathbb{R}_{\text{std}}^5$ comes with a so-called front projection $\pi_f: \mathbb{R}^5 \rightarrow \mathbb{R}^3$, and that for a generic Legendrian surface $\Lambda \subset \mathbb{R}^5$, it may be recovered from its projection to $\pi_f(\Lambda) \subset \mathbb{R}^3$, which is almost everywhere immersed except along some mild singular strata. More specifically, with respect to the typical coordinates x_1, x_2, z for $\mathbb{R}^3$, this projection is never tangent to the vertical vector ∂_z , and the remaining two coordinates are determined by the slopes of the tangent planes to $\pi_f(\Lambda)$, i.e., $y_j = \frac{\partial z}{\partial x_j}$. The underlying graded algebra of the Legendrian contact dg-algebra is then of the form

$$\mathbb{Z}[\pi_1(\Lambda)] * \mathbb{Z}\langle \mathcal{R} \rangle,$$

where $\mathcal{R}$ is the collection of Reeb chords. From the perspective of the front projection, these consist of pairs of points $(p, q) \in \pi_f(\Lambda)$ such that q lies directly above p in the z -direction, and $T_p(\pi_f(\Lambda))$ and $T_q(\pi_f(\Lambda))$ are parallel so that the y_j -coordinates of the corresponding points on Λ are equal.

In order to verify that we have the correct underlying graded algebra, we need an explicit model where we identify $\mathcal{R} \cong F \sqcup \{w, x, y, z\}$. A front projection for the Legendrian satellite of N -weaves appears in the work of Casals and Zaslow [4], though a little more care is needed to identify the Reeb chords. Assuming that we have chosen a face at infinity f_∞ , we take the front as follows; see Figure 1 for reference. First, we take an immersion of the double-branched cover of a large disk containing G_{f_∞} , branched at the vertices of G , and where the projection map $\mathbb{R}_{x_1, x_2, z}^3 \rightarrow \mathbb{R}_{x_1, x_2}^2$ recovers the branching map, and where the immersion crosses itself along the edges and is given by a so-called D_4^- singularity at vertices. We take this branched cover

to be relatively flat so that the corresponding portion of the Legendrian in $\mathbb{R}_{\text{std}}^5$ has relatively small y_1 and y_2 coordinates, and accordingly, we call this the flat portion of our model. The boundary of the flat portion consists of two circles. We then glue in two more disk pieces, sewn together along the two boundary circles along cusp edges so that the projection lifts to a smooth Legendrian. We take these two disks to be C^1 -close to each other but do not intersect each other and to be cambered upwards so that they are bell shaped, rising mostly well above the flat portion, except that the bottom disk will pass the top sheet of the flat portion near the seam. If the lower disk is the graph of a bell-shaped function $f: \mathbb{R}_{x_1, x_2}^2 \rightarrow \mathbb{R}_z$, then we may take the upper disk to be given by the graph of the nearby bell-shaped function $f + \varepsilon + \varepsilon^2 x_1$ for small enough $\varepsilon > 0$.

Let us now find the Reeb chords. By construction, we see that there are no Reeb chords between the two cambered disks. For the flat portion, we find that, for each finite face, there is a unique Reeb chord between the two sheets, given at the point where the two sheets are of maximal distance away. Because of the positioning of the seams according to the functions f and $f + \varepsilon + \varepsilon^2 x_1$, we will also find that there must be a Reeb chord appearing for the face at infinity, since the two sheets are not at maximal distance along the boundary. Finally, between the flat and cambered portions, we can choose f so that there are no Reeb chords near the seams, and so, the only other Reeb chords which occur are near the maximum of the function f , where we obtain $4 = 2 \times 2$ more Reeb chords, given by picking one of the two cambered disks and one of the two sheets below. The seed of the garden is meant to codify the approximate position of these Reeb chords.

The grading of the generators in $\mathcal{R}$ may be computed using a formula of Ekholm, Etnyre, and Sullivan [10, Lemma 3.4]. Namely, suppose that $r \in \mathcal{R}$ is a Reeb chord from r_- to r_+ , lying on sheets in the front projection which are locally the graphs of functions f_- and f_+ . Pick a generic path γ on Λ from r_+ to r_- . Let $U(\gamma)$ and $D(\gamma)$ be the number of times γ crosses the cusp edges (where the cambered disks are sewn to the flat sheets) upwards or downwards, respectively. Let $\text{Morse}(r)$ be the Morse index of the critical point of $f_+ - f_-$ at the point over which the Reeb chord r lies. Then, the degree of r is

$$|r| = D(\gamma) - U(\gamma) + \text{Morse}(r) - 1.$$

In our case, we see that $\text{Morse}(r) = 2$ for all of the Reeb chords involved. For those Reeb chords in F , we may take γ to stay within the flat portion; hence,

$$D(\gamma) = U(\gamma) = 0,$$

and hence, $|f| = 1$. For the Reeb chords w, x, y, z , we see that we can take a curve γ which crosses one cusp downwards so that $D(\gamma) = 1$ and $U(\gamma) = 0$, from which we find that $|w| = |x| = |y| = |z| = 2$. We have therefore recreated the graded algebra.

As for the differential, we appeal to the gradient flow tree technology of Ekholm [8]. Each sheet of $\pi_f(\Lambda)$ may be thought of as the graph of a function

$$\mathbb{R}_{x_1, x_2}^2 \rightarrow \mathbb{R}_z.$$

The gradient flow trees counted by Ekholm may be thought of as certain trees in $\mathbb{R}_{x_1, x_2}^2$ such that each edge of the tree is labeled by two sheets, say, sheets i and j given by graphs of functions $f_i < f_j$, and where the edges are tangent to and oriented with $-\nabla(f_j - f_i)$. The vertices of the tree are of only a handful of specified types. For example, there are Y_0 -vertices, involving three sheets which are the graphs of functions

$$f_i < f_j < f_k.$$

Near such a vertex, the tree has one input edge directed as $-\nabla(f_k - f_i)$ and two output edges directed as $-\nabla(f_j - f_i)$ and $-\nabla(f_k - f_j)$. The differential of a Reeb chord corresponds to counts of gradient flow trees which emanate from the Reeb chord, i.e., such that the tree has a positive 1-valent puncture, meaning a node corresponding to the Reeb chord and with an outgoing edge oriented along $-\nabla(f_+ - f_-)$ for the functions $f_- < f_+$ whose graphs are the sheets containing the endpoints of the Reeb chord. Many such flow trees exist, but only a finite number of them have the property that they cannot be deformed. It is these rigid flow trees which are counted in the differential.

For a Reeb chord corresponding to a face, there is a single gradient flow line to each of the vertices, and upon perturbing the front near the vertex (since the D_4^- singularity is not allowed in Ekholm's technology), one finds that each of these contributes precisely one term to the differential (though this requires additional proof via an understanding of the perturbation of the D_4^- singularity, which we do not provide). These flow lines are codified as the threads in our garden, and they are easily seen to be the only terms which can possibly contribute to the differential of the Reeb chords at the faces.

Meanwhile, for each of the four Reeb chords between the cambered and flat portions, because the slopes of the cambered portion are C^1 close and the slopes in the flat portion are C^1 close and small, the gradient trajectories emanating from these Reeb chords are all essentially just radially outward from where these Reeb chords are positioned, approximately following the trajectory of the vector field $-\nabla f$ where f is our bell-shaped function. The rigid flow trees starting at one of these four generators must pass over the face Reeb chords for degree reasons (as a 2-valent negative puncture in Ekholm's language), and there is one such trajectory of $-\nabla f$ for each face. These trajectories are codified combinatorially by the tines of the garden. But this is not the whole picture; for example, certain Y_0 vertices are allowed where the tines cross the threads. Indeed, suppose that we have a gradient flow line along the




Figure 2. A sample rigid gradient flow tree which contributes to the differential of the degree 2 generator y . Each edge of the tree is either a thread or a tine. In Ekholm’s language, the black dots are Y_0 -vertices, the green dot contained in the tree is a 2-valent (negative) puncture, the arrowheads occur at ends (after perturbing the D_4^- singularities), and the yellow point is a 1-valent (positive) puncture, which is just the Reeb chord y in this case. Each edge of the tree may be labeled with the corresponding sheets for which it is following the negative gradient flow. Along the threads, the edge is labeled by the two flat sheets. Along the tine, each edge is labeled by one of the flat sheets, either 0 or 1 for the lower or upper flat sheet, and by the lower cambered sheet.

tine, where the two sheets involved are the lower flat edge and one of the cambered edges. Then, where the tine intersects the thread, this trajectory may split: one edge will still go along the tine, but following the upper flat edge (and the same cambered edge) instead, and the other edge will go along a thread until it ends at a D_4^- vertex as we saw when looking at the differential of the chords at faces.

It is not hard to then enumerate all possible rigid gradient flow trees; essentially, they are encoded by whether or not there is a Y_0 -vertex where the tine crosses a thread. Figure 2 demonstrates an example of a rigid flow tree which counts towards the differential of the generator y , which corresponds to the Reeb chord from the top sheet of the flat portion to the lower disk of the cambered portion. The figure also has labels 0, 1 corresponding to the fact that, for each edge of the tree which is along the tine, one of the corresponding sheets is either the upper or lower of the flat sheets. Combinatorially, this data is encoded in the notion of a binary sequence, as in Definition 4.3, with each binary sequence giving precisely one term of the differential.

Verifying that this gives all of the terms in the differential is not difficult, assuming again that we have understood how threads terminate at a perturbation of the D_4^- singularities.

2. Preliminaries on stable regenerative tame dg-homomorphisms

In this short section, we prove a few algebraic preliminaries mentioned in the introduction about stable regenerative tame dg-homomorphisms. These proofs are somewhat technical, with techniques orthogonal to the main ideas of the paper, so we encourage the impatient reader to skip this section in a first read-through. They are presented for completeness.

We recall the following notions which may be found in the commutative coefficient setting in the work of Kálmán [20], extended here to non-commutative coefficients.

Definition 2.1. Suppose that (A, ∂^A) and (B, ∂^B) are dg-algebras and $\phi, \psi: (A, \partial^A) \rightarrow (B, \partial^B)$ are dg-homomorphisms. A (ϕ, ψ) -derivation is a $\mathbb{Z}$ -linear map of degree +1, $K: (A, \partial^A) \rightarrow (B, \partial^B)$, such that if $x \in A_k$ and $y \in A$, then

$$K(xy) = K(x)\psi(y) + (-1)^k \phi(x)K(y).$$

Lemma 2.2 (Cf. [20, Lemma 2.18]). *Suppose that $(A = R * \mathbb{Z}\langle S \rangle, \partial^A)$ is a semi-free dg-algebra, (B, ∂^B) is a dg-algebra, and we have dg-homomorphisms $\phi, \psi: (A, \partial^A) \rightarrow (B, \partial^B)$. If $K: S \rightarrow B$ is any map of degree +1, then there exists a unique (ϕ, ψ) -derivation $\tilde{K}: (A, \partial^A) \rightarrow (B, \partial^B)$ satisfying*

$$\tilde{K}|_R = 0 \quad \text{and} \quad \tilde{K}|_S = K.$$

Proof. All monomials in A are of the form $r_0 s_1 r_1 s_2 r_2 \cdots s_k r_k$ for some $k \geq 0$, with $r_0, \dots, r_k \in R$, and $s_1, \dots, s_k \in S^+$. Because the derivation property gives $\tilde{K}$ on products, we may prove inductively that if $\tilde{K}|_R = 0$ and $\tilde{K}|_S = K$, then on monomials, $\tilde{K}$ must take the form

$$\begin{aligned} & \tilde{K}(r_0 s_1 r_1 s_2 r_2 \cdots s_k r_k) \\ &:= \sum_{\ell=1}^k (-1)^{|s_1 \cdots s_{\ell-1}|} \phi(r_0 s_1 r_1 \cdots s_{j-1} r_{j-1}) K(s_j) \psi(r_j s_{j+1} r_{j+1} \cdots s_k r_k). \end{aligned}$$

This in turn determines $\tilde{K}$ uniquely on all of A upon extending by $\mathbb{Z}$ -linearity and is well-defined since the $\mathbb{Z}$ -linear relations between these monomials are generated by those of the form

$$\begin{aligned} & r_0 s_1 r_1 \cdots s_j (r_j + r'_j) s_{j+1} \cdots s_k r_k \\ &= r_0 s_1 r_1 \cdots s_j r_j s_{j+1} \cdots s_k r_k + r_0 s_1 r_1 \cdots s_j r'_j s_{j+1} \cdots s_k r_k. \end{aligned}$$

To check that this formula does indeed determine a (ϕ, ψ) -derivation, it suffices by $\mathbb{Z}$ -linearity to check the derivation equation on monomials, i.e., with

$$x = r_0 s_1 r_1 s_2 r_2 \cdots s_k r_k \quad \text{and} \quad y = r'_0 s'_1 r'_1 s'_2 r'_2 \cdots s'_k r'_k,$$

which we leave for the reader to verify. ■

Lemma 2.3 (Cf. [20, Lemma 2.18]). *Suppose that $(A = R * \mathbb{Z}\langle S \rangle, \partial^A)$ is a semi-free dg-algebra, (B, ∂^B) is a dg-algebra, $\phi, \psi: (A, \partial^A) \rightarrow (B, \partial^B)$ are dg-homomorphisms, and $K: (A, \partial^A) \rightarrow (B, \partial^B)$ is a (ϕ, ψ) -derivation. Then, the equation*

$$\phi - \psi = K\partial^A + \partial^B K$$

is true if and only if it holds on the restriction to both R and S .

Proof. Suppose that the equation is satisfied for elements $x \in A_k$ and $y \in A$, i.e.,

$$\phi(x) - \psi(x) = K\partial^A x + \partial^B Kx \quad \text{and} \quad \phi(y) - \psi(y) = K\partial^A y + \partial^B Ky.$$

Then, the equation also holds for the product xy by the following computation:

$$\begin{aligned} & K\partial^A(xy) + \partial^B K(xy) \\ &= K[(\partial^A x)y + (-1)^k x(\partial^A y)] + \partial^B [K(x)\psi(y) + (-1)^k \phi(x)K(y)] \\ &= K(\partial^A x)\psi(y) - (-1)^k \phi(\partial^A x)K(y) + (-1)^k K(x)\psi(\partial^A y) + \phi(x)K(\partial^A y) \\ &\quad + \partial^B K(x)\psi(y) - (-1)^k K(x)\partial^B \psi(y) + (-1)^k (\partial^B \phi(x))K(y) + \phi(x)\partial^B K(y) \\ &= (K\partial^A + \partial^B K)(x)\psi(y) + \phi(x)(K\partial^A + \partial^B K)(y) \\ &\quad + (-1)^k K(x)(\psi\partial^A - \partial^B \psi)(y) - (-1)^k (\phi\partial^A - \partial^B \phi)(x)K(y) \\ &= (\phi(x) - \psi(x))\psi(y) + \phi(x)(\phi(y) - \psi(y)) \\ &= \phi(x)\psi(y) - \psi(xy) + \phi(xy) - \phi(x)\psi(y) \\ &= \phi(xy) - \psi(xy). \end{aligned}$$

Since every element of A can be written as a $\mathbb{Z}$ -linear combination of products of elements of R and S (which are homogeneous), the result follows. ■

We arrive now at the main results of this section.

Proposition 2.4 (Cf. [15, Corollary 3.11]). *Every stable regenerative tame dg-homomorphism is a quasi-isomorphism.*

Proof. Since elementary dg-automorphisms and regenerations are dg-isomorphisms, they obviously act by quasi-isomorphisms. It suffices to check that stabilizations and destabilizations are also quasi-isomorphisms. We will check that they are inverses up to homotopy and hence inverses on homology.

Consider a stabilization and destabilization. The composition $A_S \rightarrow A_{S+} \rightarrow A_S$ is just the identity. It suffices to prove that the dg-homomorphism

$$\Phi: A_{S+} \rightarrow A_S \rightarrow A_{S+}$$

induces the identity on homology. By definition, Φ is just the identity on R and elements of S , but $\Phi(x) = 0 = \Phi(y)$. We aim now to build a chain homotopy between $\mathbb{1}$ and Φ , i.e., a $\mathbb{Z}$ -linear map $\tilde{\Psi}: A_{S+} \rightarrow A_{S+}$ of degree $+1$ such that

$$\mathbb{1} - \Phi = \tilde{\Psi}\partial^+ + \partial^+\tilde{\Psi}.$$

Consider the map $\Psi: S^+ \rightarrow A_{S+}$ given by

$$\Psi(s) = \begin{cases} 0, & s \neq x, \\ y, & s = x. \end{cases}$$

By Lemma 2.2, this extends to a $(\mathbb{1}, \Phi)$ -derivation $\tilde{\Psi}$. The chain homotopy equation is then easily verified on elements of R and S^+ , and so, by Lemma 2.3, $\tilde{\Psi}$ is a chain homotopy as desired. $\blacksquare$

Proposition 2.5. *Suppose that we have a dg-automorphism $\Phi: (A_S, \partial) \rightarrow (A_S, \partial)$ satisfying the following properties:*

- Φ fixes the underlying ring R .
- For each $s \in S$, we have $\partial s \in A_{S'}$, where $S' = \{s \in S \mid \Phi(s) = s\}$.
- For each $s \in S$, $\Phi(s) - s$ is exact.

Then, Φ is homotopic to the identity and hence induces the identity on homology.

Proof. To fix notation, suppose for each generator $s \in S \setminus S'$ that

$$s - \Phi(s) = \partial\beta_s$$

for some $\beta_s \in A_S$. Define a map $K: S \rightarrow A_S$ of degree $+1$ by

$$K(s) = \begin{cases} \beta_s, & s \notin S', \\ 0, & s \in S'. \end{cases}$$

Then, by Lemma 2.2, K extends uniquely to an $(\mathbb{1}, \Phi)$ -derivation $\tilde{K}: A_S \rightarrow A_S$ with $\tilde{K}|_R = 0$. Inductively from the defining equation for a derivation, one sees that $\tilde{K}|_{A_{S'}} = 0$. Therefore, for every $s \in S$, since $\partial s \in A_{S'}$, we have

$$(\tilde{K}\partial + \partial\tilde{K})(s) = \tilde{K}(\partial s) + \partial K(s) = 0 + \partial\beta_s = (\mathbb{1} - \Phi)(s).$$

We also have

$$(K\partial + \partial K)(r) = 0 = (\mathbb{1} - \Phi)(r)$$

for every $r \in R$. By Lemma 2.3, it follows that

$$\mathbb{1} - \Phi = \tilde{K}\partial + \partial\tilde{K},$$

and so, $\tilde{K}$ yields a chain homotopy between $\mathbb{1}$ and Φ . ■

Proposition 2.6. *Suppose that (A_S, ∂) is a semi-free dg-algebra over a ring R with stabilization (A_{S^+}, ∂^+) , with the property that all generators of S^+ (and hence also of S) have strictly positive degree. Then, for any positive integer r and field $\mathbb{F}$, the stabilization and destabilization maps induce inverse bijections*

$$\mathrm{Rep}_r(A_S, \partial; \mathbb{F}) \cong \mathrm{Rep}_r(A_{S^+}, \partial^+; \mathbb{F})$$

via pullback.

Proof. Let $i: (A_S, \partial) \rightarrow (A_{S^+}, \partial^+)$ and $\pi: (A_{S^+}, \partial^+) \rightarrow (A_S, \partial)$ be the stabilization and destabilization dg-homomorphisms. Since $\pi \circ i$ is the identity, we have

$$i^* \circ \pi^*: \mathrm{Rep}_r(A_S, \partial; \mathbb{F}) \rightarrow \mathrm{Rep}_r(A_S, \partial; \mathbb{F})$$

is the identity. On the other hand, $i \circ \pi$ is the identity on the coefficient ring, and representations are only non-trivial on the coefficient ring since all of the generators of S^+ have positive degree. Hence, we also have that its pullback

$$\pi^* \circ i^*: \mathrm{Rep}_r(A_{S^+}, \partial^+; \mathbb{F}) \rightarrow \mathrm{Rep}_r(A_{S^+}, \partial^+; \mathbb{F})$$

is the identity. ■

3. The coefficient ring

Recall from the introduction that, in generalizing the Casals–Murphy graded algebra (even before constructing the differential), there were two algebraic steps to perform, the first introducing non-commutativity between the coefficient ring and the other generators, and the second in introducing non-commutativity to the coefficient ring itself. The first will be discussed when we define the differential. But before we can even write the differential, we need a combinatorial model for the coefficient ring. We will develop a family of such combinatorial models, depending upon a choice of tree $T \subset G$ spanning all but one vertex which we will use as a base point. We will then see how to relate these models using homotopy classes of paths between the base points, a step which should feel well motivated by the statement of Theorem 1.6.




Figure 3. The green edges form a tree spanning all but one vertex, singled out also in green.

Remark 3.1. Throughout this section, we will assume that we have fixed a garden (to be defined later) and hence remove the corresponding decorations from our dg-algebras. Indeed, we are interested in the coefficient ring in this section, so it is better to keep the notation free from unnecessary clutter.

3.1. The commutative case

To begin, we consider the commutative case, as in the work of Casals and Murphy [3]. Recall from the introduction that, in the commutative case, we work over the coefficient ring $R = \mathbb{Z}[H_1(\Lambda_G)]$. The surface Λ_G is a genus g surface, so $H_1(\Lambda_G)$ is non-canonically isomorphic to $\mathbb{Z}^{2g}$. In order to have a combinatorial description for R , we therefore need to choose an isomorphism of $H_1(\Lambda_G)$ with $\mathbb{Z}^{2g}$. To do this, we make a choice of a tree $T \subset G$ spanning all but one of the vertices, as in Figure 3. Such a tree automatically contains $2g$ edges. Let $\mathbb{Z}^T$ be the free abelian group on the set of edges of T . Each edge in the graph G has preimage in Λ_G given by a circle, and that circle represents a corresponding 1-cycle in $H_1(\Lambda_G)$. (There is a minor issue of orientations which we ignore, as the geometry of Λ_G makes this choice canonical.) Furthermore, one may check that, for a tree as above, these 1-cycles form a basis for homology. One may then take the Poincaré dual basis of $H_1(\Lambda_G)$. Hence, we obtain for each edge $e \in T$ the unique primitive cycle in $H_1(\Lambda_G)$ which intersects e once (upon fixing orientations properly) and all other edges of the tree zero times. Thinking of this association as a map from T to $H_1(\Lambda_G)$, we obtain, from the universal property of free abelian groups, a group isomorphism

$$\varphi_T: \mathbb{Z}^T \xrightarrow{\sim} H_1(\Lambda_G).$$

Setting $R_T = \mathbb{Z}[\mathbb{Z}^T]$, upon applying the group ring functor, we obtain ring isomorphisms

$$\varphi_T: R_T \xrightarrow{\sim} R,$$

denoted by the same symbol by abuse of notation. Hence, for the underlying Casals–Murphy graded algebra $\mathcal{A}$, we have for each such tree T a combinatorial model $\mathcal{A}_T$, where

$$\mathcal{A}_T = R_T \langle F_{\text{fin}}, x, y, z \rangle.$$

The isomorphisms at the level of rings hence induce isomorphisms at the level of graded algebras

$$\varphi_T: \mathcal{A}_T \xrightarrow{\sim} \mathcal{A},$$

again notated by the same symbol by abuse of notation. The differential $\partial^{\mathcal{A}}$ may be computed in any of these models to yield dg-algebras $(\mathcal{A}_T, \partial_T^{\mathcal{A}})$, and we hence obtain canonical dg-isomorphisms

$$\varphi_T: (\mathcal{A}_T, \partial_T^{\mathcal{A}}) \xrightarrow{\sim} (\mathcal{A}, \partial^{\mathcal{A}}).$$

In other words, we may choose any of these combinatorial models as defining the same object. In order to change between combinatorial models, we simply apply the transition maps $\Phi_T^{T'} = \varphi_{T'}^{-1} \circ \varphi_T$.

The perspective of Casals and Murphy [3] is purely combinatorial, and hence, they work by constructing the various $(\mathcal{A}_T, \partial_T^{\mathcal{A}})$ and proving that they may all be canonically identified via canonical isomorphisms. This is surprisingly tricky, as the transition maps between R_T and $R_{T'}$ look somewhat complicated, and accounts for quite a bit of effort to achieve. In the non-commutative setting, we are automatically confronted with this issue; as described in the introduction, different choices of trees yield different choices of capping paths, and as is clear from the statement of Theorem 1.6, this makes the naturality of the corresponding transition dg-isomorphisms a bit harder. We follow the Casals–Murphy approach of working with a slightly larger dg-algebra, from which the models for each T are induced naturally. Transitioning between different choices of T in the non-commutative setting is then discussed in Section 3.4.

To this end, consider the collection E of all of the edges of the graph. Then, we may form the free abelian group $\mathbb{Z}^E$ and consider the group ring $\tilde{R} := \mathbb{Z}[\mathbb{Z}^E]$. We have natural projection maps

$$\pi_T: \mathbb{Z}^E \rightarrow \mathbb{Z}^T,$$

which upon applying the group ring functor yields projections

$$\pi_T: \tilde{R} \twoheadrightarrow R_T.$$

If we fix an explicit indexing of E such that the edges of T are listed first, this last map is just the map

$$\mathbb{Z}[e_1^{\pm 1}, \dots, e_{3g+3}^{\pm 1}] \twoheadrightarrow \mathbb{Z}[e_1^{\pm 1}, \dots, e_{2g}^{\pm 1}]$$

which evaluates a given Laurent polynomial at

$$e_{2g+1} = \cdots = e_{3g+3} = 1.$$

The strategy of Casals and Murphy [3] is to work instead over the coefficient ring $\tilde{R}$ and use the projections π_T to understand the transition maps. That is, on the graded algebra

$$\tilde{\mathcal{A}} = \tilde{R}\langle F_{\text{fin}}, x, y, z \rangle,$$

they find a differential $\tilde{\partial}^{\mathcal{A}}$. Reduction of coefficients then yields a graded algebra homomorphism

$$\pi_T: \tilde{\mathcal{A}} \rightarrow \mathcal{A}_T$$

given by projecting the coefficients along

$$\pi_T: \tilde{R} \rightarrow R_T.$$

This allows them to construct the differentials $\partial_T^{\mathcal{A}}$ via the following easy but important lemma.

Lemma 3.2. *Suppose that A_1 , A_2 , B_1 , and B_2 are objects in a category, with morphisms $\Phi^A: A_1 \rightarrow A_2$, $p_1: A_1 \rightarrow B_1$, and $p_2: A_2 \rightarrow B_2$. If p_1 is an epimorphism, then there exists at most one morphism $\Phi^B: B_1 \rightarrow B_2$ making the following diagram commute:*

$$\begin{array}{ccc} A_1 & \xrightarrow{\Phi^A} & A_2 \\ p_1 \downarrow & & \downarrow p_2 \\ B_1 & \xrightarrow{\Phi^B} & B_2 \end{array}$$

If p_1 has a right inverse $i_1: B_1 \rightarrow A_1$ (i.e., $p_1 \circ i_1 = \text{id}_{B_1}$), then such Φ^B exists if and only if

$$p_2 \circ \Phi^A \circ i_1 \circ p_1 = p_2 \circ \Phi^A,$$

in which case $\Phi^B = p_2 \circ \Phi^A \circ i_1$.

Proof. Uniqueness of Φ^B follows from the epimorphism property of p_1 . Assuming that p_1 has a right inverse, then we obtain a formula for such Φ^B , should it exist, as

$$\Phi^B = \Phi^B \circ p_1 \circ i_1 = p_2 \circ \Phi^A \circ i_1.$$

Commutativity of the diagram is just the desired equation. ■

As a special case, let us use the category of graded algebras to see how a differential may descend. Suppose that (A, ∂^A) is a dg-algebra, B is a graded algebra, $p: A \rightarrow B$ is a surjective morphism of graded algebras, and $i: B \rightarrow A$ is a right inverse

for p . Then, so long as $p \circ \partial^A \circ i \circ p = p \circ \partial^A$, by applying Lemma 3.2, we obtain $\partial^B = p \circ \partial^A \circ i$ on B fitting into the following diagram:

$$\begin{array}{ccc} A & \xrightarrow{\partial^A} & A \\ p \downarrow & & \downarrow p \\ B & \xrightarrow{\partial^B} & B \end{array}$$

We claim that (B, ∂_B) is a dg-algebra so that

$$p: (A, \partial^A) \rightarrow (B, \partial^B)$$

is a dg-homomorphism. The differential property follows because

$$(\partial^B)^2 = p \circ \partial^A \circ i \circ p \circ \partial^A \circ i = p \circ (\partial^A)^2 \circ i = 0,$$

where we have used that $p \circ \partial^A \circ i \circ p = p \circ \partial^A$. The Leibniz rule is easily checked from the fact that p and i are homomorphisms of graded algebras with $p \circ i = \text{id}$.

In the Casals–Murphy setting, one obtains the desired differentials $\partial_T^{\mathcal{A}}$ uniquely descending from $\tilde{\partial}^{\mathcal{A}}$ as above so that we have induced dg-homomorphisms

$$\pi_T: (\tilde{\mathcal{A}}, \tilde{\partial}^{\mathcal{A}}) \rightarrow (\mathcal{A}_T, \partial_T^{\mathcal{A}}).$$

To see this, one upgrades the natural right inverse $i_T: \mathbb{Z}^T \hookrightarrow \mathbb{Z}^E$ to $\pi_T: \mathbb{Z}^E \rightarrow \mathbb{Z}^T$ to the graded algebra level and checks that Lemma 3.2 applies. In this manner, $\partial_T^{\mathcal{A}}$ is actually very easy to compute. One simply looks at the combinatorial formulas for $\tilde{\partial}^{\mathcal{A}}$ (which are very explicit) and ignores any edge not in T which appears as an output, replacing it with the value 1.

3.2. The non-commutative case

In the non-commutative case, we follow the same procedure. Hence, we will need to define the ring $\tilde{R}^{nc}$ as well as the rings R_T^{nc} for each choice of tree T spanning all but one vertex v_T of G . We will prove in Section 3.3 that each of our rings R_T^{nc} is canonically isomorphic to the group ring $R_v^{nc} = \mathbb{Z}[\pi_1(\Lambda_G, \tilde{v})]$, where $v = v_T$ is the vertex missed by T .

Notice, in this case, that it is not really correct to talk about just one ring

$$R^{nc} = \mathbb{Z}[\pi_1(\Lambda_G)],$$

since we need to choose a base point in order to make this a true object. Indeed, there is no canonical isomorphism between the various R_v^{nc} as v varies in V . Although R has no non-commutative analog, we will find that $\tilde{R}$ does. Herein lies our point of entry into the theory.




Figure 4. A pictorial representation of the generators $\lambda_{e,\mathfrak{o}}$ and their inverses, in red, where the orientation $\mathfrak{o}$ on the edge e is indicated by the upwards arrow. We have also written these generators in terms of A_e and B_e , if we assume that $\mathfrak{o}$ is a constituent of a fixed orientation on the entire graph.

Let $\tilde{E}$ be the set of all edges together with an orientation. We denote an element by $(e, \mathfrak{o}) \in \tilde{E}$, where $e \in E$, and $\mathfrak{o}$ is an orientation on e . We denote the opposite orientation by $\bar{\mathfrak{o}}$. Consider the free group

$$*\tilde{E}\mathbb{Z} = \underbrace{\mathbb{Z} * \mathbb{Z} * \cdots * \mathbb{Z}}_{|\tilde{E}|=6g+6}$$

generated by elements of $\tilde{E}$. For each oriented edge $(e, \mathfrak{o}) \in \tilde{E}$, we denote by $\lambda_{e,\mathfrak{o}} \in *\tilde{E}\mathbb{Z}$ the corresponding generator.

It is worthwhile to make a visual picture for the generators $\lambda_{e,\mathfrak{o}}$. For every edge, we draw four oriented transversals, two oriented each direction, and where we think of them either from going from a state $+$ to a state $-$ or vice versa. Our $\lambda_{e,\mathfrak{o}}$ generators correspond to a transversal going from $+$ to $-$, and oriented via a right-hand rule, so that the transversal orientation followed by the edge orientation gives a positive orientation with respect to the sphere S^2 . This convention is easier illustrated than stated, as is done in Figure 4.

Elements of $*\tilde{E}\mathbb{Z}$ are just arbitrary words in the various $\lambda_{e,\mathfrak{o}}^{\pm 1}$, though for the most part, we only care about words formed from paths on S^2 as in the following definition.

Definition 3.3. A *transverse signed path* with respect to a trivalent plane graph G on S^2 is a pair (γ, σ) consisting of a curve $\gamma: [0, 1] \rightarrow S^2$ together with a sign $\sigma \in \{+, -\}$ such that

- the image of γ avoids the set of all vertices V ,
- if $\gamma(t) \in E$ lies on an edge, then $t \neq 0, 1$ and the path passes through the edge transversely.




Figure 5. The loop around a vertex which corresponds to the word w_v . If the graph is oriented as indicated by the arrows, this word is $w_v = A_{e_1} B_{e_2}^{-1} A_{e_3} B_{e_1}^{-1} A_{e_2} B_{e_3}^{-1}$.

Any transverse signed path (γ, σ) determines a unique lower-semicontinuous map $\tilde{\sigma}: [0, 1] \rightarrow \{+, -\}$ (we take $- < +$) with initial value $\tilde{\sigma}(0) = \sigma$ and which is

- locally constant away from values of t , where γ intersects an edge,
- discontinuous at values of t , where γ intersects an edge, alternating either from $+$ to $-$ or from $-$ to $+$.

In other words, we may label the path of γ , starting with the label σ , so that every time γ crosses an edge, it alternates its label. Given a transverse signed path, the map $\tilde{\sigma}$ in turn can be used to obtain an element of $*\tilde{E}\mathbb{Z}$ given by reading out, in order, the generators corresponding to every time we pass through an edge, using the convention as in Figure 4.

Definition 3.4. Any element of $*\tilde{E}\mathbb{Z}$ obtained in this way from a transverse signed path will be called a *geometric element*.

We provide not only a clarifying example, but a necessary example. Suppose that v is a vertex of the graph G , and let e_1 , e_2 , and e_3 be the incident edges, labeled cyclically counter-clockwise, as in Figure 5. For each e_i , let $\mathfrak{o}_i$ be the orientation along each edge pointing towards the vertex. (Note that it may be the case that $e_1 = e_2$, for example, in which case automatically $\mathfrak{o}_1 = \overline{\mathfrak{o}_2}$ since every edge has one source and one target.) Consider for each vertex the conjugacy class $w_v = [w_v]$ of the word

$$w_v := \lambda_{e_1, \mathfrak{o}_1} \lambda_{e_2, \overline{\mathfrak{o}_2}}^{-1} \lambda_{e_3, \mathfrak{o}_3} \lambda_{e_1, \overline{\mathfrak{o}_1}}^{-1} \lambda_{e_2, \mathfrak{o}_2} \lambda_{e_3, \overline{\mathfrak{o}_3}}^{-1},$$

which corresponds to a path going counter-clockwise around the vertex twice. We see that changing our starting basepoint will change the word, but only by cyclic permutation of its six constituent letters, and so, we will not change the conjugacy class ω_v . Let N be the normal subgroup of $*^{\tilde{E}}\mathbb{Z}$ generated by these conjugacy classes over all of the vertices. We hence obtain a group

$$\Pi := (*^{\tilde{E}}\mathbb{Z})/N.$$

In other words, Π has a presentation as $\Pi = \langle \{\lambda_{e,o}\} \mid \{w_v\} \rangle$, with $6g + 6$ generators, one for each element of $\tilde{E}$, and $2g + 2$ relations, one for each vertex $v \in V$. The image of the geometric elements of $*^{\tilde{E}}\mathbb{Z}$ in Π will be called the *geometric elements* of Π .

Remark 3.5. There is nothing special about taking our word w_v counter-clockwise. We could have instead considered the word corresponding to the composable loop going clockwise twice around the vertex

$$w'_v = \lambda_{e_1,o_1}^{-1} \lambda_{e_3,\bar{o}_3} \lambda_{e_2,o_2}^{-1} \lambda_{e_1,\bar{o}_1} \lambda_{e_3,o_3}^{-1} \lambda_{e_2,\bar{o}_2}.$$

In such a case, we have $(w'_v)^{-1} \in \omega_v$, so we could use any w'_v in place of w_v in the presentation of the group Π since the same normal subgroup will be generated.

Finally, taking the group ring, we arrive at our coefficient ring $\tilde{R}^{nc} := \mathbb{Z}[\Pi]$. There is a natural forgetful map

$$U: *^{\tilde{E}}\mathbb{Z} \twoheadrightarrow \mathbb{Z}^E,$$

given by abelianizing and forgetting about the orientation. We note that the word w_v is killed by U , yielding an induced surjective group homomorphism

$$U: \Pi \twoheadrightarrow \mathbb{Z}^E$$

also denoted U by abuse of notation. At the level of rings, we arrive at a homomorphism

$$U: \tilde{R}^{nc} \twoheadrightarrow \tilde{R}.$$

Now, suppose that $T \subset G$ is a tree spanning all but one vertex, and let $\tilde{T}$ be the collection of edges of the tree together with orientations. We have natural group homomorphisms

$$\pi_T: *^{\tilde{E}}\mathbb{Z} \twoheadrightarrow *^{\tilde{T}}\mathbb{Z}.$$

The map π_T simply sends $\lambda_{e,o}$ to the identity if $e \notin T$. Let N_T be $\pi_T(N)$, which is automatically a normal subgroup of $*^{\tilde{T}}\mathbb{Z}$ since π_T is surjective. We set

$$\Pi_T := *^{\tilde{T}}\mathbb{Z}/N_T$$

which has a presentation of the form

$$\Pi_T = \langle \{\lambda_{e, \circ}\} \mid \{\pi_T(w_v)\} \rangle,$$

with $4g$ generators corresponding to the elements $(e, \circ) \in \tilde{T}$, and $2g + 1$ relations, one for each vertex $v \in V - \{v_T\}$. (Notice that $\pi_T(w_{v_T}) = 1$, which explains why we have not included this trivial relation.) Setting $R_T = \mathbb{Z}[\Pi_T]$, we therefore have induced maps

$$\pi_T: \tilde{R}^{nc} \twoheadrightarrow R_T^{nc}.$$

As before, forgetting the orientations and abelianizing gives a surjective group homomorphism

$$U_T: \Pi_T \twoheadrightarrow \mathbb{Z}^T$$

inducing a surjective ring homomorphism

$$U_T: R_T^{nc} \twoheadrightarrow R_T.$$

It is clear that these forgetful maps commute with the projection maps, as per the following commutative diagram:

$$\begin{array}{ccc} \Pi & \xrightarrow{\pi_T} & \Pi_T \\ U \downarrow & & \downarrow U_T \\ \mathbb{Z}E & \xrightarrow{\pi_T} & \mathbb{Z}T \end{array}$$

Remark 3.6. We mentioned that we are mostly interested in the geometric elements of Π . The main reason is that every element in Π_T is a geometric element, in the sense that it lies in the image of the geometric elements of Π under the map

$$\pi_T: \Pi \rightarrow \Pi_T.$$

This point is further elucidated in Section 3.3, especially Theorem 3.10, upon the realization that transverse signed paths lift uniquely to paths on Λ_G . It is probably more mathematically optimal to use a certain path algebra instead of Π , where non-geometric elements are killed, but it is not necessary for our purposes.

Admittedly, the notation $\lambda_{e, \circ}$ is clumsy but is useful in that it indicates that $\tilde{R}^{nc}$ is completely independent of choices. However, in the case that our graph has an orientation chosen on each of its edges, which will be part of the data of a garden for other reasons, there is simpler notation which is useful for computations. Suppose that each edge e of E comes with a choice of orientation $\circ_e$. We obtain a bijection

$$\tilde{E} \cong E_+ \sqcup E_-,$$

where $E_+ \cong E$ consists of those edges matching the given orientation and $E_- \cong E$

consists of those edges which go against the given orientation. We will write

$$A_e := \lambda_{e, \circ_e}, \quad B_e := \lambda_{e, \overline{\circ_e}}.$$

We have traded the clumsiness of the notation $\lambda_{e, \circ}$ for the clumsiness of the asymmetry between A_e and B_e which is very much dependent upon the orientation chosen on each edge of G . Swapping the orientation of a single edge interchanges the roles of A and B . See Figure 4 or Figure 5 with its caption for examples. Note that the words w_v no longer necessarily alternate between A and B terms: this occurs only if all the edges are oriented away from the vertex (a source) or towards the vertex (a sink); except for very special graphs, it is impossible to choose an orientation so that every vertex is either a source or a sink.

3.3. Matching the coefficients to the geometry

We aim now to fill in the underlying geometric details of the combinatorics outlined in Section 3.2. The only geometric input we need from the contact picture is that we have a double cover

$$p: \Lambda_G \rightarrow S^2$$

branched over the set of vertices V such that, for each (open) face $f \in F$, the two connected components of $p^{-1}(f)$ are canonically labeled $+$ and $-$, respectively, with the property that if γ is a curve on Λ_G such that the curve $p \circ \gamma$ crosses precisely one edge transversely, then γ either goes from a region labeled $+$ to a region labeled $-$, or exactly the opposite, from $-$ to $+$.

It is clear from this description that transverse signed paths are identified canonically with paths on Λ_G by lifting along p and matching the starting label σ . Conversely, if we have a path in Λ_G which avoids the preimages of the vertices and passes transversely through the preimages of the edges, then its projection together with its initial label yields a transverse signed path. We call such a path on Λ_G *generic* so that transverse signed paths correspond precisely to generic paths on Λ_G .

On the other hand, transverse signed paths give the geometric elements of Π . Hence, we may think of the geometric elements as those which arise from generic paths on Λ_G . Suppose now that two generic paths with the same endpoints are homotopic to each other. We may ask if they define the same geometric element in Π . The point is that a homotopy from a generic path to another generic path may in general pass through non-generic paths. However, so long as the homotopy is itself generic, it will only pass through non-generic paths a finite number of times, and the non-generic behavior is quite mild. We have two cases:

- The family passes through a vertex $\tilde{v} \in \Lambda_G$ (corresponding to $v \in V$). In this case, so long as the homotopy is generic, this has the effect of changing the geometric

element by the vertex relation w_v , which is trivial in Π , and hence there is actually no effect.

- The family passes through a tangency to an edge. A good picture of the behavior for a generic homotopy is given by move II in Figure 7, but where the purple curve is thought of as our path at the critical moment (instead of a *tine*). In this case, the word is modified by insertion of a product $\lambda_{e,0}\lambda_{e,0}^{-1}$, which is just the identity, and hence the geometric element remains unchanged.

We have therefore proved the following theorem.

Theorem 3.7. *Suppose that $q, q' \in \Lambda_G$ are points with projection to S^2 lying in an open face of G . Suppose that γ_1 and γ_2 are two generic paths in Λ_G starting at q and ending at q' such that $[\gamma_1] = [\gamma_2] \in \pi_1(\Lambda_G, q, q')$, the homotopy classes of paths from q to q' . Then, the geometric elements $w_1, w_2 \in \Pi$ corresponding to γ_1 and γ_2 are equal, $w_1 = w_2$. We therefore obtain canonical group homomorphisms*

$$\Psi_{q,q'}: \pi_1(\Lambda_G, q, q') \rightarrow \Pi.$$

Furthermore, if $\gamma \in \pi_1(\Lambda_G, q, q')$ and $\gamma' \in \pi_1(\Lambda_G, q', q'')$, then

$$\Psi_{q,q''}(\gamma * \gamma') = \Psi_{q,q'}(\gamma) \cdot \Psi_{q',q''}(\gamma').$$

Remark 3.8. When $q = q'$, we write simply $\Psi_q := \Psi_{q,q}$.

Remark 3.9. Suppose that q_1 and q_2 are connected by a path in Λ_G which remains completely in the preimage of an open face $f \in F$; i.e., they lie in the same connected component of the preimage of f in Λ_G . Since this component is just one of the sheets lying over f , it is a topological disk, the groups $\pi_1(\Lambda_G, q_1, q')$ and $\pi_1(\Lambda_G, q_2, q')$ are canonically identified, and under this identification, $\Psi_{q_1,q'} = \Psi_{q_2,q'}$. Similarly, we have an equality $\Psi_{q,q'_1} = \Psi_{q,q'_2}$ upon identifying the domains of these maps.

Theorem 3.7 has given us an algebraic output in Π for paths on Λ_G . We will now find a natural way to recognize Π_T as identified with the fundamental group of Λ_G , with basepoint given by the unique preimage $*_T = p^{-1}(v_T)$ in Λ_G singled about the vertex v_T . Towards this purpose, notice that $p^{-1}(S^2 \setminus T)$ is a branched double cover of a topological disk, branched at a single point, so $p^{-1}(S^2 \setminus T)$ is a topological disk. Therefore, if $q \in p^{-1}(S^2 \setminus T)$, then there is a unique homotopy class of path in $\gamma_{q,T} \in \pi_1(\Lambda_G, *_T, q)$ which admits a representative whose projection avoids T . With this notation, we have the following theorem.

Theorem 3.10. *There is a canonical isomorphism*

$$\varphi_T: \Pi_T \xrightarrow{\sim} \pi_1(\Lambda_G, *_T)$$

such that, for any $q, q' \in \Lambda_G$ lying above an open face of G , the composition

$$\varphi_T \circ \pi_T \circ \Psi_{q,q'}: \pi_1(\Lambda_G, q, q') \rightarrow \pi_1(\Lambda_G, *_T)$$

is the isomorphism given by

$$[\gamma] \mapsto [\gamma_{q,T} * \gamma * (\gamma_{q',T})^{-1}].$$

In particular, for $\lambda_{e,o} \in \Pi$, writing $\lambda_{e,o} = \Psi_{q_f^+, q_g^-}(\gamma_{e,o})$ for a short path $\gamma_{e,o}$ starting at q_f^+ on the $+$ sheet above face f , ending at q_g^- on the $-$ sheet above face g , and with projection crossing only the edge e once, we have

$$(\varphi_T \circ \pi_T)(\lambda_{e,o}) = [\gamma_{q_f^+, T} * \gamma_{e,o} * (\gamma_{q_g^-, T'})^{-1}].$$

Proof. For notational convenience, around each vertex $v \in V$ (including v_T), we draw a small open disk D_v in S^2 , all of which are disjoint, and we pick any point $q_v \in \partial \overline{D_v}$ not lying on an edge of G . Let $v(V) = \bigsqcup D_v$ be the union of these disks. Also for notational convenience, we fix on the edges of T an orientation. To compute the fundamental group $\pi_1(\Lambda_G, *_T)$, we proceed in three steps.

Step 1. First, we compute the fundamental group $\pi_1(X, q_{v_T})$, where $X := S^2 \setminus v(V)$.

The space X is just a copy of S^2 with $2g + 2$ disks removed, which deformation retracts to a wedge of $2g + 1$ circles, so $\pi_1(X; q_{v_T})$ is isomorphic (non-canonically) to a free group of rank $2g + 1$. We aim to find a convenient system of generators. Since T is contractible, for each edge $e \in T$, we may pick a loop $\overline{A}_e \in \pi_1(X, q_{v_T})$ which intersects the tree only along the edge e transversely so that it is oriented in the same way that we described for $\lambda_{e,o} = A_e$ with respect to the chosen orientation on e (see again Figure 4). Also, let $\gamma \in \pi_1(X; q_{v_T})$ be the loop which follows $\partial \overline{D_{v_T}}$ around once, say, in the counter-clockwise direction (although this choice will end up being irrelevant). These $\overline{A}_e$'s, together with γ , form the generators we are after.

Lemma 3.11. *The collection $\{\overline{A}_e \mid e \in T\} \cup \{\gamma\}$ forms a free system of generators for $\pi_1(X; q_{v_T})$.*

Though the choice of $\overline{A}_e$ is not canonical, it is easy to encode the ambiguity in this choice, and this ambiguity will happily disappear in the end of our computation of $\pi_1(\Lambda_G, *_T)$. The following lemma, which makes this ambiguity explicit, follows easily from the fact that $S^2 \setminus (T \sqcup \{v_T\})$ is a punctured disk with fundamental group generated by the loop γ .

Lemma 3.12. *Suppose that $\overline{A}_e$ and $\overline{A}'_e$ are two choices of paths transversely intersecting T only along the edge e with the same orientation. Then, there exist (unique) integers $m, n \in \mathbb{Z}$ such that*

$$\overline{A}_e = \gamma^m \overline{A}'_e \gamma^n.$$

Step 2. We compute $\pi_1(Y, q_{v_T}^+)$, where $Y = p^{-1}(X)$ and $q_{v_T}^+$ is the lift of q_{v_T} to the face labeled $+$.

The projection $p: Y \rightarrow X$ induces a map

$$p_*: \pi_1(Y, q_{v_T}^+) \rightarrow \pi_1(X, q_{v_T}).$$

Since $p: Y \rightarrow X$ is a connected double cover, we automatically have a short exact sequence of groups

$$0 \rightarrow \pi_1(Y, q_{v_T}^+) \xrightarrow{p_*} \pi_1(X, q_{v_T}) \xrightarrow{\text{par}} \mathbb{Z}_2 \rightarrow 0.$$

For each class in $\pi_1(X, q_{v_T})$, a generic representative will intersect the edges transversely, and the map par is simply the parity of the number of edges that the curve intersects.

Recall the following facts.

Theorem 3.13. *Suppose that G is a finitely generated group, and $H \leq G$ is a finite index subgroup. We have the following:*

- (*Schreier Subgroup Lemma*) *Consider a right inverse $\sigma: H \setminus G \rightarrow G$ of the quotient $\pi: G \rightarrow H \setminus G$. Let S be a set of generators for G . Then, H is generated by the $|S| \cdot [G: H]$ elements*

$$g_{c,s} := \sigma(c) \cdot s \cdot (\sigma(\pi(\sigma(c) \cdot s)))^{-1}$$

for pairs $(c, s) \in H \setminus G \times S$.

- (*Nielsen–Schreier Theorem*) *If G is a free group, then H is a free group. Furthermore, if G has rank n and H is of finite index k , then H has rank $1 + k(n - 1)$.*

Since $\pi_1(X, q_{v_T})$ is a free group, these two results automatically apply to our setting. By the Nielsen–Schreier theorem, we see that $\pi_1(Y, q_{v_T}^+)$ is a free group of rank $1 + 2((2g + 1) - 1) = 1 + 4g$. On the other hand, we may use the generating set $S = \{\bar{A}_e \mid e \in T\} \cup \{\gamma\}$ and the section

$$\sigma: \mathbb{Z}_2 \rightarrow \pi_1(X, q_{v_T})$$

with

$$\sigma(0) = e, \quad \sigma(1) = \gamma$$

as input into the Schreier subgroup lemma. If we let $m_e := \text{par}(\bar{A}_e) \in \{0, 1\}$ be the parity of the loops $\bar{A}_e$ that we have chosen, then

$$\begin{cases} g_{0, \bar{A}_e} = \bar{A}_e \cdot \gamma^{-m_e}, \\ g_{0, \gamma} = \gamma \cdot \gamma^{-1} = e, \\ g_{1, \bar{A}_e} = \gamma \cdot \bar{A}_e \cdot \gamma^{m_e-1}, \\ g_{1, \gamma} = \gamma^2. \end{cases}$$

Forgetting the non-identity element, we see that we are left with precisely $4g + 1$ generators. As this matches the rank of the free group it generates, there are no relations between these generators by a result of Nielsen [24] (or its English translation [25]); cf. the work of Federer and Jónsson [16, Theorem 3.13].

Of the two generators $\bar{A}_e \cdot \gamma^{-m_e}$ and $\gamma \cdot \bar{A}_e \cdot \gamma^{m_e-1}$, precisely one of them crosses the edge e in such a way that it goes from a $+$ region to a $-$ region. We let $\tilde{A}_e$ be this generator, and we let the other be $\tilde{B}_e^{-1}$ (which matches our intuition from Figure 4). We set $\tilde{\gamma} = \gamma^2$. Hence, we have

$$\pi_1(Y, q_{v_T}^+) = \langle \tilde{\gamma}, \{\tilde{A}_e \mid e \in T\}, \{\tilde{B}_e \mid e \in T\} \rangle.$$

Again, $\tilde{A}_e$ and $\tilde{B}_e$ are not canonically defined, since our choice of $\bar{A}_e$ was only defined up to multiplying on the left and right by powers of γ . Nonetheless, it is easy to trace through how these are allowed to change. The following lemma is an easy exercise for the reader if it is not obvious.

Lemma 3.14. *For two different choices $\bar{A}_e$ and $\bar{A}'_e$, the corresponding generators $\tilde{A}_e$ and $\tilde{A}'_e$ are related by*

$$\tilde{A}_e = \tilde{\gamma}^m \tilde{A}'_e \tilde{\gamma}^n$$

for (unique) integers $m, n \in \mathbb{Z}$, and similarly for $\tilde{B}_e$ and $\tilde{B}'_e$.

We finish Step 2 by noting not only that we have a nice presentation for $\pi_1(Y, q_{v_T}^+)$, but also that, for a given element in this fundamental group, it is quite easy to read off, up to powers of $\tilde{\gamma}$, the corresponding word in these generators simply by looking at the projection to X of a generic representative. In particular, for a generic representative, this projection will cross the edges of G transversely and be disjoint from all of the vertices. From this, we have that our loop starts on the $+$ sheet, and every time it crosses an edge, it swaps between the $+$ and $-$ sheets. Suppose that this loop crosses the edges of T in the order $e_1, e_2, \dots, e_n$. Then, the corresponding word representing this loop is given by

$$\tilde{\gamma}^{m_1} \cdot g_1 \cdot \tilde{\gamma}^{m_2} \cdot g_2 \cdots \tilde{\gamma}^{m_n} \cdot g_n \cdot \tilde{\gamma}^{m_{n+1}}$$

for some integers $m_i \in \mathbb{Z}$ and where each $g_i \in \{\tilde{A}_{e_i}, \tilde{A}_{e_i}^{\pm 1}, \tilde{B}_{e_i}, \tilde{B}_{e_i}^{\pm 1}\}$ is chosen according to the convention of Figure 4.

Step 3. Finally, we glue back in $p^{-1}(v(V))$ and compute $\pi_1(\Lambda_G, *T)$.

Each $p^{-1}(D_v)$ is itself a disk, so we may attach them one by one and apply the Seifert–van Kampen theorem to compute the fundamental group. We begin with the disk $p^{-1}(D_{v_T})$, which has boundary given by the loop $\tilde{\gamma}$, so attaching this disk simply kills this generator. In particular, we have that all of the choices $\tilde{A}_e$ and $\tilde{B}_e$ already correspond to well-defined classes A_e and B_e . So, at this point, we have

$$\pi_1(Y \cup p^{-1}(D_{v_T}), q_{v_T}^+) = \langle \{A_e \mid e \in T\}, \{B_e \mid e \in T\} \rangle.$$

In fact, we see already from the description at the end of Step 2 that it is quite easy to read off the corresponding word from the projection of the loop to $X \cup D_{v_T}$, with no ambiguity in choices since we have $\tilde{\gamma} = 1$.

Now, suppose that we glue in a single disk $p^{-1}(D_v)$. This corresponds to killing the loop $\partial \overline{D_v}$, which we associate with an element of $\pi_1(Y \cup p^{-1}(D_{v_T}), q_{v_T}^+)$ by choosing an arbitrary path from $q_{v_T}^+$ to q_v^+ which does not intersect the tree. We see that the word w_v (appropriately evaluated so that all non-tree edges are given the value 1) is precisely the word generating this loop. Therefore,

$$\pi_1(\Lambda_G, q_{v_T}^+) = \langle \{A_e \mid e \in T\}, \{B_e \mid e \in T\} \mid \{\pi_T(w_v)\} \rangle,$$

which is precisely the group Π_T we described, i.e., $\Pi_T \cong \pi_1(\Lambda_G, q_{v_T}^+)$. Finally, all paths between $*_T$ and $q_{v_T}^+$ completely contained in $p^{-1}(D_{v_T})$ are homotopic, and this allows us to move our base point from $q_{v_T}^+$ to $*_T$ canonically.

Finally, to prove the expression for $\varphi_T \circ \pi_T \circ \Psi_{q,q'}$, we note that, by the composition property of Theorem 3.7, it suffices to prove the result for the short generators $\gamma_{e,o}$. But this is given by tracing through the proof and identifying the generators A_e and B_e with these explicit loops. ■

Corollary 3.15. *There exists a group homomorphism $i_T: \Pi_T \rightarrow \Pi$ which is right inverse to $\pi_T: \Pi \rightarrow \Pi_T$.*

Proof. Consider any point $q \in \Lambda_G$ lying above an open face of G . We obtain an identification $\mathcal{C}(\gamma_{q,T})\pi_1(\Lambda_G, *_T) \cong \pi_1(\Lambda_G, q)$ by conjugation by the path $\gamma_{q,T} \in \pi_1(\Lambda_G, *_T, q)$. Set

$$i_T := \Psi_q \circ \mathcal{C}(\gamma_{q,T}) \circ \varphi_T: \Pi_T \rightarrow \Pi.$$

This is a composition of group homomorphisms, hence a group homomorphism itself, and satisfies

$$\begin{aligned} \pi_T \circ i_T &= \pi_T \circ \Psi_q \circ \mathcal{C}(\gamma_{q,T}) \circ \varphi_T \\ &= \varphi_T^{-1} \circ (\varphi_T \circ \pi_T \Psi_q) \circ \mathcal{C}(\gamma_{q,T}) \circ \varphi_T \\ &= \varphi_T^{-1} \circ \mathcal{C}(\gamma_{q,T})^{-1} \circ \mathcal{C}(\gamma_{q,T}) \circ \varphi_T \\ &= \mathbb{1}, \end{aligned}$$

where we have used Theorem 3.10 to identify $\varphi_T \circ \pi_T \circ \Psi_q$ with $\mathcal{C}(\gamma_{q,T})^{-1}$. ■

To end the section, we discuss how Theorem 3.10 relates to the commutative coefficient setting. From the universal property for abelianization of groups, the map

$$U_T \circ \varphi_T^{-1}: \pi_1(\Lambda_G, *_T) \twoheadrightarrow \mathbb{Z}^T$$

must factor through the abelianization map

$$\pi_1(\Lambda_G, *_T) \twoheadrightarrow H_1(\Lambda_G),$$

yielding a surjective morphism $H_1(\Lambda_G) \twoheadrightarrow \mathbb{Z}^T$, as in the following commutative diagram:

$$\begin{array}{ccc} \Pi_T & \xrightarrow[\sim]{\varphi_T} & \pi_1(\Lambda_G, *_T) \\ U_T \downarrow & & \downarrow \text{Ab} \\ \mathbb{Z}^T & \xleftarrow{\quad} & H_1(\Lambda_G) \end{array}$$

But this surjective morphism is between two free abelian groups of the same rank $2g$ and hence is automatically an isomorphism. In other words, $\mathbb{Z}^T$ is itself the abelianization of Π_T , U_T is the abelianization map, and φ_T descends to the isomorphism

$$\varphi_T: \mathbb{Z}^T \rightarrow H_1(\Lambda_G)$$

which we used in the commutative case. We leave it to the reader to check that this definition of the isomorphism $\varphi_T: \mathbb{Z}^T \xrightarrow{\sim} H_1(\Lambda_G)$ matches our original description at the beginning of Section 3.1.

3.4. Transition maps

As discussed in Section 3.1, in order to define the Casals–Murphy dg-algebra (in the commutative setting), one uses the following steps:

- define the tilde version $(\tilde{\mathcal{A}}, \tilde{\partial}^{\mathcal{A}})$;
- plug the tilde version into Lemma 3.2 by using the maps π_T and i_T to obtain the dg-algebras $(\mathcal{A}_T, \partial_T^{\mathcal{A}})$;
- verify that the various versions with respect to different trees T are naturally identified together.

The first two steps hold over in exactly the same manner, with no changes in the non-commutative setting. It is the last step which we have not yet discussed. We will work with the (only slightly more difficult) non-commutative setting since the strategies are essentially the same; see Remark 3.18 below for a discussion of the commutative case.

Suppose that γ is a homotopy class of path on Λ_G from $*_T$ to $*_{T'}$. Then, there is an induced isomorphism $\mathcal{C}(\gamma): \pi_1(\Lambda_G, *_T) \xrightarrow{\sim} \pi_1(\Lambda_G, *_{T'})$ given by

$$\mathcal{C}(\gamma)(\eta) := \gamma^{-1} * \eta * \gamma.$$

We aim to find, for each pair of trees T and T' and homotopy class of path γ from $*_T$ to $*_{T'}$, an isomorphism $\tilde{\mathcal{C}}_T^{T'}(\gamma): \Pi \rightarrow \Pi$ fitting into the following commutative

diagram:

$$\begin{array}{ccc}
 \Pi & \xrightarrow{\tilde{\mathcal{C}}_T^{T'}(\gamma)} & \Pi \\
 \pi_T \downarrow & & \downarrow \pi_{T'} \\
 \Pi_T & \xrightarrow{\quad \mathcal{C}_T^{T'}(\gamma) \quad} & \Pi_{T'} \\
 \varphi_T \downarrow & & \downarrow \varphi_{T'} \\
 \pi_1(\Lambda_G, *_T) & \xrightarrow{\mathcal{C}(\gamma)} & \pi_1(\Lambda_G, *_{T'})
 \end{array}$$

The middle arrow is dashed to indicate that it arises from the bottom square automatically as $\mathcal{C}_T^{T'}(\gamma) := \varphi_{T'}^{-1} \circ \mathcal{C}(\gamma) \circ \varphi_T$, since all morphisms in this composition are isomorphisms. We see then that, by Lemma 3.2, the isomorphism $\mathcal{C}_T^{T'}(\gamma)$ is the unique one fitting into the top square.

Let us define $\tilde{\mathcal{C}}_T^{T'}(\gamma)$ on generators. Suppose that we have some element $\lambda_{e,o} \in \Pi$ for $e \in E$. As in the statement of Theorem 3.10, it is identified with a short transverse path $\gamma_{e,o}$ from q_f^+ on the $+$ sheet over the face f to q_g^- on the $-$ sheet over g , with projection crossing the edge e , i.e., $\lambda_{e,o} = \Psi_{q_f^+, q_g^-}(\gamma_{e,o})$. For any pairs of trees T and T' , and homotopy class of path γ from $*_T$ to $*_{T'}$, let

$$\beta_T^{T'}(f, \gamma, \pm) := \Psi_{q_f^\pm}[(\gamma_{q_f^\pm, T'})^{-1} * \gamma^{-1} * \gamma_{q_f^\pm, T}]$$

be the geometric element in Π corresponding to the homotopy class of the loop $(\gamma_{q_f^\pm, T'})^{-1} * \gamma^{-1} * \gamma_{q_f^\pm, T}$. We define

$$\begin{aligned}
 (\tilde{\mathcal{C}}_T^{T'}(\gamma))(\lambda_{e,o}) &:= \beta_T^{T'}(f, \gamma, +) \cdot \lambda_{e,o} \cdot \beta_T^{T'}(g, \gamma, -)^{-1} \\
 &= \Psi_{q_f^+, q_g^-}(\gamma_{q_f^+, T'}^{-1} * \gamma^{-1} * \gamma_{q_f^+, T} * \gamma_{e,o} * \gamma_{q_g^-, T}^{-1} * \gamma * \gamma_{q_g^-, T'}).
 \end{aligned}$$

In order to check this is well-defined, we must check that it preserves the normal subgroup N of $*^{\tilde{E}}\mathbb{Z}$ generated by the words w_v . But we see that because w_v is given as a geometric element, the inner factors of β come in pairs $\beta^{-1} \cdot \beta$ which cancel, and we are left with a conjugate of w_v . Hence, this is a well-defined group homomorphism of Π .

Lemma 3.16. *The diagram*

$$\begin{array}{ccc}
 \Pi & \xrightarrow{\tilde{\mathcal{C}}_T^{T'}(\gamma)} & \Pi \\
 \varphi_T \circ \pi_T \downarrow & & \downarrow \varphi_{T'} \circ \pi_{T'} \\
 \pi_1(\Lambda_G, *_T) & \xrightarrow{\mathcal{C}(\gamma)} & \pi_1(\Lambda_G, *_{T'})
 \end{array}$$

is commutative.

Proof. It suffices to check the result for generators. Suppose that we have $\lambda_{e,o}$, corresponding to the short transverse signed path $\gamma_{e,o}$ from q_f^+ to q_g^- . Using Theorem 3.10 to compute $\varphi_T \circ \pi_T$ and $\varphi_{T'} \circ \pi_{T'}$, we have

$$\begin{aligned}
& (\varphi_{T'} \circ \pi_{T'} \circ \tilde{\mathcal{C}}_T^{T'}(\gamma))(\lambda_{e,o}) \\
&= (\varphi_{T'} \circ \pi_{T'})(\Psi_{q_f^+, q_g^-}(\gamma_{q_f^+, T'}^{-1} * \gamma^{-1} * \gamma_{q_f^+, T} * \gamma_{e,o} * \gamma_{q_g^-, T}^{-1} * \gamma * \gamma_{q_g^-, T'})) \\
&= \gamma_{q_f^+, T'} * \gamma_{q_g^-, T'}^{-1} * \gamma^{-1} * \gamma_{q_f^+, T} * \gamma_{e,o} * \gamma_{q_g^-, T}^{-1} * \gamma * \gamma_{q_g^-, T'} * \gamma_{q_g^-, T'}^{-1} \\
&= \gamma^{-1} * \gamma_{q_f^+, T} * \gamma_{e,o} * \gamma_{q_g^-, T}^{-1} * \gamma \\
&= \mathcal{C}(\gamma)(\gamma_{q_f^+, T} * \gamma_{e,o} * \gamma_{q_g^-, T}^{-1}) \\
&= (\mathcal{C}(\gamma) \circ \varphi_T \circ \pi_T)(\lambda_{e,o}).
\end{aligned}$$

Hence, the diagram does indeed commute. $\blacksquare$

Recalling the various functoriality properties of Theorem 1.6, we will require the following lemma, where we recall again that without a tilde:

$$\mathcal{C}_T^{T'}(\gamma) = \varphi_T^{-1} \circ \mathcal{C}(\gamma) \circ \varphi_T.$$

Lemma 3.17. *The following two properties hold:*

- For all trees T and constant paths $\mathbb{1}$ at $*_T$, the maps $\tilde{\mathcal{C}}_T^T(\mathbb{1})$ and $\mathcal{C}_T^T(\mathbb{1})$ are the identity on Π and Π_T , respectively.
- For any triples of trees T , T' , and T'' and homotopy classes of paths γ from $*_T$ to $*_{T'}$ and γ' from $*_{T'}$ to $*_{T''}$, we have

$$\tilde{\mathcal{C}}_T^{T''}(\gamma * \gamma') = \tilde{\mathcal{C}}_T^{T''}(\gamma') \circ \tilde{\mathcal{C}}_T^{T'}(\gamma) \quad \text{and} \quad \mathcal{C}_T^{T''}(\gamma * \gamma') = \mathcal{C}_T^{T''}(\gamma') \circ \mathcal{C}_T^{T'}(\gamma).$$

Proof. The versions without the tildes follow from the ones with the tildes by Lemma 3.2. The first property is clear by definition. The second property follows from the fact that

$$\beta_T^{T''}(f, \gamma, \pm) = \beta_{T'}^{T''}(f, \gamma', \pm) \cdot \beta_T^{T'}(f, \gamma, \pm)$$

because the corresponding paths are homotopic:

$$\begin{aligned}
& (\gamma_{q, T''})^{-1} * (\gamma * \gamma')^{-1} * \gamma_{q, T} \\
& \sim [(\gamma_{q, T''})^{-1} * (\gamma')^{-1} * \gamma_{q, T'}] * [(\gamma_{q, T'})^{-1} * \gamma^{-1} * \gamma_{q, T}],
\end{aligned}$$

where $q = q_f^\pm$. $\blacksquare$

Unfortunately, it is not the case that the naively induced graded algebra morphisms $\tilde{\Phi}_T^{T'}(\gamma): \tilde{\mathcal{B}} \rightarrow \tilde{\mathcal{B}}$ preserve the differential $\tilde{\partial}^{\mathcal{B}}$. We will need to further act on generators

via regenerations (as suggested in Theorem 1.6). However, since we have not yet constructed the differentials, it is too early to discuss such matters. We relegate a treatment to Section 6.1.

Remark 3.18. The story is much the same in the commutative case and can be recovered from the non-commutative case upon abelianization. Namely, one finds a collection of isomorphisms

$$\tilde{\Phi}_T^{T'}: \mathbb{Z}^E \rightarrow \mathbb{Z}^E$$

with $\tilde{\Phi}_T^T = \mathbb{1}$ and $\tilde{\Phi}_T^{T'} \circ \tilde{\Phi}_{T'}^{T''} = \Phi_T^{T''}$ and fitting into the following commutative diagram:

$$\begin{array}{ccc} \mathbb{Z}^E & \xrightarrow{\tilde{\Phi}_T^{T'}} & \mathbb{Z}^E \\ \pi_T \downarrow & & \downarrow \pi_{T'} \\ \mathbb{Z}^T & \xrightarrow{\Phi_T^{T'}} & \mathbb{Z}^{T'} \\ \varphi_T \downarrow & & \downarrow \varphi_{T'} \\ H_1(\Lambda_G) & \equiv & H_1(\Lambda_G) \end{array}$$

Our construction of a combinatorial model for $\Phi_T^{T'}$ is a little different from the one implicitly presented in Casals and Murphy, who use a group action of $\mathbb{Z}^F$ on $\tilde{\mathcal{A}}[t^{\pm 1}]$ to obtain $(\mathcal{A}, \partial^{\mathcal{A}})$ as a GIT quotient. Along the way, they construct maps $s_T: \mathbb{Z}^T \rightarrow \mathbb{Z}^E$ [3, Proof of Theorem 6.6] such that one may take $\Phi_T^{T'} = \pi_{T'} \circ s_T$. It is not much harder to construct the lift $\tilde{\Phi}_T^{T'}$ as well, though this is not done in their paper.

4. The enlarged non-commutative Casals–Murphy dg-algebra

In this section, we will define a dg-algebra associated to a trivalent plane graph G and a choice of garden Γ . This dg-algebra is enlarged from the version appearing in Theorem 1.6. The functoriality properties of this enlarged dg-algebra are then proved as Theorem 5.1 in Section 5.

4.1. Gardens, binary sequences, and associated words

In order to define our dg-algebra, we first need to make a number of auxiliary choices, which are grouped together in the following convenient definition.

Definition 4.1. Suppose that G is a trivalent plane graph. A *garden* on G consists of the data of

- an *orientation*, by which we mean an orientation of all the edges of G ,

- a *centering*, consisting of, for each $f \in F$, a point c_f contained in the interior of the face f , called the *center* of f ,
- a *web*, consisting of, for each face-vertex adjacency (f, v) , a *thread* $\tau_f(v)$, which is a connected embedded arc oriented from c_f to v completely contained in f except for its endpoint at v such that no two threads intersect except at their endpoints at the centers and vertices, and such that their tangent vectors at c_f are all distinct,
- a *seed*, given by a covector in the unit cotangent bundle $\xi \in S^*_s S^2$ for a *base point* $s \in S^2$,
- a *rake*, consisting of oriented embedded loops γ_f , called *tines*, one for each $f \in F$, which start and end at the base point s , pass through c_f and no other center, have initial and terminal tangent vectors positive with respect to ξ , and intersect each other only at the base point s .

The face $f_\Gamma \in F$ to which the base point s belongs is called the *face at infinity with respect to Γ* .

An oriented path $\gamma: [a, b] \rightarrow S^2$ is said to be *non-degenerate* with respect to a garden Γ if

- its endpoints do not lie on a vertex, edge, center, or thread,
- it passes through edges and threads transversely at interior points,
- it passes through centers so that it is not tangent to any thread exiting that center.

The garden is said to be *non-degenerate* if the tines of the rake are all non-degenerate, and *degenerate* otherwise.

A *homotopy* of gardens is a smoothly varying path of gardens Γ_t , $0 \leq t \leq 1$. We always allow our homotopies to pass through degenerate gardens in which the tines may be degenerate; on the other hand, homotopies must still be through gardens. Hence, for example, the base point of the seed cannot pass through a thread or edge, since then it would not be a garden at that time. Similarly, a tine γ_f cannot homotope to non-trivially intersect another tine nor to pass through a face other than c_f .

Remark 4.2. Our definition of a garden is slightly different from the version of Casals–Murphy [3]; their version is essentially the same, except that they take the base point s of the seed to lie in a chosen face at infinity, and they exclude a tine through the face at infinity. We will handle this slightly differently when we define finite-type gardens in Definition 6.2. Our definition is essentially equivalent to theirs but places it in our context in which all of the faces are treated equally and where the base point is free to move around.

See Figure 6 for an example of a non-degenerate garden.




Figure 6. A non-degenerate garden. We have elected, for simplicity, to stereographically project from the base point so that s is the point at infinity; the seed ξ is chosen so that all tines, represented by the magenta lines, are oriented from the bottom of the picture to the top. (See Figure 9 for how we picture the seed, when needed.) The graph, together with the orientations of the edges, are notated in black. The green circles are the centers, and the web is given by the blue threads. Since the threads are always oriented from the centers to the vertices, we exclude their orientations. On the fourth tine is a sample binary sequence from 0 to 0, which contributes a term to the differential of z of the form $(H_1)(A_2)(H_3)(-B_4)(-H_5)(-B_6)(f_7)(A_8)(H_9)(A_{10})$, where the A and B terms correspond to where the tine passes an edge and are short for the corresponding elements A_e and B_e in W , the H terms are shorthand for the element $H_f(v)$ for where the tine switches at a thread, and f_7 is the face through which the tine passes.

Up to homotopy, any two non-degenerate gardens are equivalent via swapping the orientation on edges together with finite sequences of seven classes of moves, moves I–VII. The first five, moves I–V, were studied by Casals and Murphy [3]. (Our version of move V is presented in a slightly modified form.) Remark 4.2 explains why moves VI and VII appear here but not in the original study of Casals and Murphy.

moves I–IV are illustrated in Figure 7. Move I forms or removes a bigon between a tine and a thread, whereas move II likewise forms or removes a bigon between a tine and an edge. Move III passes a tine through a vertex of the graph. Move IV swings the tangent vector to a thread at a vertex past the tangent vector, the tine. If Γ and Γ' are any two homotopic non-degenerate gardens, then one may always find a homotopy between them which is a composition of isotopies of the garden on the graph and a composition of moves I–IV.




Figure 7. Moves I–IV. Implicit are all possible orientations. For example, for move IV, we can also swing the tangent vector to a thread around in the clockwise direction as opposed to the counter-clockwise direction drawn.




Figure 8. Move V. Suppose that we have two adjacent tines, and consider a curve from one center to the next which intersects no tines but may of course intersect edges or threads. This is depicted as an orange arc in the left diagram. Note that in the dark gray region (and outside of the local picture), the garden may be arbitrary. We have removed the threads emanating from the centers for simplicity; they may be moved arbitrarily by move IV. Then, we may take the tines and perform a positive half-twist of the two centers supported in a neighborhood of a connecting curve. In the diagram, the tines tunnel through the light gray rectangle so that they swap centers.

Move V is demonstrated in Figure 8. Move V interchanges centers along a generic curve connecting the two centers. A single application of move V will change the homotopy type of the garden, though there are also non-trivial sequences of applications of move V which preserve the homotopy type. These matters will be discussed in great detail later.

Finally, moves VI and VII are illustrated in Figure 9. They involve passing a seed through a thread or an edge. Again, a single application of move VI or VII will in general change the homotopy type of the garden, as such a homotopy restricts the




Figure 9. Moves VI and VII. We pass the seed of the garden through a thread or an edge, respectively. The base point of the seed is located at the center of each diagram, and the seed itself is represented by the small vertical line and arrow, representing the kernel of the seed ξ and its coorientation. Hence, in the diagrams presented, all the tines are oriented from right to left.

base point to lie in the same region of S^2 bounded by the graph and the web (just a triangle). However, a non-trivial sequence of applications of moves VI and VII may in general preserve a garden, as will also be discussed in great detail later.

Every non-degenerate path has three types of critical behaviors, passing through either a center c_f , an edge of the graph, or the interior of a thread. Note that we are not interested in intersections of such paths with the tines, which do not appear in the definition of a non-degenerate path. Also, notice that the tines of a non-degenerate garden are themselves non-degenerate paths. For each type of intersection point, we study a certain type of combinatorial interaction by way of binary sequences. The following definition appears in the commutative setting of Casals and Murphy [3] where it was originally defined only for tines. We make the same definition more generally for non-degenerate paths.

Definition 4.3. Suppose that γ is an embedded path in S^2 which is non-degenerate with respect to a trivalent plane graph G and garden Γ . A *binary sequence along γ* is a lower semicontinuous function $B: [a, b] \rightarrow \{0, 1\}$ with the following properties:

- B is locally constant away from the three types of intersection points.
- At a center, B is discontinuous and must switch from 0 to 1.
- At an edge, B is discontinuous and must switch either from 0 to 1 or from 1 to 0.
- At the interior of a thread, B is either locally constant or switches from 0 to 1.

For $i, j \in \{0, 1\}$, the set of binary sequences along γ starting at i and ending at j is denoted by $\mathcal{BS}_{i,j}(\gamma)$.

To each binary sequence B along γ , we encode a monomial

$$w_B \in \tilde{R}^{nc} * \mathbb{Z}\langle F \rangle = \mathbb{Z}[\Pi] * \mathbb{Z}\langle F \rangle,$$

with coefficient ± 1 , determined as follows. To determine w_B , we consider the set D_B of discontinuities of the binary sequence, labeled in the order that they appear on the

tine. Each element of D_B corresponds either to an edge, a thread, or a face. Note that D_B does *not* in general contain every thread the path γ intersects, but only those where the binary sequence is discontinuous. For each element of $d \in D_B$, we determine a monomial in $\tilde{R}^{nc} * \mathbb{Z}\langle F \rangle$ with coefficient ± 1 , described in the next paragraphs, and set w_B to be the product of these monomials in order. Let us now describe the monomial associated to each discontinuity. As we proceed, the reader should verify that our conventions match with the example in Figure 6, which shows a sample binary sequence along a tine, and which has the associated word described in the caption; recall again that the notion we are defining works for all non-degenerate paths and not just tines.

For an edge, the binary sequence always switches, either from 0 to 1, or vice versa. Associating the label 0 with the label $-$ and the label 1 with the label $+$, we have that the path crosses the edge in a manner which matches one of the four diagrams of Figure 4, where our path γ is taken in place of the tine. Accordingly, we associate to this discontinuity the corresponding generator A_e , A_e^{-1} , $-B_e$, or $-B_e^{-1}$, where we note that we have included a negative sign in two of these cases.

For a discontinuity with a thread $\tau_f(v)$, we associate with it a monomial

$$H(\tau_f(v)) \in \tilde{R}^{nc} = \mathbb{Z}[\Pi]$$

defined as follows. Consider a counter-clockwise loop around the vertex v based at c_f which goes from the label $-$ to the label $+$, and where the B and B^{-1} terms come with negative signs (as we did for edges). In other words, if we look at Figure 5, if the center c_f is in the region with the red circle and the thread flows from the red circle to the vertex v , then we have $H(\tau_f(v)) = B_{e_1}^{-1} A_{e_2} B_{e_3}^{-1}$ by following a single loop counter-clockwise around v from $-$ to $+$, and where the sign is positive because there are two B terms in this case. Notice that the sign of $H(\tau_f(v))$ is automatically $(-1)^{r_v}$, where r_v is the number of outgoing edges. As for our construction of w_B , we include another sign, using $\pm H(\tau_f(v))$ for a discontinuity at $\tau_f(v)$, where the sign is determined by the following right-hand rule: the sign is $+1$ or -1 according to whether the orientation of the thread followed by the orientation of the path gives a frame matching the underlying orientation on the plane. The signs are indicated in Figure 10.

Remark 4.4. In our convention for $H(\tau_f(v))$, we have used the geometric element associated to a counter-clockwise loop around v . We could have alternatively used clockwise loops in our convention, which has the effect of simply negating $H(\tau_f(v))$, which would be equivalent to using the counter-clockwise convention but for the opposite orientation on G .

Finally, if we have a discontinuity at a center c_f , we simply replace it with $f \in F$, considered as a monomial in $\tilde{R}^{nc} * \mathbb{Z}\langle F \rangle$. This finishes the definition of w_B ; we




Figure 10. The sign placed in front of $H(\tau_f(v))$ for a binary sequence which changes from 0 to 1 at a thread. The horizontal lines are threads and the vertical lines represent our path γ to which we are associating our monomial.

repeat, as we said at the beginning of this discussion, that an example is found in Figure 6 and its caption.

Remark 4.5. If $D_B = \emptyset$, then we set $w_B = 1$.

To each non-degenerate immersed path γ and each $i, j \in \{0, 1\}$, we may therefore define four associated words $w_{i,j} \in \tilde{R}^{nc} = \mathbb{Z}[\Pi]$, given by

$$w_{i,j}(\gamma) := \sum_{B \in \mathcal{BS}_{i,j}(\gamma)} w_B.$$

Remark 4.6. If $\mathcal{BS}_{i,j}(\gamma) = \emptyset$, then $w_{i,j}(\gamma) = 0$.

We end this subsection by establishing two fundamental properties about the words $w_{i,j}(\gamma)$. The first is a statement about locality, showing that breaking up non-degenerate paths into smaller non-degenerate paths plays nicely with the associated word construction. The second involves invariance of the associated word under mild homotopies. Both of these are used in the commutative coefficient case by Casals and Murphy [3] within various proofs; here, we single them out in separate lemmas.

Lemma 4.7. *Suppose that γ_0 and γ_1 are non-degenerate paths such that the terminal point of γ_0 is the initial point of γ_1 . Let $\gamma = \gamma_0 * \gamma_1$ be the composition of paths, smoothed along where they glue. Then,*

$$w_{i,j}(\gamma) = \sum_{k=0}^1 w_{i,k}(\gamma_0) \cdot w_{k,j}(\gamma_1).$$

Proof. We have obvious bijections

$$\mathcal{BS}_{i,j}(\gamma) = (\mathcal{BS}_{i,0}(\gamma_0) \times \mathcal{BS}_{0,j}(\gamma_1)) \sqcup (\mathcal{BS}_{i,1}(\gamma_0) \times \mathcal{BS}_{1,j}(\gamma_1)),$$

and the two sides of the equation in the lemma are just given by writing out the corresponding words according to these two associations. ■

With this lemma proved, it is convenient from now on, for a non-degenerate path γ , to write

$$W(\gamma) := \begin{pmatrix} w_{0,0}(\gamma) & w_{0,1}(\gamma) \\ w_{1,0}(\gamma) & w_{1,1}(\gamma) \end{pmatrix} \in \text{Mat}_{2 \times 2}(\tilde{R}^{\text{nc}} * \mathbb{Z}\langle F \rangle).$$

We will call this the *matrix of words* for γ . Lemma 4.7 implies that a composition of non-degenerate paths is obtained by matrix multiplication of the corresponding matrices of words, i.e., $W(\gamma_1 * \gamma_2) = W(\gamma_1) \cdot W(\gamma_2)$. Even more is true; Lemma 4.8 below implies that the matrix of words is an invariant of non-degenerate paths relative to homotopies (through possibly degenerate paths) not crossing the centering of the garden.

Lemma 4.8. *Suppose that γ_0 and γ_1 are two non-degenerate paths with the same endpoints which are homotopic via a homotopy γ_t (of not necessarily non-degenerate paths) which never intersects the centering of the garden. Then, the matrices of words are equal,*

$$W(\gamma_0) = W(\gamma_1).$$

Additionally, if γ_0 and γ_1 are two non-degenerate paths differing by an application of move IV, then again $W(\gamma_0) = W(\gamma_1)$.

Proof. Any such homotopy may be generically perturbed so that it is given (up to isotopy of S^2) by a finite sequence of moves I–III as in Figure 7 and creation/removal of self-intersections, with the understanding that our paths take the place of the tines. (Indeed, this lemma will primarily be applied to tines later on.) Since self-intersections do not affect the matrix of words, it suffices to check invariance under any of moves I–IV. By Lemma 4.7, this is just a local check. To be precise, let γ and γ' be the before and after for each curve. Then, we may decompose

$$\begin{aligned} \gamma &= \eta * \gamma_\ell * \mu, \\ \gamma' &= \eta * \gamma'_\ell * \mu, \end{aligned}$$

where γ_ℓ and γ'_ℓ represent the local pieces on the left and right of the corresponding figure. It suffices to prove that

$$W(\gamma_\ell) = W(\gamma'_\ell),$$

i.e., that $w_{i,j}(\gamma_\ell) = w_{i,j}(\gamma'_\ell)$ for each $i, j \in \{0, 1\}$, since then

$$W(\gamma) = W(\eta)W(\gamma_\ell)W(\mu) = W(\eta)W(\gamma'_\ell)W(\mu) = W(\gamma').$$

For move I, notice on the right that because there are only thread crossings, the only local binary sequences for which there is a discontinuity are from 0 to 1, and hence,

$$\begin{cases} w_{0,0}(\gamma_\ell) = 1 = w_{0,0}(\gamma'_\ell), \\ w_{1,1}(\gamma_\ell) = 1 = w_{1,1}(\gamma'_\ell), \\ w_{1,0}(\gamma_\ell) = 0 = w_{1,0}(\gamma'_\ell). \end{cases}$$

On the other hand, whereas $\mathcal{BS}_{0,1}(\gamma_\ell) = \emptyset$, we have $\mathcal{BS}_{0,1}(\gamma'_\ell)$ consisting of two elements, depending upon which thread crossing is chosen to give a discontinuity. However, those two binary sequences come with opposite thread orientations, and hence,

$$w_{0,1}(\gamma'_\ell) = \pm H(\tau) \mp H(\tau) = 0 = w_{0,1}(\gamma_\ell).$$

For move II, we have each $W(\gamma_\ell) = I$ being the identity matrix since there are no crossings in the left diagram. On the other hand, on the right, both crossings must be discontinuities and must switch the label. Clearly, $w_{1,0}(\gamma'_\ell) = w_{0,1}(\gamma'_\ell) = 1$ since there are no corresponding binary sequences. But also, it is easy to see that we have chosen our orientations so that also

$$w_{0,0}(\gamma'_\ell) = w_{1,1}(\gamma'_\ell) = 1.$$

For example, if both the edge and the tine are oriented upwards in the diagram, then

$$w_{1,1}(\gamma'_\ell) = (-B_e)(-B_e^{-1}) = 1,$$

where e is the edge.

For move III, suppose without loss of generality that our paths are directed from bottom left to top right. Starting from the bottom left endpoint of the local path and working clockwise, suppose that the edges and tines are labeled $e_1, \tau_1, e_2, \tau_2, e_3, \tau_3$. Suppose that the orientations of all the edges are inwards towards the vertex; we leave the other orientations to the reader. (The orientations of the threads are always inwards pointing, so we need not specify them.) One checks that, with our orientations, $H(\tau_i) = B_{e_i}^{-1} A_{e_{i-1}} B_{e_{i-2}}^{-1}$, where the indices are taken modulo 3. The following relations are easily checked, since the corresponding local binary sequence sets are singletons:

$$\begin{cases} w_{0,0}(\gamma_\ell) = (A_{e_1}^{-1})(-B_{e_2}) = (-H(\tau_3))A_{e_3} = w_{0,0}(\gamma'_\ell), \\ w_{1,1}(\gamma_\ell) = (-B_{e_1})(A_{e_2}^{-1}) = A_{e_3}(-H(\tau_2)) = w_{1,1}(\gamma'_\ell), \\ w_{1,0}(\gamma_\ell) = (-B_{e_1})(H(\tau_1))(-B_{e_2}) = A_{e_3} = w_{1,0}(\gamma'_\ell). \end{cases}$$

For the final case, we have $\mathcal{BS}_{0,1}(\gamma_\ell) = \emptyset$ so that $w_{0,1}(\gamma_\ell) = 0$. On the other hand, $\mathcal{BS}_{0,1}(\gamma'_\ell)$ consists of two elements: one which crosses only at the edge, and the other

which also crosses at both threads. We see that

$$w_{0,1}(\gamma'_\ell) = -B_{e_3}^{-1} + (-H(\tau_3))(A_{e_3})(-H(\tau_2)) = 0.$$

For move IV, the fact that the binary sequence must change from 0 to 1 at the center severely restricts the possible local binary sequences. In particular, on the right-hand side, if the path is oriented upwards, then the binary sequence cannot switch at the thread because it would have to change from 0 to 1, but must be 0 when it enters the face. Similar reasoning applies if the path is oriented downwards. Regardless, we find that

$$\begin{cases} w_{0,0}(\gamma_\ell) = 0 = w_{0,0}(\gamma'_\ell), \\ w_{1,1}(\gamma_\ell) = 0 = w_{1,1}(\gamma'_\ell), \\ w_{1,0}(\gamma_\ell) = 0 = w_{1,0}(\gamma'_\ell), \\ w_{0,1}(\gamma_\ell) = f = w_{0,1}(\gamma'_\ell), \end{cases}$$

where f is the face corresponding to the center c_f in the diagram through which the path passes. ■

4.2. The enlarged dg-algebra

Define

$$\tilde{\mathcal{B}}_G^+ := \tilde{R}^{\text{nc}} * \mathbb{Z}\langle F, x, y, z, w \rangle,$$

the free associative graded algebra over the ring $\tilde{R}^{\text{nc}}$ generated by all faces F (including any chosen face at infinity) of degree 1 and four extra generators x, y, z, w of degree 2. This is of course a slightly different object than that appears in Theorems 1.4 and 1.6, but we will be able to use Lemma 3.2 together with the transition maps of Section 3.4 to eventually prove our desired results. The tilde in the notation refers to the coefficient ring $\tilde{R}^{\text{nc}}$ (as opposed to R_T), and the plus in the notation refers to the inclusion of the face at infinity and the generator w which do not appear in the original formulation of Casals and Murphy in the commutative coefficient setting [3]. We will discuss removing these enlargements in Section 6.

For a given garden Γ on G , we define a linear degree -1 endomorphism $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ by specifying its action on generators and extending via the Leibniz rule

$$\tilde{\partial}(a \cdot b) = (\tilde{\partial}a) \cdot b + (-1)^{|a|} a \cdot (\tilde{\partial}b)$$

for homogeneous elements $a, b \in \tilde{\mathcal{B}}_G^+$. To simplify notation, let

$$X = \begin{pmatrix} z & w \\ y & x \end{pmatrix}$$

be the matrix consisting of the four non-face generators of $\tilde{\mathcal{B}}_G^+$. Then, $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ acts on the generators by

$$\begin{aligned}\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} f &= \sum_{v \in f} H(\tau_f(v)), \quad \forall f \in F, \\ \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} X &= \sum_{f \in F} W(\gamma_f).\end{aligned}$$

Remark 4.9. A complete example of an enlarged dg-algebra for a specific pair (G, Γ) is included in Appendix A.

It is not immediately clear that $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ is actually a differential. Since we have defined it so that it is of degree -1 and is extended via the Leibniz rule, it suffices to check that $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ squares to zero. We begin with a lemma which allows us to extend Lemma 4.8 to the setting in which a homotopy of non-degenerate paths passes through a center.

Lemma 4.10. *Suppose that γ_0 and γ_1 are two non-degenerate paths with the same endpoints which are homotopic via a homotopy γ_t (of not necessarily non-degenerate paths) such that there is a single such path γ_{t_0} (with $t_0 \neq 0, 1$) passing through a single center so that the frame given by the vector $\frac{d}{dt}|_{t=t_0} \gamma_t$ followed by the tangent vector to γ_{t_0} forms a positively oriented frame at this center. Then, for any garden Γ on G ,*

$$W(\gamma_1) - W(\gamma_0) = \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} W(\gamma_{t_0}).$$

Proof. Let us write

$$\gamma_{t_0} := \gamma^s * \gamma^c * \gamma^t,$$

where γ^c represents a small portion of γ^{t_0} which crosses through the center and none of the edges or threads, and γ^s and γ^t represent the starting and terminal portions of the curve. By homotoping the homotopy γ_t , we may assume without loss of generality that, for some small $\varepsilon > 0$, we may write

$$\gamma_{t_0-\varepsilon} = \gamma^s * \gamma^{c,L} * \gamma^t \quad \text{and} \quad \gamma_{t_0+\varepsilon} = \gamma^s * \gamma^{c,R} * \gamma^t,$$

where $\gamma^{c,L}$ and $\gamma^{c,R}$ are small arcs with the same endpoints at γ^c but passing on opposite sides of the corresponding center. We notice that the curves $\gamma^{c,L}$ and $\gamma^{c,R}$ have corresponding matrices of words of the form

$$W(\gamma^{c,L}) = \begin{pmatrix} 1 & w_{0,1}(\gamma^{c,L}) \\ 0 & 1 \end{pmatrix}, \quad W(\gamma^{c,R}) = \begin{pmatrix} 1 & w_{0,1}(\gamma^{c,R}) \\ 0 & 1 \end{pmatrix}$$

since the corresponding arcs only interact with the threads of the center. In fact, each thread τ around this center contributes a unique local binary sequence from 0 to 1 for

either $\gamma^{c,L}$ or $\gamma^{c,R}$, and the corresponding monomial $H(\tau)$ comes with a minus sign for $\gamma^{c,L}$ and a plus sign for $\gamma^{c,R}$. That is, if the center is c_f , then

$$w_{0,1}(\gamma_k^{c,R}) - w_{0,1}(\gamma_k^{c,L}) = \sum_{v \in f} H(\tau_f(v)) = \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} f.$$

Using Lemma 4.8, we have $W(\gamma_0) = W(\gamma_{t_0-\varepsilon})$ and $W(\gamma_1) = W(\gamma_{t_0+\varepsilon})$. Therefore, from our computation, and from the fact that γ^s and γ^t pass through no centers so that

$$\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} W(\gamma^s) = 0 = \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} W(\gamma^t),$$

we find that

$$\begin{aligned} W(\gamma_1) - W(\gamma_0) &= W(\gamma_{t_0+\varepsilon}) - W(\gamma_{t_0-\varepsilon}) \\ &= W(\gamma^s)W(\gamma^{c,R})W(\gamma^t) - W(\gamma^s)W(\gamma^{c,L})W(\gamma^t) \\ &= W(\gamma^s)[W(\gamma^{c,R}) - W(\gamma^{c,L})]W(\gamma^t) \\ &= W(\gamma^s) \begin{pmatrix} 0 & \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} f \\ 0 & 0 \end{pmatrix} W(\gamma^t) \\ &= W(\gamma^s) \left[\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} \right] W(\gamma^t) \\ &= W(\gamma^s) [\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} W(\gamma^c)] W(\gamma^t) \\ &= \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} [W(\gamma^s)W(\gamma^c)W(\gamma^t)] \\ &= \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} W(\gamma_{t_0}). \end{aligned} \quad \blacksquare$$

We may now prove the differential property.

Proposition 4.11. *For any non-degenerate garden Γ for a trivalent plane graph G , the linear maps $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ are differentials, i.e.,*

$$(\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})^2 = 0.$$

Hence, $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})$ is a dg-algebra with respect to any choice of non-degenerate garden.

Proof. It suffices to check that $(\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})^2 = 0$ on generators. For degree reasons, this holds on faces, and it suffices to check the relation on the generators x , y , z , and w , which we conveniently group as the matrix X .

For the given garden Γ , let $f_1, \dots, f_{g+3}$ be the enumeration of the faces given by enumerating the tines in clockwise order as they emanate from the seed. For example, in Figure 6, the face at infinity is given as f_2 ; alternatively, in any of the diagrams in

Figure 9, the tines are ordered from bottom to top, both as they leave the seed and as they exit the seed. Choose non-degenerate loops $\eta_0, \eta_1, \dots, \eta_{g+3}$ based at the seed such that γ_{f_k} lies between η_{k-1} and η_k . By Lemma 4.10, we find that

$$W(\eta_k) - W(\eta_{k-1}) = \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} W(\gamma_{f_k}).$$

Summing all of these up, we find that

$$W(\eta_{g+3}) - W(\eta_0) = \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} \sum_{k=1}^{g+3} W(\gamma_{f_k}) = (\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})^2 X.$$

But notice that both η_0 and η_{g+3} may be contracted to a point so that

$$W(\eta_0) = I = W(\eta_{g+3})$$

is just the identity matrix. Hence, indeed, $(\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})^2 X = 0$. ■

Remark 4.12. Our proof of Proposition 4.11 is simpler than the one presented by Casals and Murphy [3]. They build an involution of pairs consisting of binary sequences and threads, each of which contributes a term to $(\tilde{\partial}_{G,\Gamma}^{\mathcal{A}})^2$ (recall that they do not include the enlargement in which we include two extra generators, so we have not included a $+$ in the notation), and check that the signs of these terms are opposite of each other. One could reason in precisely the same way in the non-commutative setting with no complication. Their argument is closer to the contact-geometric reason why the differential should square to zero: canceling pairs identified via this involution are rigid gradient flow trees at the boundaries of moduli spaces of (non-rigid) gradient flow trees of dimension 1.

5. Functoriality for the enlarged dg-algebra

At this point, we have defined an enlarged version $(\mathcal{B}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})$ of the non-commutative Casals–Murphy dg-algebra. Some of the functoriality properties of Theorem 1.6 will not appear at the enlarged level, since there is no choice of tree, and since we are not restricting to finite-type gardens. However, the composition property, changing the orientation, and changing the garden do make perfect sense at the enlarged level. The goal in this section is to prove the following theorem.

Theorem 5.1. *For any choice of garden Γ on a trivalent plane graph G of genus g , there exists a differential $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ on the graded algebra $\tilde{\mathcal{B}}_G^+$, making it into a dg-algebra. For any two gardens which are homotopic, the differentials are canonically identified.*

There exists a group action of $\mathcal{H}(G) := (\mathbb{Z}_2)^E \times F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$ on the space of gardens up to homotopy. For each $\zeta \in \mathcal{H}(G)$ and homotopy class of garden Γ on G , with $\Gamma' = \zeta \cdot \Gamma$, there is a dg-isomorphism

$$\Phi_{\Gamma'}^{\Gamma'}(\zeta): (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\zeta \cdot \Gamma}^{\mathcal{B}^+})$$

satisfying the following properties:

- *Composition property:* the dg-isomorphisms compose naturally, in the sense that

$$\Phi_{\Gamma'}^{\Gamma'}(\zeta) \circ \Phi_{\Gamma'}^{\Gamma'}(\theta) = \Phi_{\Gamma'}^{\Gamma'}(\zeta\theta).$$

- *Orientation changes:* the action of the factor $(\mathbb{Z}_2)^E$ is induced by ring automorphisms of $\mathbb{Z}[\Pi]$.
- *Modifying the garden:* for any $\zeta \in F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$ (i.e., trivial in the $(\mathbb{Z}_2)^E$ factor), the map $\Phi_{\Gamma'}^{\Gamma'}(\zeta)$ is a regenerative tame dg-isomorphism. Furthermore, if ζ stabilizes Γ , i.e., $\zeta \cdot \Gamma = \Gamma$, then $\Phi_{\Gamma'}^{\Gamma'}(\zeta)$ is homotopic to the identity and hence acts by the identity in homology.

We have already verified that $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ is a differential which is unaffected by moves I–IV, and hence, our dg-isomorphisms will be constructed for orientation changes along with each of moves V, VI, and VII. We proceed in the following order, subsection by subsection:

- *Changing Orientations:* we construct a dg-isomorphism every time we change the orientation of an edge simply by negating the corresponding generators in the coefficient ring.
- *Invariance under move V:* we construct a tame dg-isomorphism for each application of move V. Further, we show that move V yields a natural free transitive braid group action Br_{g+3} on the homotopy classes of the space of gardens with a fixed seed. We have a surjection $F_{g+2} \twoheadrightarrow \text{Br}_{g+3}$, the kernel of which acts by dg-automorphisms homotopic to the identity.
- *Exactness of the full twist:* the full twist in the braid group preserves the homotopy class of the garden, as it may be realized by just twisting the seed by 360 degrees. We prove that the full twist acts by a dg-automorphism homotopic to the identity.
- *Invariance under moves VI and VII:* we construct a dg-isomorphism for each application of moves VI and VII. We show that these commute with move V, and further show that moves VI and VII yield a $\mathbb{Z}_2 * \mathbb{Z}_3$ action on the homotopy classes of gardens, and understand the corresponding regenerative tame dg-isomorphisms.
- *Proof of Theorem 5.1:* in this short subsection, we simply summarize the results we have obtained, which proves the desired theorem.

5.1. Changing orientations

We have already seen as a consequence of Lemma 4.8 that moves I–IV preserve the differential on the nose. In other words, if two gardens Γ and Γ' differ by a composition of isotopies and applications of moves I–IV, then $\tilde{\partial}_\Gamma = \tilde{\partial}_{\Gamma'}$, and hence, we may take $\Phi_\Gamma^{\Gamma'} = 1$. We turn our attention now to the remaining ways in which we may change a garden, by changing the orientation of an edge, or by moves V–VII. In this subsection, we focus on changing orientations.

There is an obvious $(\mathbb{Z}_2)^E$ action on the homotopy classes of gardens given by swapping orientations at corresponding edges. Let us consider the simplest case, in which we swap a single edge; $e_0 \in E$ is an edge. Let us define a ring automorphism $\Phi_{e_0}: \mathbb{Z}[\Pi] \rightarrow \mathbb{Z}[\Pi]$ on generators as follows. Recall that Π is generated by elements of the form $\lambda_{e,o}$, where $(e, o) \in \tilde{E}$. Define Φ_{e_0} on these generators by

$$\Phi_{e_0}(\lambda_{e,o}) = \begin{cases} \lambda_{e,o}, & e \neq e_0, \\ -\lambda_{e,o}, & e = e_0. \end{cases}$$

In turn, this induces an graded algebra automorphism of $\tilde{\mathcal{B}}_G^+$, acting by the identity on generators. Furthermore, it is obvious that $(\Phi_{e_0})^2 = 1$, and that Φ_{e_0} and Φ_{e_1} commute for any $e_0, e_1 \in E$. Hence, we also have a homomorphism

$$(\mathbb{Z}_2)^E \rightarrow \text{Aut}(\tilde{\mathcal{B}}_G^+).$$

For $\iota \in (\mathbb{Z}_2)^E$ (the character ι is used because we have an *involution* on the orientations of edges), we write Φ_ι for the induced automorphism of $\tilde{\mathcal{B}}_G^+$.

Lemma 5.2. *For any homotopy class of garden Γ and $\iota \in (\mathbb{Z}_2)^E$, the graded algebra automorphism $\Phi_\iota: \tilde{\mathcal{B}}_G^+ \rightarrow \tilde{\mathcal{B}}_G^+$ intertwines the differentials $\tilde{\partial}_{G,\Gamma}^+$ and $\tilde{\partial}_{G,\iota,\Gamma}^+$, yielding a dg-isomorphism*

$$\Phi_\iota: (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\iota,\Gamma}^+).$$

Proof. It suffices to check the equation

$$\Phi_\iota \circ \tilde{\partial}_{G,\Gamma}^+ = \tilde{\partial}_{G,\iota,\Gamma}^+ \circ \Phi_\iota$$

on generators. If we apply both sides to a face f , then this reads

$$\sum_f \Phi_\iota(H^\Gamma(\tau_f(v))) = \sum_f H^{\iota,\Gamma}(\tau_f(v)),$$

where we have augmented the notation $H(\tau_f(v))$ to point out that it depends upon a choice of orientation of the edges. But this relation is clear term by term for each f by definition of $H(\tau_f(v))$, i.e.,

$$\Phi_\iota(H^\Gamma(\tau_f(v))) = H^{\iota,\Gamma}(\tau_f(v)),$$

since both sides are $\pm H^\Gamma(\tau_f(v))$ depending upon the parity of the number of edges incident to v which change orientation. Similarly, applying both sides to one of the generators x , y , z , or w , it suffices to check that

$$\Phi_\iota(W^\Gamma(\gamma_f)) = W^{\iota\Gamma}(\gamma_f)$$

for each tine γ_f . In turn, this resorts to checking that, for each binary sequence B along γ_f , we have $\Phi_\iota(w_B^\Gamma) = w_B^{\iota\Gamma}$. Each w_B is a product of monomials for each discontinuity of B , and it suffices to check for each such discontinuity. For threads, we have already verified the relation. For centers, there is no sign. For an edge, the corresponding terms $\pm \lambda_{e,0}^{\pm 1}$ in w_B^Γ and $w_B^{\iota\Gamma}$ are the same if ι leaves the orientation of the edge unfixed and negatives of each other if ι reverses the orientation of the edge. But this negation is the same as the action of Φ_ι . ■

5.2. Invariance under move V

We turn now to understanding move V. We refer to Figure 8 again, with the convention that the tines are oriented upwards, c_f is the center on the left, and c_g is the center on the right. Let γ_f and γ_g be the tines in the left figure and γ'_f and γ'_g the tines in the right figure. Recall in the proof of Proposition 4.11 that we wrote each tine as a composition of starting, central, and terminal portions, e.g.,

$$\gamma_f = \gamma_f^s * \gamma_f^c * \gamma_f^t.$$

Both γ_f^s and $(\gamma'_f)^s$ have the same start point (the seed) and endpoint (just under c_g), and we see that they are in fact homotopic with fixed endpoints without passing through any centers, since move V is applied to neighboring tines. Hence, by Lemma 4.8,

$$W(\gamma_f^s) = W((\gamma'_f)^s).$$

On the other hand, γ_f^t and $(\gamma'_f)^t$ have the same start point (just above c_f) and endpoint (the seed), and we see that there is a natural homotopy between them which passes through the center c_g . We saw in Lemma 4.10 how a homotopy passing through a center affects the corresponding words. In particular, if η is the orange curve in the figure, oriented from c_f to c_g , then

$$W((\gamma'_f)^t) - W(\gamma_f^t) = W(\eta) \begin{pmatrix} 0 & \tilde{\partial}_{G,\Gamma}^+ g \\ 0 & 0 \end{pmatrix} W(\gamma_g^t).$$

By similar reasoning,

$$W(\gamma_g^t) = W((\gamma'_g)^t)$$

and

$$W(\gamma_g^s) - W((\gamma')_g^t) = W(\gamma_f^s) \begin{pmatrix} 0 & \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} f \\ 0 & 0 \end{pmatrix} W(\eta).$$

Now, if we compute the differentials of X with respect to the two gardens Γ and Γ' , we see that the contributions only differ along these two tines. Therefore,

$$(\tilde{\partial}_{G,\Gamma'}^{\mathcal{B}^+} - \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})X = W((\gamma')_f) + W((\gamma')_g) - W(\gamma_f) - W(\gamma_g).$$

But we have by our relations above that

$$\begin{aligned} W((\gamma')_f) - W(\gamma_f) &= W((\gamma')_f^s) \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} W((\gamma')_f^t) - W(\gamma_f^s) \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} W(\gamma_f^t) \\ &= W(\gamma_f^s) \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} W(\eta) \begin{pmatrix} 0 & \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} g \\ 0 & 0 \end{pmatrix} W(\gamma_g^t). \end{aligned}$$

Similarly,

$$W((\gamma')_g) - W(\gamma_g) = -W(\gamma_f^s) \begin{pmatrix} 0 & \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} f \\ 0 & 0 \end{pmatrix} W(\eta) \begin{pmatrix} 0 & g \\ 0 & 0 \end{pmatrix} W(\gamma_g^t).$$

Adding these together, we find that

$$(\tilde{\partial}_{G,\Gamma'}^{\mathcal{B}^+} - \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})X = -\tilde{\partial}^{\mathcal{B}^+} W(\gamma_{f,g}^\Gamma),$$

where the differential on the right-hand side is given no subscript since it may be taken with respect to either Γ or Γ' , and where $\gamma_{f,g}^\Gamma$ is the path given by following γ_f^s , then passing through c_f , then following η , then passing through c_g , and finally following γ_g^t so that

$$W(\gamma_{f,g}^\Gamma) = W(\gamma_f^s) \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} W(\eta) \begin{pmatrix} 0 & g \\ 0 & 0 \end{pmatrix} W(\gamma_g^t).$$

This path is actually well defined up to homotopy, assuming that the garden has been fixed up to homotopy. Suppose that Γ is a garden, and suppose that γ_f is tine adjacent and to the left of γ_g (as we have been assuming). Then, $\gamma_{f,g}^\Gamma$ may be defined as the (unique homotopy class of) path from the base point to itself, with initial and tangent vectors positive with respect to the seed, which stays between γ_f and γ_g , and such that it passes through the centers c_f and c_g once in that order.

Our discussion has proved the following proposition.

Proposition 5.3. *Suppose that Γ and Γ' are gardens which differ from each up to homotopies fixing the seed by an application of move V as in Figure 8 applied to*

neighboring tines γ_f and γ_g in Γ . Let $\gamma_{f,g}^\Gamma$ be the path described above. Then, the graded algebra homomorphism $\Phi_\Gamma^{\Gamma'}: \tilde{\mathcal{B}}^+ \rightarrow \tilde{\mathcal{B}}^+$ defined on generators by

$$\begin{aligned}\Phi_\Gamma^{\Gamma'}(f) &= f, \\ \Phi_\Gamma^{\Gamma'}(X) &= X + W(\gamma_{f,g}^\Gamma)\end{aligned}$$

is a tame dg-isomorphism with respect to $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$ and $\tilde{\partial}_{G,\Gamma'}^{\mathcal{B}^+}$.

Proof. This is just a summary of our discussion. Tameness is clear, as this dg-isomorphism is just a composition of four elementary dg-automorphisms, one for each of the generators x, y, z , and w which have been collected in the matrix X . ■

Remark 5.4. Notice that move V comes with a handedness, in which we perform a clockwise half-twist on the centers. We notice that, after move V, the new tine γ'_g is now to the left of γ'_f , and one might be tempted to say that the inverse of move V is given by taking $X \mapsto X + W(\gamma_{g,f}^{\Gamma'})$. But this is not correct, since this second application of move V actually yields a different garden, up to homotopy, given by a full twist of the two centers. Instead, we artificially define the dg-isomorphism for the counter-clockwise twist to just be the inverse map. In particular, it has the slightly funky form as follows: suppose that Γ' is given from Γ by an inverse application of move V, with a counter-clockwise half twist. If f is to the left of g in Γ , then g is to the left of f in Γ' , and our dg-isomorphism takes $X \mapsto X - W(\gamma_{f,g}^{\Gamma'})$. It is worth noting that $\gamma_{f,g}^{\Gamma'}$ may be defined with respect to Γ as well: it is the (unique homotopy class of) path from the base point to itself, positive with respect to the seed, which remains between the two tines γ'_g and γ'_f , and crosses the faces c_f and c_g once in that order. Notice here that it crosses c_f first, but now c_f is to the right of c_g .

Proposition 5.3 demonstrates that move V preserves the tame dg-isomorphism type of our dg-algebra. But actually, there is more structure. Suppose that we have trivalent plane graph G . Notice that, up to isotopy, the centering and web of any two gardens are the same, and hence, we may assume this data is fixed. We will also fix an underlying orientation of the edges of G , as we have described the effect of changing this orientation in Section 5.1. Let $\mathcal{G}(G)$ denote the space of (not necessarily non-degenerate) gardens for this fixed centering, web, and orientation, all of which we omit from the notation. If $\xi \in S^*S^2$ is a possible seed, meaning that its base point s is not on the vertices, edges, centering, or web, then denote by $\mathcal{G}(G; \xi)$ the space of gardens with this fixed centering, web, and orientation, but with ξ as its seed. We may realize this as the fiber of a Serre fibration.

Let $\mathcal{S}_n(\xi)$ denote the space of n -tuples of embedded loops on S^2 starting and ending at ξ with initial and terminal tangent vectors positive with respect to ξ such that the strands are disjoint except at the base point s and with a chosen point on

each loop other than s . We note that $\mathcal{S}_1(\xi)$ is contractible with respect to the C^∞ topology; cf. the work of Eliashberg and Polterovich [14, Proposition 1.4.A] for the contractibility of a slightly different space, from which contractibility of our $\mathcal{S}_1(\xi)$ easily follows. (Their space is a little more rigid than ours near the seed and does not include the choice of point.) Furthermore, for each element of $\mathcal{S}_n(\xi)$, because S^2 is oriented, the loops are ordered, and we note that the map $\mathcal{S}_{n+1}(\xi) \rightarrow \mathcal{S}_n(\xi)$ given by forgetting the first loop is a Serre fibration with fiber given by a space equivalent to $\mathcal{S}_1(\xi)$. Hence, inductively, each $\mathcal{S}_n(\xi)$ is contractible. Additionally, for any such element of $\mathcal{S}_n(\xi)$, we may extract from it an (unordered) n -tuple of points, one for each loop. This is a Serre fibration

$$\pi: \mathcal{S}_n(\xi) \rightarrow \text{UConf}_n(S^2 \setminus \{s\})$$

to the n th unordered configuration space. When $n = g + 3$, the space $\mathcal{G}(G; \xi)$ is just the fiber over the centering. Hence, from the long exact sequence of homotopy groups, and since $\mathcal{S}_n(\xi)$ is contractible, we find that the connected components of $\mathcal{G}(G; \xi)$ are in bijection with the braid group on $g + 3$ elements:

$$\pi_0(\mathcal{G}(G; \xi)) = \pi_1(\text{UConf}_{g+3}(S^2 \setminus \{s\})) \cong \text{Br}_{g+3}.$$

As moves I–IV leave the differential invariant, we see that the dg-algebra is constant on each connected component of $\mathcal{G}(G; \xi)$. On the other hand, we see that move V consists of taking two neighboring centers and twisting them around each other; this is precisely a braid group element. That is, $\pi_0(\mathcal{G}(G; \xi))$ naturally forms a Br_{g+3} -torsor.

To be more explicit, recall that, as a group, the braid group on n strands has the natural presentation

$$\text{Br}_n = \left\langle \sigma_1, \dots, \sigma_{n-1} \mid \begin{array}{ll} \sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}, & 1 \leq i \leq n-2 \\ \sigma_i \sigma_j = \sigma_j \sigma_i, & |i - j| \geq 2 \end{array} \right\rangle.$$

Suppose that we have fixed a garden with seed ξ . Then, the tines may be ordered $f_1, \dots, f_{g+3}$ based upon their initial tangent vectors at the base point. The element σ_i in Br_{g+3} then acts via move V applied to the tines γ_{f_i} and $\gamma_{f_{i+1}}$. In Proposition 5.3, we have already discussed how move V acts by dg-isomorphism on the dg-algebra. In fact, these dg-isomorphisms are coherent, in the sense that the braid relations are satisfied up to exact tame dg-isomorphisms, as we now verify.

Let F_{g+2} be the free group on $g + 2$ generators, denoted by $\tilde{\sigma}_1, \dots, \tilde{\sigma}_{g+2}$. Let

$$\pi: F_{g+2} \twoheadrightarrow \text{Br}_{g+3}$$

be the natural projection, with $\pi(\tilde{\sigma}_k) = \sigma_k$. Then, the action of Br_{g+3} on $\pi_0(\mathcal{G}(G; \xi))$ yields in turn an action of F_{g+2} . Under this action, we have the following lemma.

Lemma 5.5. *Suppose that $\Gamma \in \pi_0(\mathcal{G}(G; \xi))$ and $\zeta \in F_{g+2}$. Then, there exist tame dg-isomorphisms*

$$\Phi_\Gamma^\zeta: (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\xi,\Gamma}^+)$$

such that

$$\Phi_\Gamma^1 = \mathbb{1}$$

and

$$\Phi_\Gamma^{\zeta\theta} = \Phi_{\theta,\Gamma}^\zeta \circ \Phi_\Gamma^\theta.$$

Furthermore, if $\pi(\zeta) = \mathbb{1} \in \text{Br}_{g+3}$, then Φ_Γ^ζ is homotopic to the identity and hence induces the identity on homology.

Proof. The construction of the dg-isomorphisms Φ_Γ^ζ is clear; the element $\tilde{\sigma}_k$ acts via the action of move V as applied to the faces f_k and f_{k+1} via the formula of Proposition 5.3. The composition rule is given by construction since we are using a free group and have defined the action on a minimal generating set.

The kernel K of the projection $\pi: F_{g+2} \twoheadrightarrow \text{Br}_{g+3}$ is the normal subgroup generated by the braid relations. We check the extent to which the braid relations are satisfied by the morphisms constructed by move V. Notice that it suffices to prove that Φ_Γ^ζ is homotopic to the identity for these braid relations which normally generate K . Indeed, suppose that for ζ one such generator we have Ψ_Γ^ζ as our desired chain homotopy, i.e., with

$$\Psi_\Gamma^\zeta \circ \tilde{\partial}_{G,\Gamma}^+ - \tilde{\partial}_{G,\xi,\Gamma}^+ \circ \Psi_\Gamma^\zeta = \mathbb{1} - \Phi_\Gamma^\zeta.$$

Then, for $\zeta, \theta \in K$, we obtain a chain homotopy for $\Phi_\Gamma^{\zeta\theta^{-1}}$ by taking

$$\Psi_\Gamma^{\zeta\theta^{-1}} = (\Psi_{\theta^{-1},\Gamma}^\zeta - \Psi_\Gamma^\theta) \circ \Phi_{\theta^{-1},\Gamma}^{\theta^{-1}},$$

and if $\theta \in F_{g+2}$ and $\zeta \in K$, then we may take

$$\Psi_\Gamma^{\theta^{-1}\zeta\theta} = (\Phi_\Gamma^\theta)^{-1} \Psi_{\theta,\Gamma}^\zeta \Phi_\Gamma^\theta.$$

First, consider the braid relations $\sigma_i \sigma_j = \sigma_j \sigma_i$ for $|i - j| \geq 2$. Notice that move V applied to f_i and f_{i+1} does not affect $\gamma_{f_j, f_{j+1}}$ at all, and vice versa. In other words,

$$\begin{aligned} \gamma_{f_j, f_{j+1}}^{\sigma_i \cdot \Gamma} &= \gamma_{f_j, f_{j+1}}^\Gamma, \\ \gamma_{f_i, f_{i+1}}^{\sigma_j \cdot \Gamma} &= \gamma_{f_i, f_{i+1}}^\Gamma. \end{aligned}$$

We find that both $\tilde{\sigma}_i \tilde{\sigma}_j$ and $\tilde{\sigma}_j \tilde{\sigma}_i$ therefore both act on the degree 2 generators by

$$X \mapsto X + W(\gamma_{f_i, f_{i+1}}^\Gamma) + W(\gamma_{f_j, f_{j+1}}^\Gamma).$$




Figure 11. All of the tines are oriented upwards, and the curves η_1 and η_2 are oriented from left to right.

Meanwhile, they act by the identity on the degree 1 generators. Hence, since they act the same way on all generators, they are identical. In particular, we have

$$\Phi_{\Gamma}^{\tilde{\sigma}_i \tilde{\sigma}_j (\tilde{\sigma}_j \tilde{\sigma}_i)^{-1}} = 1.$$

The second braid relation $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$ is a little more tricky, and we will need to choose nonzero Ψ . To simplify notation, we will suppose that $i = 1$. The argument is the same for any other triple of adjacent tines, so this may be assumed without loss of generality. The tines will simply be referred to as γ_1 , γ_2 , and γ_3 passing through the centers c_1 , c_2 , and c_3 , respectively, and we will use paths η_1 from c_1 to c_2 and η_2 from c_2 to c_3 completely contained between neighboring tines. As usual, for any tine, we use the decomposition

$$\gamma_f = \gamma_f^s * \gamma_f^c * \gamma_f^t.$$

See Figure 11.

Let us consider the action of $\sigma_1 \sigma_2 \sigma_1$. As discussed, the action of σ_1 is given on the generators x , y , z , and w , by adding the corresponding associated words for $\gamma_{1,2}$, which goes along γ_1^s through f_1 , then along η_1 , through f_2 , and ends with γ_2^t . Once we have performed this operation, we are now in a position to apply σ_2 , as in Figure 12.

We see in this case that the connecting curve goes under the second center f_2 . By homotoping the curves using moves I–IV, we see that move V acts at this stage by the word associated to $\gamma_1^s * (\eta_{1,2}^D) * \gamma_3^t$, where the superscript D denotes that the curve goes down under f_2 . Finally, after this application of move V, the picture now looks




Figure 12. After performing move V according to σ_1 , we are now in a position to apply move V along the orange curve, i.e., by σ_2 .




Figure 13. The diagram after an application of σ_i followed by an application of σ_{i+1} . The left and right figures are the same up to a number of applications of moves I and II. A final application of σ_i may now be performed by considering move V along the orange curve of the right diagram.

as in the left diagram of Figure 13, which we homotope to the right diagram. In total, we find that $\tilde{\sigma}_1 \tilde{\sigma}_2 \tilde{\sigma}_1$ acts via

$$X \mapsto X + W(A) + W(B) + W(C),$$

where A , B , and C are the curves in the top row of Figure 14. Similarly, we find that $\tilde{\sigma}_2 \tilde{\sigma}_1 \tilde{\sigma}_2$ acts via

$$X \mapsto X + W(A') + W(B') + W(C')$$

and similarly on the other generators, where A' , B' , and C' are on the bottom row of Figure 14. Up to homotopy through moves I–IV, the following can be seen:

- A and C' are the same curve, except that C' is pushed to go around f_3 to the right, whereas A goes around it to the left,
- B and B' are the same curve, except that B' is pushed to go over f_2 , whereas B goes under it,




Figure 14. The curves corresponding to the words added by the action of $\sigma_1 \sigma_2 \sigma_1$, in the top row, and $\sigma_2 \sigma_1 \sigma_2$, in the bottom row.

- C and A' are the same, except that C is pushed to go around f_1 to the left, whereas A' goes around it to the right.

We described in Lemma 4.10 how pushing a non-degenerate curve through a center changes the associated word by an exact element, with primitive involving the face through which the curve is pushed. One easily sees that

$$(W(A') + W(B') + W(C')) - (W(A) + W(B) + W(C)) = \tilde{\partial}^{\mathcal{B}^+} W(\gamma_{f_1, f_2, f_3}^\Gamma),$$

where $\gamma_{f_1, f_2, f_3}^\Gamma$ is the path which follows γ_1^s , then passes through f_1 , then follows η_1 , then passes through f_2 , then follows η_2 , then passes through f_3 , and finally traverses γ_3^t . (Notice that we have removed the garden from the notation on the right since for all such choices, the differential acts in the same way on the matrix of words, since the words only involve the degree 1 face generators.) In other words, we have

$$(\Phi_{\Gamma}^{\tilde{\sigma}_2 \tilde{\sigma}_1 \tilde{\sigma}_2} - \Phi_{\Gamma}^{\tilde{\sigma}_1 \tilde{\sigma}_2 \tilde{\sigma}_1})(X) = \tilde{\partial}^{\mathcal{B}^+} W(\gamma_{f_1, f_2, f_3}^\Gamma).$$

Applying $(\Phi_{\Gamma}^{\tilde{\sigma}_1 \tilde{\sigma}_2 \tilde{\sigma}_1})^{-1}$, which acts trivially on the right-hand side since degree 1 generators are unharmed, yields

$$\Phi_{\Gamma}^{\tilde{\sigma}_1^{-1} \tilde{\sigma}_2^{-1} \tilde{\sigma}_1^{-1} \tilde{\sigma}_2 \tilde{\sigma}_1 \tilde{\sigma}_2}(X) = X + \tilde{\partial}^{\mathcal{B}^+} W(\gamma_{f_1, f_2, f_3}^\Gamma).$$

This is homotopic to the identity because it satisfies the hypotheses of Proposition 2.5. ■

Remark 5.6. From the proof, we see that everywhere we use the action of the free group F_{g+2} , we could instead use the action of the quotient by the normal subgroup generated by the commutators $[\sigma_i, \sigma_j]$ for $|i - j| \geq 2$.

5.3. The full twist

So far, we have understood the path components $\pi_0(\mathcal{G}(G; \xi))$ of gardens where we have fixed the seed as a Br_{g+3} -torsor. Let us now study $\pi_0(\mathcal{G}(G))$. (Again, implicitly, we have fixed an orientation in this discussion.) We have a Serre fibration

$$\pi: \mathcal{G}(G) \rightarrow S^*(S^2 \setminus (G \cup W))$$

given by projecting to the seed. Here, G is the union of the vertices and edges of the graph, and W (for *web*) consists of the threads and centers. By definition, $\pi^{-1}(\xi) = \mathcal{G}(G; \xi)$. On the other hand, we have $S^2 \setminus (G \cup W)$ consisting of a union of open triangles, each with one vertex at a center opposite an edge of the triangle which is just an edge of the graph, and two vertices on the vertices of the graph, each opposite an edge lying along a thread of the web. Since each edge of the graph has a triangle on either side of it, there are $2|E| = 6g + 6$ such triangles. Let $\text{Tri}(G)$ be the collection of these triangles, just a finite set with $6g + 6$ elements. Let $\nu \in \text{Tri}(G)$ be one such triangle. Then, we obtain a sequence

$$\pi_1(S^*\nu; \xi) \rightarrow \pi_0(\mathcal{G}(G; \xi)) \rightarrow \pi_0(\mathcal{G}(G))$$

which is exact at the central term. Notice that $S^*\nu \cong \nu \times S^1 \sim S^1$, so $\pi_1(S^*\nu; \xi) \cong \mathbb{Z}$. Geometrically, this represents the free action given by twisting ξ by a full 360 degrees. This $\mathbb{Z}$ -action on $\pi_0(\mathcal{G}(G; \xi))$ corresponds to the full braid twist in Br_{g+3} , which is the generator of the center of Br_{g+3} . We hence find that $\pi_0(\mathcal{G}(G))$ consists of a collection of $6g + 6$ $\text{Br}_{g+3}/\mathbb{Z}$ -torsors, indexed by $\text{Tri}(G)$ based on in which triangle the base point lies. We will study $\text{Tri}(G)$ further in Section 5.4.

Even though the full twist acts non-trivially on $\pi_0(\mathcal{G}(G; \xi))$, it acts trivially on $\pi_0(\mathcal{G})$. Comparing with Theorem 1.6, we must therefore prove that the full twist is homotopic to the identity.

Lemma 5.7. *Suppose that G is a trivalent plane graph with garden $\Gamma \in \pi_0(\mathcal{G}(G; \xi))$. Then, any lift $\tilde{\Delta} \in F_{g+2}$ of the full twist $\Delta = (\sigma_{g+2} \cdots \sigma_1)^{g+3} \in \text{Br}_{g+3}$ acts on $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+)$ via a tame dg-automorphism homotopic to the identity.*

Proof. It suffices to prove the result for one lift, since a composition of dg-automorphisms homotopic to the identity is itself homotopic to the identity, and any two lifts differ by a dg-automorphism homotopic to the identity by Lemma 5.5. We use the lift

$$\tilde{\Delta} = (\tilde{\sigma}_{g+2} \cdots \tilde{\sigma}_1)^{g+3}.$$

Let $\tilde{C} = \tilde{\sigma}_{g+2} \cdots \tilde{\sigma}_1$ be a lift of $C \in \text{Br}_{g+3}$, so $\tilde{\Delta} = \tilde{C}^{g+3}$. Identifying

$$\text{Br}_{g+3} \cong \pi_1(\text{UConf}_{g+3}(S^2 \setminus \{s\})),$$




Figure 15. The action of the cyclical twist C on the tines. In general, there may be more tines, but we only draw three for convenience. The disk $\mathbb{D}_-$ is shaded in gray, and the centers occur in the order from left to right across the top of the disk. The seed is off at infinity. We have drawn the action of C so that it is supported in a neighborhood of the circle separating $\mathbb{D}_+$ and $\mathbb{D}_-$.

we see C acts on the garden as follows. Up to homotopy, there is a unique decomposition $S^2 = \mathbb{D}_+ \cup_{S^1} \mathbb{D}_-$ into two disks with the property that $\mathbb{D}_+$ contains the seed, and that the centers appear along $S^1 = \mathbb{D}_+ \cap \mathbb{D}_-$ in order according to the boundary orientation of $\mathbb{D}_+$. Then, $\tilde{C}$ acts on the tines via a homotopy supported near the gluing S^1 so that the centers are rotated from one to the next. See Figure 15.

Let us consider the action of $\tilde{C}$ on the dg-algebra. Notice that curves A and B in Figure 14 already correspond to the curves $\gamma_{f,g}$ involved in $\tilde{\sigma}_1$ followed by $\tilde{\sigma}_2$, respectively. It is not hard to see how this pattern continues as we apply $\tilde{C}$: $\tilde{\sigma}_k$ corresponds to the curve which goes through f_1 , under f_2 through f_k and then through f_{k+1} . Let us denote these curves by $\gamma_{1,k+1}^\Gamma$.

More generally, we may define curves $\gamma_{i,j}^\Gamma$ as the homotopy class of loop based at the base point of Γ given by first following a path from the base point to c_i through $\mathbb{D}_+$, then a path from c_i to c_j through $\mathbb{D}_-$, and then a path from c_j back to the base point through $\mathbb{D}_+$. Notice that this description completely determines a unique curve up to homotopy since $\mathbb{D}_+$ and $\mathbb{D}_-$ are contractible. Further, notice that when $i = 1$ and $j = k + 1$, this matches the description above. Finally, up to homotopy, we have simply

$$\gamma_{i,j}^\Gamma \sim (\gamma_i^\Gamma)^{-1} * \gamma_j^\Gamma,$$

where γ_k^Γ is the k th tine. The curves $\gamma_{i,j}^\Gamma$ are shown in Figure 16.

Notice that, either by inspection of Figure 15 or by noting that the decomposition $S^2 = \mathbb{D}_+ \cup \mathbb{D}_-$ is compatible with both Γ and $C \cdot \Gamma$, we have

$$\gamma_{i,j}^{C \cdot \Gamma} = \gamma_{i+1,j+1}^\Gamma,$$




Figure 16. The curves $\gamma_{i,j}^\Gamma$, where $\mathbb{D}_-$ is highlighted in gray. In this case, imagining the centers are labeled in this case from 1 to 5 from left to right, then the curve depicted is either $\gamma_{2,5}^\Gamma$ or $\gamma_{5,2}^\Gamma$, depending upon the orientation of the curve. Here, we have moved the base point in from infinity for a slightly different perspective.

with the convention that the indices are taken modulo $g + 3$. In fact, explicitly on tines, one sees that

$$\gamma_i^{C \cdot \Gamma} = \eta_1^{-1} * \gamma_{i+1}^\Gamma,$$

where η_1 is a loop lying between γ_1^Γ and γ_2^Γ (as in Proposition 4.11).

On the other hand, the action of $\tilde{C}$ on the degree 2 generators is just

$$\tilde{C} \cdot X = X + \sum_{k=1}^{g+2} W(\gamma_{1,k+1}^\Gamma).$$

Hence, applying $\tilde{C}$ a total of $g + 3$ times, we find that

$$\begin{aligned} \tilde{C}^{g+3} \cdot X &= X + \sum_{j=0}^{g+2} \sum_{k=2}^{g+3} W(\gamma_{1,k}^{C^j \cdot \Gamma}) \\ &= \sum_{\substack{1 \leq i, j \leq g+3 \\ i \neq j}} W(\gamma_{i,j}^\Gamma) \\ &= \sum_{\substack{1 \leq i, j \leq g+3 \\ i \neq j}} W((\gamma_i^\Gamma)^{-1}) \cdot W(\gamma_j^\Gamma). \end{aligned}$$

But even further, notice that if $i = j$, then $\gamma_{i,i}$ can be homotoped so that it hits c_{f_i} twice in a row, without intersecting a thread or edge in between. Hence, $W(\gamma_{i,i}) = 0$, since no binary sequence can switch from 0 to 1 at both points it intersects c_{f_i} since it does not have a chance to go from 1 to 0 in between. Hence, we may more simply write

$$\tilde{\mathcal{C}}^{g+3} \cdot X = X + \left(\sum_{i=1}^{g+3} W(\gamma_i^{-1}) \right) \cdot \left(\sum_{i=1}^{g+3} W(\gamma_j) \right).$$

By Proposition 2.5, we see that the action of $\tilde{\Delta}$ satisfies the required conditions so long as we can prove that $(\sum_{i=1}^{g+3} W(\gamma_i^{-1})) \cdot (\sum_{j=1}^{g+3} W(\gamma_j))$ is exact. To this end, let $\bar{\Gamma}$ be the garden obtained from Γ by switching the orientation of all the tines simultaneously and negating ξ . Then, since the differentials of Γ and $\bar{\Gamma}$ are the same on faces, and hence on matrices of words,

$$\begin{aligned} \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} \left(\sum_{i=1}^{g+3} W(\gamma_i^{-1}) \right) &= \tilde{\partial}_{G,\bar{\Gamma}}^{\mathcal{B}^+} \left(\sum_{i=1}^{g+3} W(\gamma_i^{-1}) \right) \\ &= (\tilde{\partial}_{G,\bar{\Gamma}}^{\mathcal{B}^+})^2 X \\ &= 0. \end{aligned}$$

On the other hand, $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} X = \sum_{j=1}^{g+3} W(\gamma_j)$. Hence, we have

$$\left(\sum_{i=1}^{g+3} W(\gamma_i^{-1}) \right) \cdot \left(\sum_{j=1}^{g+3} W(\gamma_j) \right) = \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} \left[\left(\sum_{i=1}^{g+3} W(\gamma_i^{-1}) \right) \cdot X \right]$$

being exact as desired. ■

5.4. Invariance under moves VI and VII

Recall that moves VI and VII pass a seed through a thread or edge, respectively, as in Figure 9. We now construct tame dg-automorphisms which intertwine the respective differentials for the gardens in two separate lemmas.

Lemma 5.8. *Suppose that Γ and Γ' are two gardens which differ by move VI as in the left side of Figure 9, and suppose that τ is the thread appearing in the diagram. Let σ be $+1$ if the thread is oriented upwards and -1 if the thread is oriented downwards, and let*

$$\tilde{H}(\tau) = \sigma \cdot H(\tau)$$

be the monomial associated to the thread, with an appropriate sign. Then, we have a tame dg-automorphism

$$\Phi_{\Gamma'}^{\Gamma'} : (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma'}^{\mathcal{B}^+})$$

defined on generators by

$$\begin{aligned}\Phi_{\Gamma}^{\Gamma'}(f) &= f, \\ \Phi_{\Gamma}^{\Gamma'}(X) &= \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} X \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}.\end{aligned}$$

Proof. Suppose that γ is any curve on S^2 from the base point s back to itself with initial and final tangent vectors matching the coorientation of the seed ξ , and suppose that the thread passes through ξ according to move VI. Notice that before the thread passes through, we may write

$$\gamma = \gamma^{\tau} * \gamma',$$

whereas after, we have

$$\gamma = \gamma' * \gamma^{\tau},$$

up to homotopy of endpoints not interacting with the garden Γ in the first case and with the garden Γ' in the second. In particular, in this decomposition, the binary sequences associated to γ' and γ^{τ} are identified, so we may write $W(\gamma')$ and $W(\gamma^{\tau})$ without respect to the garden, even though we must write $W^{\Gamma}(\gamma)$ and $W^{\Gamma'}(\gamma)$ accordingly. The matrix of words for γ^{τ} is

$$W(\gamma^{\tau}) = \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}.$$

Therefore, we find

$$W^{\Gamma}(\gamma) = \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} W(\gamma'), \quad W^{\Gamma'}(\gamma) = W(\gamma') \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}.$$

Since $\begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}$, we arrive at

$$W^{\Gamma}(\gamma) = \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} W^{\Gamma'}(\gamma) \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}.$$

Summing these relations over all tines γ_f , we find the desired result

$$\begin{aligned}\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}(X) &= \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} \cdot (\tilde{\partial}_{G,\Gamma'}^{\mathcal{B}^+}(X)) \cdot \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} \\ &= \tilde{\partial}_{G,\Gamma'}^{\mathcal{B}^+} \left(\begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} \cdot X \cdot \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} \right).\end{aligned}$$

This verifies that $\Phi_{\Gamma}^{\Gamma'}$, as in the statement in the lemma, is a dg-homomorphism. Writing these expressions out in terms of x , y , z , and w , we have

$$\begin{cases} \Phi(x) = x - y \cdot \tilde{H}(\tau), \\ \Phi(y) = y, \\ \Phi(z) = z + \tilde{H}(\tau) \cdot y, \\ \Phi(w) = w + \tilde{H}(\tau) \cdot x - \tilde{H}(\tau) \cdot y \cdot \tilde{H}(\tau) - z \cdot \tilde{H}(\tau), \end{cases}$$

from which it follows that the constructed dg-isomorphism is tame, given as a composition of four elementary dg-automorphisms. ■

Lemma 5.9. *Suppose that Γ and Γ' are two gardens which differ by move VII as in the right side of Figure 9, and suppose that e is the edge appearing in the diagram. Let $\circ$ be the orientation of the edge upwards (regardless of the orientation induced by the garden), and set $\lambda := \lambda_{e,\circ}$ and $\bar{\lambda} := \lambda_{e,\bar{\circ}}$. Then, we have a regeneration*

$$\Phi_{\Gamma}^{\Gamma'}: (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma'}^+)$$

defined on generators by

$$\begin{aligned} \Phi_{\Gamma}^{\Gamma'}(f) &= f, \\ \Phi_{\Gamma}^{\Gamma'}(X) &= \begin{pmatrix} 0 & -\lambda^{-1} \\ \bar{\lambda} & 0 \end{pmatrix} X \begin{pmatrix} 0 & \bar{\lambda}^{-1} \\ -\lambda & 0 \end{pmatrix}. \end{aligned}$$

Proof. The proof is the same as in Lemma 5.8, but where now we use γ^e instead of γ^{τ} , with

$$W(\gamma^e) = \begin{pmatrix} 0 & \pm\lambda^{-1} \\ \mp\bar{\lambda} & 0 \end{pmatrix},$$

where the signs depend upon whether the orientation of the edge in the garden points upwards or downwards. Either choice of sign yields the same end result. ■

Notice that, in the proof of Lemmas 5.8 and 5.9, the key point is that the formulas for X actually hold for any non-degenerate path from the seed back to itself. This is useful for the following reason. Suppose that Γ is a garden, and suppose that Γ' is a resulting garden after move VI; move VII is handled similarly. Consider the application of move V, according to $\tilde{\sigma}_i$, on either Γ or Γ' . Then, we have

$$W(\gamma_{f_i, f_{i+1}}^{\Gamma}) = \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} W(\gamma_{f_i, f_{i+1}}^{\Gamma'}) \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}.$$

It follows, regardless of whether we apply move V followed by move VI or move VI

followed by move V, the image of the matrix of generators X is given by the left- and right-hand sides of the following equality:

$$\begin{aligned} & \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} X \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} + W(\gamma_{f_i, f_{i+1}}^\Gamma), \\ &= \begin{pmatrix} 1 & \tilde{H}(\tau) \\ 0 & 1 \end{pmatrix} (X + W(\gamma_{f_i, f_{i+1}}^{\Gamma'})) \begin{pmatrix} 1 & -\tilde{H}(\tau) \\ 0 & 1 \end{pmatrix}. \end{aligned}$$

In other words, the dg-isomorphisms by which moves V and VI act commute. By the same reasoning, moves V and VII also commute. Hence, we may essentially study moves VI and VII independently from move V, which we already understand quite well, as summarized by Lemma 5.5.

Recall from Section 5.3 that $\pi_0(\mathcal{G}(G))$ is given by a collection of $\text{Br}_{g+3}/\mathbb{Z}$ -torsors over the collection $\text{Tri}(G)$ of triangles obtained by cutting S^2 along the graph and the web (which we have fixed up to isotopy, since the web is a contractible choice). Moves VI and VII are precisely the moves which allow us to move from one triangle to the next. In fact, suppose that triangles $\nu_1, \nu_2 \in \text{Tri}(G)$ are adjacent to each other, sharing an edge which is either an edge of the graph or a thread. If we have a garden Γ with base point in ν_1 , then we may homotope all of the tines (via moves I–IV) so that they all pass through the edge shared between ν_1 and ν_2 before interacting with any other edges, threads, or centers. Then, we can push the seed from ν_1 to ν_2 using move VI or VII accordingly. Since these moves commute with move V, this yields a bijection between the corresponding $\text{Br}_{g+3}/\mathbb{Z}$ -torsors of homotopy classes of gardens. In fact, this bijection is completely canonical, irrespective of how we applied moves I–IV in order to arrive at a situation in which we could push the seed through an edge or thread via move VI or VII. It so happens also that the formulas of Lemmas 5.8 and 5.9 only depend upon the orientation through which the seed is pushed. Hence, if Γ_1 is a homotopy class of garden with base point on ν_1 and Γ_2 is the corresponding homotopy class of garden with base point on ν_2 given by an application of move VI or VII, then the corresponding formula for move VI or VII is a canonical tame dg-automorphism or regeneration from $(\tilde{\mathcal{B}}^+, \tilde{\partial}_{G, \Gamma_1}^{\mathcal{B}^+})$ to $(\tilde{\mathcal{B}}^+, \tilde{\partial}_{G, \Gamma_2}^{\mathcal{B}^+})$.

In this way, for each triangle $\nu \in \text{Tri}(G)$ and each of its edges (which is either an edge of G or a thread in the web), we may associate a dg-automorphism corresponding to moving the seed through that edge. Furthermore, the triangle comes with a canonical orientation, and hence, we can distinguish canonically between the two edges of the triangle which are actually threads. We write L, R, S , for *left*, *right*, and *straight*, to denote the corresponding move VI and VII actions, where the labels L, R , and S appear in the order of the natural boundary orientation of the triangle, and where S is placed on the unique edge of the triangle which is an edge of the graph. In this way, we naturally find an $F_3 = \langle L, R, S \rangle$ -action on gardens up to homotopy

which lift canonically to regenerative tame dg-automorphisms. That is, for each word $\zeta \in F_3$ and $\Gamma \in \pi_0(\mathcal{G}(G))$, we obtain canonical regenerative tame dg-automorphisms

$$\Phi_{\Gamma}^{\zeta} : (\tilde{\mathcal{B}}_{\Gamma}^+, \tilde{\partial}_{\Gamma}^{\mathcal{B}^+}) \rightarrow (\tilde{\mathcal{B}}_{\Gamma}^+, \tilde{\partial}_{\Gamma, \zeta}^{\mathcal{B}^+})$$

which compose in the obvious way

$$\Phi_{\Gamma}^{\xi\theta} = \Phi_{\theta, \Gamma}^{\xi} \circ \Phi_{\Gamma}^{\theta}.$$

As is suggested in Theorem 1.6, we would like to understand the stabilizer of the action on F_3 on gardens. If we push a seed back and forth through the same edge, then we come back to the same garden up to homotopy, i.e., $\Gamma = LR \cdot \Gamma = S^2 \cdot \Gamma = RL \cdot \Gamma$. In fact, this is reflected at the level of dg-automorphisms.

Lemma 5.10. *We have that L and R act by inverses, and S acts as its own inverse, i.e.,*

$$\mathbb{1} = \Phi_{R\Gamma}^L \Phi_{\Gamma}^R = \Phi_{S\Gamma}^S \Phi_{\Gamma}^S = \Phi_{L\Gamma}^R \Phi_{\Gamma}^L.$$

Proof. This is obvious from Lemmas 5.8 and 5.9 and the fact that we pass the thread or edge back and forth with opposite orientation. ■

If a thread moves around a vertex of the graph, i.e., via the word $(SL)^3$, then the garden is again preserved up to homotopy, since the seed has just gone around a contractible loop in S^2 around a vertex. This is also reflected at the level of dg-automorphism.

Lemma 5.11. *We have*

$$\mathbb{1} = \Phi_{LSLSL\Gamma}^S \circ \Phi_{SLSL\Gamma}^L \circ \Phi_{LSL\Gamma}^S \circ \Phi_{SL\Gamma}^L \circ \Phi_{L\Gamma}^S \circ \Phi_{\Gamma}^L.$$

Proof. Upon applying $(SL)^3$, we traverse three threads and three edges. Let us label them $\tau_1, e_1, \tau_2, e_2, \tau_3, e_3$ in the order that they appear. Let $\lambda_k = A_{e_k}$ and $\bar{\lambda}_k = B_{e_k}$ with respect to the inward pointing orientation of the edge. Let

$$Z := \begin{pmatrix} 1 & H(\tau_1) \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & -\lambda_1^{-1} \\ \bar{\lambda}_1 & 0 \end{pmatrix} \begin{pmatrix} 1 & H(\tau_2) \\ 0 & 1 \end{pmatrix} \\ \cdot \begin{pmatrix} 0 & -\lambda_2^{-1} \\ \bar{\lambda}_2 & 0 \end{pmatrix} \begin{pmatrix} 1 & H(\tau_3) \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & -\lambda_3^{-1} \\ \bar{\lambda}_3 & 0 \end{pmatrix}.$$

The right-hand side of the desired equation preserves the degree 1 generators and acts on the degree 2 generators as

$$X \mapsto ZXZ^{-1}.$$

But notice that $Z = \pm W(\gamma_v)$, where γ_v is the loop going around a vertex once, and $W(\gamma_v) = I$, since γ_v can be contracted to a point under move IV. ■

Our action of $F_3 = \langle L, R, S \rangle$ is actually larger than we need, since even on the dg-algebra level, we have $LR = S^2 = (SL)^3 = 1$. We may therefore instead use the group

$$\mathbb{Z}_2 * \mathbb{Z}_3 \cong \langle S, (SL) \mid S^2, (SL)^3 \rangle \cong \langle L, R, S \mid LR, S^2, (SL)^3 \rangle.$$

In other words, we arrive at the following.

Lemma 5.12. *There is an action of $F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$ on homotopy classes of gardens (with the same orientation). For each homotopy class of garden Γ and element $\tau \in F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$, there exists a regenerative tame dg-isomorphism*

$$\Phi_\Gamma^\tau: (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\tau,\Gamma}^+).$$

These satisfy the composition property

$$\Phi_\Gamma^{\tau\tau'} = \Phi_{\tau',\Gamma}^{\tau'} \circ \Phi_\Gamma^\tau.$$

Proof. We have defined an action of $\mathbb{Z}_2 * \mathbb{Z}_3$ on the dg-algebras which commute with the action of F_{g+3} via move V. In other words, we may write

$$\Phi_\Gamma^\tau := \Phi_{\tau_2,\Gamma}^{\tau_1} \circ \Phi_\Gamma^{\tau_2} = \Phi_{\tau_1,\Gamma}^{\tau_2} \circ \Phi_\Gamma^{\tau_1},$$

where $\tau = (\tau_1, \tau_2) \in F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$. The commutativity allows us to group like terms. In other words, if $\tau = (\tau_1, \tau_2)$ and $\tau' = (\tau'_1, \tau'_2)$, then

$$\begin{aligned} \Phi_{\tau',\Gamma}^{\tau'} \circ \Phi_\Gamma^\tau &= \Phi_{(\tau_2\tau')\cdot\Gamma}^{\tau_1} \circ \Phi_{\tau',\Gamma}^{\tau_2} \circ \Phi_{\tau'_2,\Gamma}^{\tau'_1} \circ \Phi_\Gamma^{\tau'_2} \\ &= \Phi_{(\tau_2\tau'_1\tau'_2)\cdot\Gamma}^{\tau_1} \circ \Phi_{(\tau'_1\tau'_2)\cdot\Gamma}^{\tau_2} \circ \Phi_{\tau'_2,\Gamma}^{\tau'_1} \circ \Phi_\Gamma^{\tau'_2} \\ &= \Phi_{(\tau_1\tau'_1\tau'_2)\cdot\Gamma}^{\tau_2} \circ \Phi_{(\tau'_1\tau'_2)\cdot\Gamma}^{\tau_1} \circ \Phi_{\tau'_2,\Gamma}^{\tau'_1} \circ \Phi_\Gamma^{\tau'_2} \\ &= \Phi_{(\tau_1\tau'_1\tau'_2)\cdot\Gamma}^{\tau_2} \circ \Phi_{\tau'_2,\Gamma}^{\tau_1\tau'_1} \circ \Phi_\Gamma^{\tau'_2} \\ &= \Phi_{(\tau'_2\tau'_1\tau'_1)\cdot\Gamma}^{\tau_2} \circ \Phi_{(\tau_1\tau'_1)\cdot\Gamma}^{\tau'_2} \circ \Phi_\Gamma^{\tau_1\tau'_1} \\ &= \Phi_{(\tau_1\tau'_1)\cdot\Gamma}^{\tau_2\tau'_2} \circ \Phi_\Gamma^{\tau_1\tau'_1} \\ &= \Phi_\Gamma^{(\tau_1\tau'_1,\tau'_2\tau'_2)} \\ &= \Phi_\Gamma^{\tau\tau'}. \end{aligned}$$

■

One of the items in Theorem 1.6 asks us to prove that the stabilizer at Γ of the action of $\mathcal{H}(G)$ on homotopy classes of gardens acts in a manner homotopic to the identity. From our description, we are now at a point where we may understand this

stabilizer. We notice that since the orientation of the edges of G is invariant under homotopies, any element of the stabilizer is trivial in the factor $(\mathbb{Z}_2)^E$.

Suppose that $\Gamma \in \pi_0(\mathcal{G}(G))$, and let us consider its stabilizer. Suppose that we fix $s \in S^2$. Let C be the union of the centering of a garden, which we continue to assume that fixed. Then, there is an injective homomorphism

$$\phi: \pi_1(S^2 \setminus C; s) \rightarrow \mathbb{Z}_2 * \mathbb{Z}_3$$

as follows. If we choose a representative non-degenerate loop, then we may label it with a word in $F_3 = \langle L, R, S \rangle$ according to how it passes through threads and edges of the graph. Any two non-degenerate representatives of the same element in the fundamental group are equivalent up to the relations $LR = S^2 = (SL)^3 = 1$, as we can form a homotopy which cancels out going back and forth through the same edge or going around a vertex. The image of ϕ consists of those elements of $\mathbb{Z}_2 * \mathbb{Z}_3$ which preserve gardens with base point in the triangle in $\text{Tri}(G)$ containing s .

Notice that $S^2 \setminus C$ is homotopy equivalent to a wedge of $g + 2$ circles. Hence, $\pi_1(S^2 \setminus C; s) \cong F_{g+2}$. One could use a canonical set of generators given by the curves $\eta_1, \dots, \eta_{g+2}$ appearing in the proof of Proposition 4.11, with η_k lying between the tines γ_{f_k} and $\gamma_{f_{k+1}}$. One could study the action of moves VI and VII as applied along η_k , though we take an alternative approach.

Suppose that c_f is one of the centers in C , and suppose that the corresponding face f has n_f vertices, or equivalently, c_f has n_f triangles with c_f as a vertex. If $\nu \in \text{Tri}(G)$ has c_f as a vertex and Γ is a homotopy class of garden with base point in ν , then $L^{n_f} \cdot \Gamma$ is a garden again with base point in ν , though not the same garden. However, there will be some element $\tau_f \in \text{Br}_{g+3}/\mathbb{Z}$ such that $\tau_f L^{n_f} \cdot \Gamma = \Gamma$.

In our presentation, however, we are interested in studying a garden Γ with base point s which is not necessarily in a triangle with vertex c_f . Suppose that we choose a path from s to (some fixed point in) ν , where ν is a triangle with vertex c_f . The path corresponds to a word $P_f \in \mathbb{Z}_2 * \mathbb{Z}_3$. The loops based at s corresponding to $P_f L^{n_f} P_f^{-1}$ generate $\pi_1(S^2 \setminus C; s)$. (In fact, any collection of all but one of them are a free set of generators.) It follows, therefore, that if Γ is a garden with base point s , then its stabilizer in $\mathcal{H}(G)$ is

$$\text{Stab}_{\mathcal{H}(G)}(\Gamma) = \langle P_f \tilde{\tau}_f L^{n_f} P_f^{-1} \rangle,$$

the subgroup generated by the collection $\{P_f \tilde{\tau}_f L^{n_f} P_f^{-1}\}$ over all faces f and lifts $\tilde{\tau}_f \in F_{g+2}$ of $\tau_f \in \text{Br}_{g+3}/\mathbb{Z}$. (This is true regardless of the choices of paths P_f , precisely because the loops corresponding to $P_f L^{n_f} P_f^{-1}$ generate $\pi_1(S^2 \setminus C; s)$, and any other choice of P_f differs by left multiplication by some element of this group.) Our goal, cf. Theorem 1.6, is to prove that the various $P_f \tilde{\tau}_f L^{n_f} P_f^{-1}$ act $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+)$ via dg-automorphisms homotopic to the identity.

Lemma 5.13. *In the setup just described, the dg-automorphisms $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+)$ induced by the elements $P_f \tilde{\tau}_f L^{n_f} P_f^{-1}$ are homotopic to the identity.*

Proof. It suffices to prove that $\tilde{\tau}_f L^{n_f}$ is homotopic to the identity whenever Γ has base point in one of the triangles with vertex c_f , since conjugation by P_f does not affect this. In fact, for the same reason, by conjugating by a sequence of actions of move V, we may assume without loss of generality that Γ is such that the tine γ_f passing through c_f is the right-most tine exiting the seed. In other words, if we enumerate the faces $f_1, \dots, f_{g+3}$ according to the order of the tines, then $f = f_{g+3}$ is the specific face we are focusing on. We label the centers $c_1, \dots, c_{g+3}$ and the tines $\gamma_1, \dots, \gamma_{g+3}$ accordingly. We write n_{g+3} instead of $n_{f_{g+3}}$ and τ for $\tau_{f_{g+3}}$. Finally, we note that it suffices to prove the result for a single lift $\tilde{\tau}$ of τ , since any two lifts differ by a sequence of full twists, which acts by a dg-automorphism itself homotopic to the identity.

With these assumptions out of the way, we must first understand the element $\tau \in \text{Br}_{g+3}/\mathbb{Z}$. Consider a decomposition of $S^2 = \mathbb{D}_+ \cup_{S^1} \mathbb{D}_-$ into two disks, similar to that in Lemma 5.7, but now where both the base point and c_{g+3} are contained in the interior of $\mathbb{D}_+$, and c_1 through c_{g+2} now occur in order around the boundary of $\mathbb{D}_+$, as in Figure 17. Notice that if we apply L^{n_f} to the garden, it has the effect given by twisting $\mathbb{D}_+$ clockwise by a full twist. Correspondingly, τ must be a counter-clockwise twist of $\mathbb{D}_-$, since this undoes the clockwise twist of $\mathbb{D}_+$. That is,

$$\tau = (\sigma_{g+1} \cdots \sigma_2 \sigma_1)^{g+2}.$$

We may therefore use the lift

$$\tilde{\tau} := (\tilde{\sigma}_{g+1} \cdots \tilde{\sigma}_2 \tilde{\sigma}_1)^{g+2}.$$

Let $\tilde{D} := \tilde{\sigma}_{g+1} \cdots \tilde{\sigma}_2 \tilde{\sigma}_1$ so that $\tilde{\tau} = \tilde{D}^{g+2}$. Let $D = \pi(\tilde{\tau})$ so that $\tau = D^{g+2}$. Then, the action of $\tilde{\tau}$ is given by preserving the degree 1 generators and acting on the degree 2 generators by

$$\tilde{\tau} \cdot X = X + \sum_{j=0}^{g+1} \sum_{k=2}^{g+2} W(\gamma_{1,k}^{D^j \cdot \Gamma}),$$

where we have used the same notation $\gamma_{i,j}^\Gamma \sim (\gamma_i^\Gamma)^{-1} * \gamma_j^\Gamma$ as in the proof of Lemma 5.7. However, the sum is less simple here because D does not simply shift indices by 1. Indeed, by inspection of when we apply the action of D ,

$$\gamma_i^{D \cdot \Gamma} \sim \begin{cases} (\eta_1)^{-1} * \gamma_{i+1,j+1}^\Gamma, & i, j \neq g+2, g+3, \\ (\eta_1)^{-1} * \gamma_1^\Gamma * \eta_{g+2}, & i = g+2, \\ \gamma_{g+3}^\Gamma, & i = g+3, \end{cases}$$




Figure 17. Our setup for Lemma 5.13, where $\mathbb{D}_-$ is shaded in gray, and the center c_{g+3} is now in the interior of $\mathbb{D}_+$. For simplicity, we have drawn the situation $g = 1$ so that there are only four tines. The dotted black arrow indicates that $L^{n_{g+3}}$ acts by homotoping the seed around this loop (under the assumption that the base point is located in the face corresponding to f_{g+3}), which is realized by a full clockwise rotation of $\mathbb{D}_+$.

where η_1 is the path between γ_1^Γ and γ_2^Γ and η_{g+2} is the path between γ_{g+2}^Γ and γ_{g+3}^Γ , as in the proof of Proposition 4.11. (Notice that η_{g+2} appears in Figure 17 as the dotted black line.) It follows that the action of $\tilde{\tau}$ on the generators X may be written as

$$\tilde{\tau} \cdot X = X + \sum_{1 \leq i < j \leq g+2i \neq j} W(\gamma_{i,j}^\Gamma) + \sum_{1 \leq j < i \leq g+2} W(\gamma_{i,j}^\Gamma * \eta_{g+2}^{-1}).$$

On the other hand, the action of $L^{n_{g+3}}$ (as the seed moves counter-clockwise around c_{g+3}) is given by

$$L^{n_{g+3}} \cdot X = \begin{pmatrix} 1 & H_1 \\ 0 & 1 \end{pmatrix} \cdots \begin{pmatrix} 1 & H_{n_{g+3}} \\ 0 & 1 \end{pmatrix} X \begin{pmatrix} 1 & -H_{n_{g+3}} \\ 0 & 1 \end{pmatrix} \cdots \begin{pmatrix} 1 & -H_1 \\ 0 & 1 \end{pmatrix},$$

where H_k are the monomials associated to the threads emanating from c_{g+3} in order. (This formula holds for any Γ with base point in a triangle with vertex c_{g+3} .) But multiplying all of these out, this is just

$$L^{n_{g+3}} \cdot X = \begin{pmatrix} 1 & \tilde{\partial}^{\mathcal{B}^+} f_{g+3} \\ 0 & 1 \end{pmatrix} X \begin{pmatrix} 1 & -\tilde{\partial}^{\mathcal{B}^+} f_{g+3} \\ 0 & 1 \end{pmatrix}.$$

For simplicity, let us write

$$\begin{aligned} W_i &:= W(\gamma_i), \quad \bar{W}_i := W(\gamma_i^{-1}), \\ W &:= \sum_{i=1}^{g+3} W_i, \quad \bar{W} := \sum_{i=1}^{g+3} \bar{W}_i. \end{aligned}$$

We have already seen, in the course of the proof of Lemma 5.7, that

$$\begin{aligned} \tilde{\partial}^{\mathcal{B}^+} \bar{W} &= \tilde{\partial}^{\mathcal{B}^+} W = 0, \\ \tilde{\partial}^{\mathcal{B}^+} X &= W. \end{aligned}$$

Notice also that

$$W_{g+3} = \begin{pmatrix} 0 & f_{g+3} \\ 0 & 0 \end{pmatrix} = \bar{W}_{g+3}$$

since they both matrices of words associated to out-and-back paths from the seed to c_{g+3} , which lie in the same face. Finally, notice that

$$W(\eta_{g+2}) = \begin{pmatrix} 1 & -\tilde{\partial}^{\mathcal{B}^+} f_{g+3} \\ 0 & 1 \end{pmatrix} = I - \tilde{\partial}^{\mathcal{B}^+} W_{g+3}.$$

This allows us to notationally simplify both the actions of $\tilde{\tau}$ and $L^{n_{g+3}}$. For the action of $\tilde{\tau}$, notice that if we include terms with $i = j$, then those terms are 0. So, we have

$$\begin{aligned} \tilde{\tau} \cdot X &= X + \sum_{1 \leq i < j \leq g+2} \bar{W}_i W_j + \sum_{1 \leq j < i \leq g+2} \bar{W}_i W_j (I - \tilde{\partial}^{\mathcal{B}^+} W_{g+3}) \\ &= X + \left(\sum_{i=1}^{g+2} \bar{W}_i \right) \left(\sum_{j=1}^{g+2} W_j \right) - \sum_{1 \leq j < i \leq g+2} \bar{W}_i W_j (\tilde{\partial}^{\mathcal{B}^+} W_{g+3}) \\ &= X + (\bar{W} - W_{g+3})(W - W_{g+3}) + W_{g+3} W (\tilde{\partial}^{\mathcal{B}^+} W_{g+3}) \\ &\quad - \sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j (\tilde{\partial}^{\mathcal{B}^+} W_{g+3}) \\ &= X + \bar{W} W - \bar{W} W_{g+3} - W_{g+3} W + W_{g+3} W (\tilde{\partial}^{\mathcal{B}^+} W_{g+3}) \\ &\quad - \sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j (\tilde{\partial}^{\mathcal{B}^+} W_{g+3}), \end{aligned}$$

where the tricky term $W_{g+3}W(\tilde{\partial}^{\mathcal{B}^+}W_{g+3})$ arises because we include in our sum the case $i = g + 3$. The action of $L^{n_{g+3}}$ simplifies to

$$\begin{aligned} L^{n_{g+3}} \cdot X &= (I + \tilde{\partial}^{\mathcal{B}^+}W_{g+3})X(I - \tilde{\partial}^{\mathcal{B}^+}W_{g+3}) \\ &= X + (\tilde{\partial}^{\mathcal{B}^+}W_{g+3}) \cdot X - X(\tilde{\partial}^{\mathcal{B}^+}W_{g+3}) - (\tilde{\partial}^{\mathcal{B}^+}W_{g+3})X(\tilde{\partial}^{\mathcal{B}^+}W_{g+3}). \end{aligned}$$

Combining these actions, we have the following, upon reordering terms:

$$\begin{aligned} \tilde{\tau}L^{n_{g+3}} \cdot X - X &= \bar{W}W \\ &\quad + (\tilde{\partial}^{\mathcal{B}^+}W_{g+3}) \cdot X - W_{g+3} \cdot W \\ &\quad + W_{g+3}W(\tilde{\partial}^{\mathcal{B}^+}W_{g+3}) - (\tilde{\partial}^{\mathcal{B}^+}W_{g+3})X(\tilde{\partial}^{\mathcal{B}^+}W_{g+3}) \\ &\quad - \left(X + \sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j\right)(\tilde{\partial}^{\mathcal{B}^+}W_{g+3}) - \bar{W}W_{g+3}. \end{aligned}$$

The first three lines on the right are exact:

$$\begin{aligned} \tilde{\tau}L^{n_{g+3}} \cdot X - X &= \tilde{\partial}^{\mathcal{B}^+}(-\bar{W}X + W_{g+3}X - W_{g+3}X(\tilde{\partial}^{\mathcal{B}^+}W_{g+3})) \\ &\quad - \left(X + \sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j\right)(\tilde{\partial}^{\mathcal{B}^+}W_{g+3}) - \bar{W}W_{g+3}. \end{aligned}$$

By Proposition 2.5, it suffices to prove that the last line is exact. For this, it suffices to prove that

$$\tilde{\partial}^{\mathcal{B}^+} \left(\sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j \right) = \bar{W} - W,$$

since then the last line is just $\tilde{\partial}^{\mathcal{B}^+}[(\sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j + X)W_{g+3}]$.

Let us prove this final claim. Let η_0 through η_{g+3} be as in the proof of Proposition 4.11 so that η_k is between the times γ_k^Γ and γ_{k+1}^Γ . Then, as was proved in that proposition,

$$\tilde{\partial}^{\mathcal{B}^+}W_j = W(\eta_j) - W(\eta_{j-1}),$$

and similarly, we may conclude that

$$\tilde{\partial}^{\mathcal{B}^+}\bar{W}_i = W((\eta_{i-1})^{-1}) - W((\eta_i)^{-1}).$$

Therefore,

$$\begin{aligned} &\tilde{\partial}^{\mathcal{B}^+} \left(\sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j \right) \\ &= \sum_{1 \leq j < i \leq g+3} [(W((\eta_{i-1})^{-1}) - W((\eta_i)^{-1}))W_j - \bar{W}_i(W(\eta_j) - W(\eta_{j-1}))] \\ &= \sum_{j=1}^{g+3} (W((\eta_j)^{-1}) - W(\eta_{g+3}^{-1}))W_j - \sum_{i=1}^{g+3} \bar{W}_i(W(\eta_{i-1}) - W(\eta_0)) \end{aligned}$$

by summing over one index. But η_0 and η_{g+3} are contractible via homotopies not interacting with the centering, and so, $W(\eta_0) = W((\eta_{g+3})^{-1}) = I$. We find, therefore, that

$$\tilde{\partial}^{\mathcal{B}^+} \left(\sum_{1 \leq j < i \leq g+3} \bar{W}_i W_j \right) = \bar{W} - W + \sum_{j=1}^{g+3} (W(\eta_j^{-1} * \gamma_j) - W(\gamma_j^{-1} * \eta_{j-1})).$$

But $\eta_j^{-1} * \gamma_j \sim \gamma_j^{-1} * \eta_{j-1}$ are both homotopic to the out-and-back paths from the base point to the center c_j , so the last sum cancels, and we have the desired claim. ■

5.5. Proof of Theorem 5.1

In order to prove Theorem 5.1, we essentially just need to check that orientation changes commute with moves V–VII. We include an explicit proof for completeness.

Proof of Theorem 5.1. Suppose that $\zeta = (\iota, \tau) \in (\mathbb{Z}_2)^E \times (F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3))$, and Γ is a homotopy class of garden. Then, we may take

$$\Phi_{\Gamma}^{\zeta \cdot \Gamma}(\tau): (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\tau \cdot \Gamma}^{\mathcal{B}^+})$$

as the composition

$$\Phi_{\Gamma}^{\zeta \cdot \Gamma}(\tau) := \Phi_{\iota} \circ \Phi_{\Gamma}^{\tau},$$

where Φ_{ι} is as in Lemma 5.2 describing the action of changing orientations, and where Φ_{Γ}^{τ} is as in Lemma 5.12 describing the action of moves V, VI, and VII together. We now verify the desired properties.

- *Composition Property:* it suffices to check that orientation changes commute with moves V–VII, in which case we have that

$$\Phi_{\Gamma}^{\zeta \cdot \Gamma}(\zeta) = \Phi_{\iota} \circ \Phi_{\Gamma}^{\tau} = \Phi_{\iota \cdot \Gamma}^{\tau} \circ \Phi_{\iota}.$$

Once we establish this claim, the proof of the composition property is now the same as in the proof of Lemma 5.12, where we use the fact that move V and moves VI and VII commuted in the same way. We work towards establishing the claim.

First, note that if γ is a non-degenerate path, then

$$\Phi_{\iota}(W^{\Gamma}(\gamma)) = W^{\iota \cdot \Gamma}(\gamma).$$

Indeed, we may check this discontinuity by discontinuity for each binary sequence. Equality at centers through which γ passes is clear because

$$\Phi^{\iota} \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix}.$$

Similarly, for discontinuities at edges of the graph, since our sign conventions are precisely so that A_e and B_e come with opposite signs, the result is clear. Finally, for discontinuities at threads, we have already discussed the relevant statement in the proof of Lemma 5.2.

With this in mind, it is clear that orientation changes commute with move V. This is clear on the ring and on the degree 1 generators. On the degree 2 generators, this follows from the formula of move V in Proposition 5.3 combined with the fact just discussed, which implies that

$$\Phi_t(W^\Gamma(\gamma_{f,g})) = W^{\iota\Gamma}(\gamma_{f,g}),$$

where $\gamma_{f,g} = \gamma_{f,g}^\Gamma = \gamma_{f,g}^{\iota\Gamma}$ since the curve is independent of the orientations of the edges of the graph.

Move VI, with formula as in Lemma 5.8, is similar, since $\tilde{H}(\tau)$ appearing in the formula transforms in the same way. As for move VII, the matrix $\begin{pmatrix} 0 & -\lambda^{-1} \\ \lambda & 0 \end{pmatrix}$ is completely independent of the orientation of the edge, whereas applying Φ_t either preserves or negates this matrix, depending upon whether the corresponding edge has its orientation reversed. Either way, since we conjugate by this matrix, even if the orientation is reversed, the minus signs cancel, and we have that changing orientation commutes with move VII.

- *Orientation changes*: this is just a part of Lemma 5.2.
- *Modifying the garden*: this follows from Lemmas 5.12 and 5.13 (with the understanding that the stabilizer of Γ is generated by the elements occurring in Lemma 5.13, as discussed in Section 5.4). ■

6. Removing the enlargements

We now wish to prove Theorem 1.6, our starting point being Theorem 5.1 of the last section. We must form the following two operations:

- We must shrink the ring from $\tilde{R}^{nc} = \mathbb{Z}[\Pi]$ to $R_T^{nc} = \mathbb{Z}[\Pi_T]$, using functoriality properties from $\tilde{\mathcal{B}}_G^+$ (with its differentials depending upon a garden) to yield functoriality properties on $\mathcal{B}_{G,T}^+$ (with its differentials depending upon a garden as well as the tree T).
- In the special case that the garden Γ is of finite type (Definition 6.2 below), we must further show that we can destabilize to cancel the generators f_Γ and w .

We spend a subsection on each of these and a final short subsection to check that we have proved Theorems 1.4 and 1.6.

6.1. Shrinking the ring

The projections $\pi_T: \Pi \rightarrow \Pi_T$ from Section 3.2 induce graded algebra morphisms

$$\pi_T: \tilde{\mathcal{B}}_G^+ \rightarrow \mathcal{B}_{G,T}^+,$$

simply by reducing all of the coefficients. On the other hand, we have right inverses $i_T: \Pi_T \rightarrow \Pi$ constructed in Corollary 3.15, which induces graded algebra morphisms

$$i_T: \mathcal{B}_{G,T}^+ \rightarrow \tilde{\mathcal{B}}_G^+.$$

We may now simply check that

$$\pi_T \circ \tilde{\partial}_{G,\Gamma}^+ \circ i_T \circ \pi_T = \pi_T \circ \tilde{\partial}_{G,\Gamma}^+$$

by checking on generators (as an algebra) and extending to the entire algebra by using the Leibniz rule of the differential together with $\mathbb{Z}$ -linearity. By Lemma 3.2, the differentials $\tilde{\partial}_{G,\Gamma}^+$ hence induce differentials

$$\partial_{G,\Gamma,T}^+ = \pi_T \circ \tilde{\partial}_{G,\Gamma}^+ \circ i_T: \mathcal{B}_{G,T}^+ \rightarrow \mathcal{B}_{G,T}^+.$$

In practice, the new differential $\partial_{G,\Gamma,T}^+$ acts in the same way as it usually would, but every time we see a generator $\lambda_{e,o}$ for an edge e not in T arising from the differential of an element of $F \cup \{x, y, z, w\}$, we simply set it equal to 1.

Recall that we defined morphisms

$$\tilde{\mathcal{C}}_T^{T'}(\gamma): \Pi \rightarrow \Pi$$

given on generators $\lambda_{e,o}$, which we may think of as a transverse signed path from c_f^+ to c_g^- in Λ_G , by

$$(\tilde{\mathcal{C}}_T^{T'}(\gamma))(\lambda_{e,o}) = \beta_T^{T'}(f, \gamma, +) \cdot \lambda_{e,o} \cdot \beta_T^{T'}(f, \gamma, -)^{-1}.$$

In turn, this induces a graded algebra isomorphism

$$\tilde{\mathcal{C}}_T^{T'}(\gamma): \tilde{\mathcal{B}}_G^+ \rightarrow \tilde{\mathcal{B}}_G^+,$$

with the caveat that it does not intertwine the differential $\tilde{\partial}_{G,\Gamma}^+$. In order to fix this, we may now consider the regeneration

$$\tilde{\mathcal{R}}_T^{T'}(\gamma): \tilde{\mathcal{B}}_G^+ \rightarrow \tilde{\mathcal{B}}_G^+$$

defined on generators by

$$\begin{aligned} (\tilde{\mathcal{R}}_T^{T'}(\gamma))(f) &= \beta_T^{T'}(f, \gamma, -) \cdot f \cdot \beta_T^{T'}(f, \gamma, +)^{-1}, \\ (\tilde{\mathcal{R}}_T^{T'}(\gamma))(x) &= \beta_T^{T'}(f_\Gamma, \gamma, +) \cdot x \cdot \beta_T^{T'}(f_\Gamma, \gamma, +)^{-1}, \end{aligned}$$

$$\begin{aligned}
(\tilde{\mathcal{R}}_T^{T'}(\gamma))(y) &= \beta_T^{T'}(f_\Gamma, \gamma, +) \cdot y \cdot \beta_T^{T'}(f_\Gamma, \gamma, -)^{-1}, \\
(\tilde{\mathcal{R}}_T^{T'}(\gamma))(z) &= \beta_T^{T'}(f_\Gamma, \gamma, -) \cdot z \cdot \beta_T^{T'}(f_\Gamma, \gamma, +)^{-1}, \\
(\tilde{\mathcal{R}}_T^{T'}(\gamma))(w) &= \beta_T^{T'}(f_\Gamma, \gamma, -) \cdot w \cdot \beta_T^{T'}(f_\Gamma, \gamma, +)^{-1}.
\end{aligned}$$

In other words, we simply treat f as though it were a path from c_f^- to c_f^+ , and similarly, the generators in X as paths from $c_{f_\Gamma}^\pm$ to $c_{f_\Gamma}^\pm$, depending upon which particular generator we are looking at. (For example, the generator w , corresponding to binary sequences of type $(1, 0)$, are thought of as paths from $c_{f_\Gamma}^+$ to $c_{f_\Gamma}^-$.)

Proposition 6.1. *For a fixed garden Γ , the graded algebra isomorphism*

$$\tilde{\Phi}_T^{T'}(\gamma) := \tilde{\mathcal{R}}_T^{T'}(\gamma) \circ \tilde{\mathcal{C}}_T^{T'}(\gamma): \tilde{\mathcal{B}}_G^+ \rightarrow \tilde{\mathcal{B}}_G^+$$

is a dg-automorphism with respect to the differential $\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}$. Furthermore, these satisfy the composition rule

$$\tilde{\Phi}_T^{T''}(\gamma') \circ \tilde{\Phi}_T^{T'}(\gamma) = \tilde{\Phi}_T^{T''}(\gamma * \gamma').$$

Proof. We first check that $\tilde{\Phi}_T^{T'}(\gamma)$ intertwines the differential by checking this property on generators. For the degree 1 generators f , we have

$$\begin{aligned}
(\tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} \circ \tilde{\Phi}_T^{T'}(\gamma))(f) &= \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}(\beta_T^{T'}(f, \gamma, -) \cdot f \cdot \beta_T^{T'}(f, \gamma, +)^{-1}) \\
&= \beta_T^{T'}(f, \gamma, -) \cdot \sum_{v \in f} H(\tau_f(v)) \cdot \beta_T^{T'}(f, \gamma, +)^{-1},
\end{aligned}$$

whereas

$$(\tilde{\Phi}_T^{T'}(\gamma) \circ \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+})(f) = \tilde{\mathcal{C}}_T^{T'}(\gamma) \left(\sum_{v \in f} H(\tau_f(v)) \right).$$

If we write out $H(\tau_f(v))$ as a (signed) product of three generators in W , we see that

$$\tilde{\mathcal{C}}_T^{T'}(\gamma)(H(\tau_f(v))) = \beta_T^{T'}(f, \gamma, -) \cdot H(\tau_f(v)) \cdot \beta_T^{T'}(f, \gamma, +)^{-1}$$

since the internal β terms cancel out.

Similarly, for the degree 2 generators, we have the same principle, but where we notice that our binary sequences are based at curves starting and ending at the seed, located in the face f_Γ . For example, for the generator x , we have

$$\partial_{G,\Gamma}^{\mathcal{B}^+}(x) = \sum_f w_{1,0}(\gamma_f^s) \cdot f \cdot w_{1,1}(\gamma_f^t),$$

and we notice correspondingly that each word comprising $w_{1,0}(\gamma_f^s)$ corresponds to a geometric element of the group Π which originates with label $+$ in the face f_Γ , and similarly, $w_{1,1}(\gamma_f^t)$ terminates with label $+$ in f_Γ .

Finally, for the composition property, we may check it on the ring and on the generators. On the ring, we just have $\tilde{\mathcal{C}}_T^{T'}(\gamma)$ which satisfy the desired composition property. On face generators, we have

$$\begin{aligned}
 & (\tilde{\Phi}_T^{T''}(\gamma') \circ \tilde{\Phi}_T^{T'}(\gamma))(f) \\
 &= (\tilde{\Phi}_T^{T''}(\gamma'))(\beta_T^{T'}(f, \gamma, -) \cdot f \cdot \beta_T^{T'}(f, \gamma, +)^{-1}) \\
 &= (\tilde{\mathcal{C}}_T^{T''}(\gamma'))(\beta_T^{T'}(f, \gamma, -)) \cdot (\tilde{\mathcal{R}}_T^{T''}(\gamma))(f) \cdot (\tilde{\mathcal{C}}_T^{T''}(\gamma'))(\beta_T^{T'}(f, \gamma, +)^{-1}) \\
 &= (\beta_T^{T''}(f, \gamma', -) \cdot \beta_T^{T'}(f, \gamma, -) \cdot \beta_T^{T''}(f, \gamma', -)^{-1}) \\
 &\quad \cdot (\beta_T^{T''}(f, \gamma', -) \cdot f \cdot \beta_T^{T''}(f, \gamma', +)^{-1}) \\
 &\quad \cdot (\beta_T^{T''}(f, \gamma', +) \cdot \beta_T^{T'}(f, \gamma, +) \cdot \beta_T^{T''}(f, \gamma', +)^{-1}) \\
 &= \beta_T^{T''}(f, \gamma * \gamma', -) \cdot f \cdot \beta_T^{T''}(f, \gamma * \gamma', +)^{-1} \\
 &= (\tilde{\Phi}_T^{T''}(\gamma * \gamma'))(f).
 \end{aligned}$$

A similar computation works for the degree 2 generators. ■

Upon checking that

$$\pi_{T'} \circ \tilde{\Phi}_T^{T'}(\gamma) \circ i_T \circ \pi_T = \pi_{T'} \circ \tilde{\Phi}_T^{T'}(\gamma)$$

on generators and hence on the entire graded algebra, we may apply Lemma 3.2 so that, for fixed garden Γ , the dg-isomorphisms $\tilde{\Phi}_T^{T'}(\gamma)$ induce dg-isomorphisms

$$\Phi_T^{T'}(\gamma): (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}) \rightarrow (\mathcal{B}_{G,T'}^+, \partial_{G,\Gamma,T'}^{\mathcal{B}^+}).$$

The fact that they intertwine the differentials also follows from Lemma 3.2, since the morphisms $\Phi_T^{T'}(\gamma) \circ \partial_{G,\Gamma,T}^{\mathcal{B}^+}$ and $\partial_{G,\Gamma,T'}^{\mathcal{B}^+} \circ \Phi_T^{T'}(\gamma)$ are induced by the same morphism upstairs given by

$$\tilde{\Phi}_T^{T'}(\gamma) \circ \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} = \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+} \circ \tilde{\Phi}_T^{T'}(\gamma)$$

and hence are equal.

6.2. Destabilization for finite-type gardens

Recall from Remark 4.2 that the Casals–Murphy definition of a garden excludes a tine passing through the face at infinity. We use the following slightly different definition.

Definition 6.2. A garden Γ on a trivalent plane graph G is said to be of *finite type* if the tine γ_{f_∞} passing through the center of the face at infinity c_{f_Γ} comes either first or last as it exits the seed.

Notice that if a garden is of finite type, then the tine γ_{f_∞} may be homotoped so that it is completely contained inside of the face f_∞ . After this homotopy, the tine only interacts with the rest of the garden at threads and the center c_{f_∞} . It follows that there is a unique binary sequence in $\mathcal{BS}(\gamma_{f_\infty})$, since any binary sequence must go from 0 to 1 as the tine passes through f_∞ , and can only go from 0 to 1 as it passes through threads. In other words, for a finite-type garden Γ , the differential on $\tilde{\mathcal{B}}_G^+$ on degree 2 generators is of the form

$$\tilde{\partial}_{G,\Gamma}^+ X = \begin{pmatrix} 0 & f_\Gamma \\ 0 & 0 \end{pmatrix} + \sum_{f \in F_{\text{fin}}} W(\gamma_f) = \begin{pmatrix} 0 & f_\Gamma \\ 0 & 0 \end{pmatrix} + \sum_{f \in F_{\text{fin}}} W(\gamma_f^s) \cdot \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} \cdot W(\gamma_f^T).$$

In other words, the differentials of the generators x , y , and z only see terms corresponding to the tines through the finite faces, whereas the differential of w picks up f_Γ for the face at infinity. This already explains algebraically why, when Casals and Murphy defined a garden, they were able to simply not include a tine through the face at infinity—they were only looking at the generators x , y , and z !

Projecting under π_T , we have that for a finite-type garden,

$$\partial_{G,\Gamma,T}^+ X = \begin{pmatrix} 0 & f_\Gamma \\ 0 & 0 \end{pmatrix} + \sum_{f \in F_{\text{fin}}} \pi_T(W(\gamma_f^s)) \cdot \begin{pmatrix} 0 & f \\ 0 & 0 \end{pmatrix} \cdot (W(\gamma_f^T)).$$

Singling out the generator w ,

$$\partial_{G,\Gamma,T}^+ w = f_\Gamma + \sum_{f \in F_{\text{fin}}} \pi_T(w_{0,0}(\gamma_f^s)) \cdot f \cdot \pi_T(w_{1,1}(\gamma_f^T)).$$

Consider now the elementary automorphism $\mu_{G,\Gamma,T}$ of $\mathcal{B}_{G,T}^+$ given by

$$\mu_{G,\Gamma,T}(f_\Gamma) = f_\Gamma - \sum_{f \in F_{\text{fin}}} \pi_T(w_{0,0}(\gamma_f^s)) \cdot f \cdot \pi_T(w_{1,1}(\gamma_f^T))$$

(and the identity on all other generators), inducing a dg-isomorphism

$$\mu_{G,\Gamma,T}: (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+) \xrightarrow{\sim} (\mathcal{B}_{G,T}^+, (\partial_{G,\Gamma,T}^+)',)$$

with induced differential the same on all generators except w and f_Γ , for which

$$\begin{aligned} (\partial_{G,\Gamma,T}^+)' w &= f_\Gamma, \\ (\partial_{G,\Gamma,T}^+)' f_\Gamma &= 0. \end{aligned}$$

We may therefore destabilize, yielding the desired dg-algebra $(\mathcal{B}_{G,f_\Gamma,T}, \partial_{G,\Gamma,T}^{\mathcal{B}})$. Explicitly, the formula for $\partial_{G,\Gamma,T}^{\mathcal{B}}$ on the generators in F_{fin} and x , y , and z is just the same as the formula for $\partial_{G,\Gamma,T}^+$, where we notice that none of these differentials include a term f_Γ or w by inspection.

6.3. Proof of Theorems 1.4 and 1.6

We now have all of the ingredients to prove our desired theorems.

Proof of Theorem 1.6. We have constructed the dg-algebras $(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+})$ in Section 6.1. For a fixed garden, we have discussed changing trees in the form of the morphisms $\tilde{\Phi}_T^{T'}(\gamma)$ in Proposition 6.1 and its induced morphisms $\Phi_T^{T'}(\gamma)$ downstairs. On the other hand, by Lemma 3.2 (along with the usual check that the lemma can be applied), the morphisms $\Phi_\Gamma^{\Gamma'}(\zeta)$ constructed in Section 5.1 via the actions of moves V–VII, where

$$\Gamma' = \zeta \cdot \Gamma,$$

descend from $\tilde{\mathcal{B}}_G^+$ under π_T ; i.e., we obtain

$$\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta): (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}) \rightarrow (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma',T}^{\mathcal{B}^+})$$

arising as the unique graded algebra morphism fitting into the diagram

$$\begin{array}{ccc} \tilde{\mathcal{B}}_G^+ & \xrightarrow{\Phi_\Gamma^{\Gamma'}(\zeta)} & \tilde{\mathcal{B}}_G^+ \\ \pi_T \downarrow & & \downarrow \pi_T \\ \mathcal{B}_{G,T}^+ & \xrightarrow[\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta)]{-} & \mathcal{B}_{G,T}^+ \end{array}$$

By Lemma 3.2, the fact that it intertwines the differentials follows from the same property upstairs. We may therefore define

$$\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma) := \Phi_T^{T'}(\gamma) \circ \Phi_{\Gamma,T}^{\Gamma',T'}(\zeta).$$

We now check the desired properties.

- *Composition property:* just like in the proof of Theorem 5.1, it suffices to check that the two factors $\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta)$ and $\Phi_T^{T'}(\gamma)$ commute with each other, in the sense that

$$\Phi_T^{T'}(\gamma) \circ \Phi_{\Gamma,T}^{\Gamma',T'}(\zeta) = \Phi_{\Gamma,T'}^{\Gamma',T'}(\zeta) \circ \Phi_T^{T'}(\gamma),$$

since if this is satisfied, then the composition property of Theorem 5.1 and the composition property of Proposition 6.1 yield the desired result. It suffices to check this composition property upstairs, i.e., that

$$\tilde{\Phi}_T^{T'}(\gamma) \circ \Phi_\Gamma^{\Gamma'}(\zeta) = \Phi_\Gamma^{\Gamma'}(\zeta) \circ \tilde{\Phi}_T^{T'}(\gamma).$$

Because $\Phi_\Gamma^{\Gamma'}(\zeta)$ is itself constructed as a composition of orientation changes and moves V–VII, it suffices to check this commutativity for each of these types of

moves, each of which is easy to check. For example, if ζ acts via an application of move V, then since it preserves all degree 1 generators, it suffices to check commutativity on the degree 2 generators. That is, if the application of move V is with respect to the faces f and g , we have, for example,

$$\begin{aligned}
 (\tilde{\Phi}_T^{T'}(\gamma) \circ \Phi_{\Gamma}^{\Gamma'}(\zeta))(x) &= \tilde{\Phi}_T^{T'}(\gamma)(x + w_{1,1}(\gamma_{f,g}^{\Gamma})) \\
 &= (\tilde{\Phi}_T^{T'}(\gamma))(x) + (\tilde{\Phi}_T^{T'}(\gamma))w_{1,1}(\gamma_{f,g}^{\Gamma}) \\
 &= \beta_T^{T'}(f_{\Gamma}, \gamma, +) \cdot x \cdot \beta_T^{T'}(f_{\Gamma}, \gamma, +)^{-1} \\
 &\quad + \beta_T^{T'}(f_{\Gamma}, \gamma, +) \cdot w_{1,1}(\gamma_{f,g}^{\Gamma}) \cdot \beta_T^{T'}(f_{\Gamma}, \gamma, +)^{-1} \\
 &= \beta_T^{T'}(f_{\Gamma}, \gamma, +) \cdot (x + w_{1,1}(\gamma_{f,g}^{\Gamma})) \cdot \beta_T^{T'}(f_{\Gamma}, \gamma, +)^{-1} \\
 &= \beta_T^{T'}(f_{\Gamma}, \gamma, +) \cdot ((\Phi_{\Gamma}^{\Gamma'})(\zeta))(x) \cdot \beta_T^{T'}(f_{\Gamma}, \gamma, +)^{-1} \\
 &= (\Phi_{\Gamma}^{\Gamma'}(\zeta))(\beta_T^{T'}(f_{\Gamma}, \gamma, +) \cdot x \cdot \beta_T^{T'}(f_{\Gamma}, \gamma, +)^{-1}) \\
 &= (\Phi_{\Gamma}^{\Gamma'}(\zeta) \circ \tilde{\Phi}_T^{T'}(\gamma))(x).
 \end{aligned}$$

The other generators y, z, w are analogous, and we leave moves VI and VII for the reader.

- *Orientation changes:* this descends from the same property in Theorem 5.1 under Lemma 3.2.
- *Fixing the trees and capping paths; modifying the garden:* this also descends from the enlarged property.
- *Fixing the garden; modifying the trees and capping paths:* this follows from Proposition 6.1.
- *Finite-type gardens:* we have defined the differential on $\partial_{G,\Gamma,T}^{\mathcal{B}}$ precisely in this way in Section 6.2. ■

Proof of Theorem 1.4. We have already found the desired differentials yielding the dg-algebras $(\mathcal{B}_{G,f_{\Gamma},T}, \partial_{G,\Gamma,T}^{\mathcal{B}})$. The fact that these partially abelianize to the Casals–Murphy dg-algebras is clear by comparing our construction with their construction.

Suppose that Γ and Γ' are two finite-type gardens and

$$T, T' \subset G$$

are two trees spanning all but one vertex. For any homotopy class of path γ from $*_T$ to $*_{T'}$ and element $\zeta \in \mathcal{H}(G)$ such that

$$\zeta \cdot \Gamma = \Gamma',$$

we then have a quasi-isomorphism induced under destabilization as in Lemma 3.2 as in the following diagram:

$$\begin{array}{ccc}
 (\mathcal{B}_{G,f_{\Gamma},T}, \partial_{\mathcal{B}_{G,\Gamma},T}^{\mathcal{B}}) & \xleftarrow{\text{destab}} & (\mathcal{B}_{G,T}^+, (\partial_{\mathcal{B}_{G,\Gamma},T}^{\mathcal{B}})') \\
 \downarrow \text{dashed} & & \downarrow \mu_{G,\Gamma,T}^{-1} \\
 & & (\mathcal{B}_{G,T}^+, \partial_{\mathcal{B}_{G,\Gamma},T}^{\mathcal{B}}) \\
 & & \downarrow \Phi_{\Gamma,T'}^{\Gamma',T'}(\gamma) \\
 & & (\mathcal{B}_{G,T'}^+, \partial_{\mathcal{B}_{G,\Gamma',T'}}^{\mathcal{B}}) \\
 & & \downarrow \mu_{G,\Gamma',T'} \\
 (\mathcal{B}_{G,f_{\Gamma'},T'}, \partial_{\mathcal{B}_{G,\Gamma'},T'}^{\mathcal{B}}) & \xleftarrow{\text{destab}} & (\mathcal{B}_{G,T'}^+, (\partial_{\mathcal{B}_{G,\Gamma',T'}}^{\mathcal{B}})')
 \end{array}$$

This is clearly a dg-isomorphism, since its inverse is constructed in the same way but swapping (Γ, T) with (Γ', T') and using γ^{-1} in place of γ . ■

7. Representations and rank r face colorings

The goal of this section is to prove Theorem 1.12.

7.1. Representations-to-colors for the enlarged dg-algebra

As with functoriality between the various dg-algebras $(\mathcal{B}_{G,T}^+, \partial_{\mathcal{B}_{G,\Gamma},T}^{\mathcal{B}})$, when studying representations, it is convenient to begin with $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{\mathcal{B}_{G,\Gamma}}^{\mathcal{B}})$, and to therefore verify a version of Theorem 1.12 but with respect to the dg-isomorphism $\Phi_{\Gamma}^{\Gamma'}(\zeta)$ appearing in Theorem 5.1. Let us therefore begin with a study of $\text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{\mathcal{B}_{G,\Gamma}}^{\mathcal{B}}; \mathbb{F})$ and see how we can extract from it a rank r face coloring.

For degree reasons, a representation $\varepsilon \in \text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{\mathcal{B}_{G,\Gamma}}^{\mathcal{B}}; \mathbb{F})$ kills all of the generators in nonzero degree, and hence, we may identify ε with a representation of Π , which we think of as a ring homomorphism

$$\varepsilon_{\Pi}: \mathbb{Z}[\Pi] \rightarrow \text{Mat}_{r \times r}(\mathbb{F}).$$

The fact that ε is a representation of the dg-algebra is just the requirement that

$$\varepsilon_{\Pi}(\tilde{\partial}_{\mathcal{B}_{G,\Gamma}}^{\mathcal{B}} f) = 0$$

for all $f \in F$. Writing this out, we are simply looking at representations of Π with the property that

$$\sum_{v \in f} \varepsilon_{\Pi}(H(\tau_f(v))) = 0.$$

Recall that associated to any path γ on the sphere S^2 , non-degenerate with respect to the graph G and its web, there is an associated word

$$W(\gamma) \in \text{Mat}_{2 \times 2}(\mathbb{Z}[\Pi] * \mathbb{Z}\langle F \rangle).$$

If the path γ furthermore is disjoint from the face centers, then all of the faces disappear from the expression, and $W(\gamma) \in \text{Mat}_{2 \times 2}(\mathbb{Z}[\Pi])$, in which case $W(\gamma)$ is a composition of matrices either of one of the following two forms:

$$\begin{pmatrix} 1 & H \\ 0 & 1 \end{pmatrix} \quad \pm \begin{pmatrix} 0 & \lambda \\ -\bar{\lambda}^{-1} & 0 \end{pmatrix}.$$

Here, H is a monomial in Π , with sign ± 1 , and λ and $\bar{\lambda}$ are just elements of Π . Hence, under ε_{Π} , the matrices

$$\begin{aligned} \varepsilon_{\Pi} \begin{pmatrix} 1 & H \\ 0 & 1 \end{pmatrix} &= \begin{pmatrix} I & \varepsilon_{\Pi}(H) \\ 0 & I \end{pmatrix}, \\ \varepsilon_{\Pi} \begin{pmatrix} 0 & \lambda \\ -\bar{\lambda}^{-1} & 0 \end{pmatrix} &= \begin{pmatrix} 0 & \varepsilon_{\Pi}(\lambda) \\ -\varepsilon_{\Pi}(\bar{\lambda})^{-1} & 0 \end{pmatrix} \end{aligned}$$

are both invertible matrices; i.e., they live in $\text{GL}_{2r}(\mathbb{F})$.

Suppose that γ is a small loop around c_f . Then,

$$W(\gamma) = \begin{pmatrix} 1 & \tilde{\partial}_{G,\Gamma}^+ f \\ 0 & 1 \end{pmatrix}.$$

The condition that ε is a representation implies that

$$\varepsilon_{\Pi}(W(\gamma)) = I_{2r} \in \text{GL}_{2r}(\mathbb{F}).$$

Hence, even though $W(\gamma)$ is invariant only up to homotopies (with fixed endpoints) not passing through the centering, the matrix $\varepsilon_{\Pi}(W(\gamma))$ is now independent under all homotopies (with fixed endpoints). Since our paths lie on S^2 , any two paths with the same endpoints are homotopic, so the corresponding element $\varepsilon_{\Pi}(W(\gamma)) \in \text{GL}_{2r}(\mathbb{F})$ only depends upon the endpoints of γ . Furthermore, two points in the same triangle $\text{Tri}(G)$ are essentially indistinguishable. We therefore obtain, for any two triangles $v_1, v_2 \in \text{Tri}(G)$, an element

$$M^{\varepsilon}(v_1, v_2) \in \text{GL}_{2r}(\mathbb{F}).$$

These satisfy the following properties:

- For all $v_1, v_2, v_3 \in \text{Tri}(G)$, we have

$$M^\varepsilon(v_1, v_2) \cdot M^\varepsilon(v_2, v_3) = M^\varepsilon(v_1, v_3).$$

- For all $v \in \text{Tri}(G)$, we have $M^\varepsilon(v, v) = I_{2r}$.

Recall now that as part of the data of a garden Γ , we have a base point, which lives in one of the triangles $v_\Gamma \in \text{Tri}(G)$. Hence, for each $v \in \text{Tri}(G)$, we may define

$$M^{\varepsilon, \Gamma}(v) := M^\varepsilon(v_\Gamma, v) \in \text{GL}_{2r}(\mathbb{F}).$$

Consider the $2r \times r$ matrix $\begin{pmatrix} I \\ 0 \end{pmatrix}$. The product $M^{\varepsilon, \Gamma}(v) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix}$ is the $2r \times r$ matrix consisting of the first r columns of $M^{\varepsilon, \Gamma}(v)$. Suppose that v_1 and v_2 are two triangles sharing an edge along a thread of the graph. Then, we have that, for some $H \in \text{GL}_r(\mathbb{F})$,

$$\begin{aligned} M^{\varepsilon, \Gamma}(v_2) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix} &= M^{\varepsilon, \Gamma}(v_1) \cdot M^\varepsilon(v_1, v_2) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix} \\ &= M^{\varepsilon, \Gamma}(v_1) \cdot \begin{pmatrix} I & H \\ 0 & I \end{pmatrix} \cdot \begin{pmatrix} I \\ 0 \end{pmatrix} = M^{\varepsilon, \Gamma}(v_1) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix}. \end{aligned}$$

Therefore, for each $f \in F$, we obtain a canonical $2r \times r$ matrix $M^{\varepsilon, \Gamma}(f)$ given as

$$M^{\varepsilon, \Gamma}(f) := M^{\varepsilon, \Gamma}(v) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix}$$

for any triangle $v \in \text{Tri}(G)$ contained in f . In fact, this matrix is of rank r , on account of the fact that the columns of $M^{\varepsilon, \Gamma}(v)$ are linearly independent, a fortiori the first r columns are linearly independent. Taking the span of the columns of $M^{\varepsilon, \Gamma}(f)$ yields an element $\chi^{\varepsilon, \Gamma}(f) \in \text{Gr}(r, 2r; \mathbb{F})$.

Similarly, suppose that $v_1, v_2 \in \text{Tri}(G)$ are two triangles sharing an edge which is also an edge of the graph. Let us write $f_1, f_2 \in F$ for the two faces containing v_1 and v_2 , respectively. Then, we have that there are some $\Lambda, \Lambda' \in \text{GL}_r(\mathbb{F})$ such that

$$\begin{aligned} M^{\varepsilon, \Gamma}(f_2) &= M^{\varepsilon, \Gamma}(v_2) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix} = M^{\varepsilon, \Gamma}(v_1) \cdot M^\varepsilon(v_1, v_2) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix} \\ &= M^{\varepsilon, \Gamma}(v_1) \cdot \begin{pmatrix} 0 & \Lambda' \\ \Lambda & 0 \end{pmatrix} \cdot \begin{pmatrix} I \\ 0 \end{pmatrix} = M^{\varepsilon, \Gamma}(v_1) \cdot \begin{pmatrix} 0 \\ \Lambda \end{pmatrix} \\ &= M^{\varepsilon, \Gamma}(v_1) \cdot \begin{pmatrix} 0 \\ I \end{pmatrix} \cdot \Lambda. \end{aligned}$$

Notice that $M^{\varepsilon, \Gamma}(v_1) \cdot \begin{pmatrix} 0 \\ I \end{pmatrix}$ consists of the last r column vectors of $M^{\varepsilon, \Gamma}(v_1)$, which are linearly independent with span transverse to the span of the first r columns, i.e., $\chi^{\varepsilon, \Gamma}(f_1)$. Multiplying by Λ on the right does not change the span. Therefore, we see that $\chi^{\varepsilon, \Gamma}(f_1)$ and $\chi^{\varepsilon, \Gamma}(f_2)$ are transverse.

Let us be explicit about what our construction has given. In the first step of the construction, we had a morphism

$$M^\Gamma: \text{Rep}_r((\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G, \Gamma}^+); \mathbb{F}) \rightarrow (\text{GL}_{2r}(\mathbb{F}))^{\text{Tri}(G)}.$$

For each element of $\text{GL}_{2r}(\mathbb{F})$, we then multiply on the right by $\begin{pmatrix} I \\ 0 \end{pmatrix}$ to obtain, for each $v \in \text{Tri}(G)$, an element of $\text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F})$, by which we mean a $2r \times r$ matrix of rank r , yielding the composition

$$\text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G, \Gamma}^+; \mathbb{F}) \xrightarrow{M^\Gamma} (\text{GL}_{2r}(\mathbb{F}))^{\text{Tri}(G)} \xrightarrow{\bullet \begin{pmatrix} I \\ 0 \end{pmatrix}} (\text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F}))^{\text{Tri}(G)}.$$

The image of this composition has the property that it lands inside of $(\text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F}))^F \subset (\text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F}))^{\text{Tri}(G)}$, those elements for which all triangles in the same face have the same matrix. Taking the span of the column vectors yields a map

$$\text{span}: \text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F}) \rightarrow \text{Gr}(r, 2r; \mathbb{F}),$$

which we may now apply to each face to yield a morphism

$$\text{Span}: (\text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F}))^F \rightarrow (\text{Gr}(r, 2r; \mathbb{F}))^F.$$

If we trace this through, starting at a representation, we end up with an element of $\text{Col}_r(G; \mathbb{F})$.

Summarizing, we have therefore constructed a map

$$\tilde{\Psi}_{G, \Gamma}: \text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G, \Gamma}^+) \rightarrow \text{Col}_r(G; \mathbb{F})$$

given by

$$\tilde{\Psi}_{G, \Gamma}(\varepsilon) := (\chi^{\varepsilon, \Gamma}(f))_{f \in F}$$

which fits into the following diagram, where the dashed arrows indicate that these exist (and are obviously uniquely determined):

$$\begin{array}{ccccc} (\text{GL}_{2r}(\mathbb{F}))^{\text{Tri}(G)} & \xrightarrow{\bullet \begin{pmatrix} I \\ 0 \end{pmatrix}} & (\text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F}))^{\text{Tri}(G)} & & \\ \uparrow M^\Gamma & & \uparrow & & \\ \text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G, \Gamma}^+; \mathbb{F}) & \dashrightarrow & (\text{Mat}_{2r \times r}^{\text{rk}=r}(\mathbb{F}))^F & \xrightarrow{\text{Span}} & (\text{Gr}(r, 2r; \mathbb{F}))^F \\ & \searrow \tilde{\Psi}_{G, \Gamma} & & & \uparrow \\ & & & & \text{Col}_r(G; \mathbb{F}) \end{array}$$

Let us trace the conjugation action of $\mathrm{GL}_r(\mathbb{F})$ on $\mathrm{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}; \mathbb{F})$ through this diagram. Suppose that $g \in \mathrm{GL}_r(\mathbb{F})$. Then, by a simple computation,

$$(g \cdot \varepsilon)_\Pi \begin{pmatrix} 1 & H \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} g & 0 \\ 0 & g \end{pmatrix} \begin{pmatrix} 1 & \varepsilon_\Pi(H) \\ 0 & 1 \end{pmatrix} \begin{pmatrix} g^{-1} & 0 \\ 0 & g^{-1} \end{pmatrix}$$

and

$$(g \cdot \varepsilon)_\Pi \begin{pmatrix} 0 & \lambda \\ -\bar{\lambda}^{-1} & 0 \end{pmatrix} = \begin{pmatrix} g & 0 \\ 0 & g \end{pmatrix} \begin{pmatrix} 0 & \varepsilon_\Pi(\lambda) \\ -\varepsilon_\Pi(\bar{\lambda})^{-1} & 0 \end{pmatrix} \begin{pmatrix} g^{-1} & 0 \\ 0 & g^{-1} \end{pmatrix}.$$

We find, therefore, that since our map M^Γ was given as a composition of these matrices, that

$$M^{g \cdot \varepsilon, \Gamma}(v) = \begin{pmatrix} g & 0 \\ 0 & g \end{pmatrix} M^{\varepsilon, \Gamma}(v) \begin{pmatrix} g^{-1} & 0 \\ 0 & g^{-1} \end{pmatrix}.$$

In other words, M^Γ intertwines the $\mathrm{GL}_r(\mathbb{F})$ action on representations with simultaneous conjugation by the matrix $\begin{pmatrix} g & 0 \\ 0 & g \end{pmatrix}$. Multiplying by the matrix $\begin{pmatrix} I & \\ & 0 \end{pmatrix}$ on the right, we see that the unlabeled dotted map in the previous diagram intertwines the $\mathrm{GL}_r(\mathbb{F})$ -conjugation on representations with the action of $\mathrm{GL}_r(\mathbb{F})$ on $(\mathrm{Mat}_{2r \times r}^{\mathrm{rk}=r}(\mathbb{F}))^F$ given by simultaneous action on each component by

$$g \cdot M = \begin{pmatrix} g & 0 \\ 0 & g \end{pmatrix} \cdot M \cdot g^{-1}$$

for $M \in \mathrm{Mat}_{2r \times r}^{\mathrm{rk}=r}(\mathbb{F})$. Finally, applying Span, we find that the multiplication by g^{-1} on the right in the above expression does not change the span, and we are hence left with the fact that $\tilde{\Psi}_{G,\Gamma}$ intertwines the conjugation action of $\mathrm{GL}_r(\mathbb{F})$ with the action of $\mathrm{PGL}_{2r}(\mathbb{F})$ on $\mathrm{Col}_r(G; \mathbb{F})$ via the map from $\mathrm{GL}_r(\mathbb{F})$ to $\mathrm{PGL}_{2r}(\mathbb{F})$ given by

$$g \mapsto \begin{bmatrix} g & 0 \\ 0 & g \end{bmatrix}.$$

In conclusion, because $\tilde{\Psi}_{G,\Gamma}$ has this compatibility between the conjugation action on representations by elements of $\mathrm{GL}_r(\mathbb{F})$ and the action on colorings by elements of $\mathrm{PGL}_{2r}(\mathbb{F})$, it descends to the corresponding moduli spaces as

$$\Psi_{G,\Gamma}: \mathcal{M}_r^{\mathrm{Rep}}(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^{\mathcal{B}^+}; \mathbb{F}) \rightarrow \mathcal{M}_r^{\mathrm{Col}}(G; \mathbb{F}).$$

7.2. Functoriality for the representations-to-colors map

Notice that the maps $\bullet \left(\begin{smallmatrix} I \\ 0 \end{smallmatrix} \right)$ and Span both intertwine a natural simultaneous left action of $\mathrm{GL}_{2r}(\mathbb{F})$ which reduces to the desired $\mathrm{PGL}_{2r}(\mathbb{F})$ -action on $\mathrm{Color}_r(G; \mathbb{F})$. We wish

to intertwine the functoriality properties for the dg-algebra $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+)$ appearing in Theorem 5.1 with the $\mathrm{PGL}_{2r}(\mathbb{F})$ -action on $\mathrm{Color}_r(G; \mathbb{F})$.

Suppose that $\zeta \in \mathcal{H}(G) = (\mathbb{Z}_2)^E \times F_{g+2} \times (\mathbb{Z}_2 * \mathbb{Z}_3)$, and let Γ and Γ' be two gardens with $\Gamma' = \zeta \cdot \Gamma$. Then, we have maps

$$\Phi_{\Gamma'}^{\Gamma'}(\zeta): (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+) \rightarrow (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma'}^+),$$

inducing pullbacks

$$(\Phi_{\Gamma'}^{\Gamma'}(\zeta))^*: \mathrm{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma'}^+; \mathbb{F}) \rightarrow \mathrm{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F}).$$

The following proposition yields the desired functoriality for the maps $\Psi_{G,\Gamma}$.

Proposition 7.1. *For each Γ, Γ', ζ as above, we have*

$$\Psi_{G,\Gamma'} = \Psi_{G,\Gamma} \circ (\Phi_{\Gamma'}^{\Gamma'}(\zeta))^*.$$

Proof. Let us be more explicit about what we wish to prove. Suppose that we are given

$$\varepsilon \in \mathrm{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma'}^+; \mathbb{F}).$$

Then, we have

$$\begin{cases} M^{\Gamma'}(\varepsilon) \in (\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)} \\ M^{\Gamma}((\Phi_{\Gamma'}^{\Gamma'}(\zeta))^* \varepsilon) \in (\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)} \end{cases}.$$

It suffices to find an element $Z_{\Gamma'}^{\Gamma'}(\zeta, \varepsilon) \in \mathrm{GL}_{2r}(\mathbb{F})$ such that

$$Z_{\Gamma'}^{\Gamma'}(\zeta, \varepsilon) \cdot M^{\Gamma}((\Phi_{\Gamma'}^{\Gamma'}(\zeta))^* \varepsilon) = M^{\Gamma'}(\varepsilon),$$

since then we see that

$$\tilde{\Psi}_{G,\Gamma'}(\varepsilon) = [Z_{\Gamma'}^{\Gamma'}(\zeta, \varepsilon)] \cdot \tilde{\Psi}_{G,\Gamma}((\Phi_{\Gamma'}^{\Gamma'}(\zeta))^* \varepsilon),$$

where $[Z_{\Gamma'}^{\Gamma'}(\zeta, \varepsilon)]$ is the image of $Z_{\Gamma'}^{\Gamma'}(\zeta, \varepsilon)$ under the projection

$$\mathrm{GL}_{2r}(\mathbb{F}) \rightarrow \mathrm{PGL}_{2r}(\mathbb{F}).$$

Recall that $\Phi_{\Gamma'}^{\Gamma'}(\gamma)$ arises via a composition of three types of behavior: orientation changes, changing the tines (move V), and moving the base point (moves VI–VII). It suffices, therefore, to study each type of behavior, since we may simply compose the various $Z_{\Gamma'}^{\Gamma'}(\zeta, \varepsilon)$ under compositions. Let us check each of these individually.

- *Orientation changes:* we have already seen how orientation changes affect the elements $H(\tau)$ (the proof of Lemma 5.2) and $\lambda_{e,o}$ (by definition); it follows that we may take each

$$Z_{\Gamma'}^{\Gamma'}(\zeta, \varepsilon) = I$$

for any ε when $\zeta \in (\mathbb{Z}_2)^E$.

- *Tine changes (move V)*: Tine changes only affect the differentials of the generators of degree 2 and hence do not affect the representations at all. Again, we may take

$$Z_{\Gamma}^{\Gamma'}(\zeta, \varepsilon) = I$$

in this setting as well.

- *Moving the base point (moves VI–VII)*: move VI or VII simply moves the base point through a tine or thread. In particular, it follows that we may take

$$Z_{\Gamma}^{\Gamma'}(\zeta, \varepsilon) := M^{\varepsilon}(\nu_{\Gamma'}, \nu_{\Gamma}) = M^{\Phi^* \varepsilon}(\nu_{\Gamma'}, \nu_{\Gamma}),$$

where the two matrices on the right are the same because the representations ε and $\Phi^* \varepsilon$ take the same values on elements of the coefficient ring because Φ preserves the degree 1 generators (i.e., the maps ε_{Π} and $(\Phi^* \varepsilon)_{\Pi}$ are the same, as ring homomorphisms from $\mathbb{Z}[\Pi]$ to $\text{Mat}_{r \times r}(\mathbb{F})$). ■

Given a tree T , we have a projection $\pi_T: (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+) \rightarrow (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+)$ inducing pullbacks

$$(\pi_T)^*: \text{Rep}_r(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+; \mathbb{F}) \rightarrow \text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F})$$

and

$$(\pi_T)^*: \mathcal{M}_r^{\text{Rep}}(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+; \mathbb{F}) \rightarrow \mathcal{M}_r^{\text{Rep}}(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F}).$$

In particular, we obtain the maps

$$\tilde{\Psi}_{G,\Gamma,T} := \tilde{\Psi}_{G,\Gamma} \circ (\pi_T)^*$$

and

$$\Psi_{G,\Gamma,T} := \Psi_{G,\Gamma} \circ (\pi_T)^*,$$

the latter of which is the one occurring in Theorem 1.12.

Recall that if we wish to change the tree, then we have dg-isomorphisms

$$\Phi_T^{T'}(\gamma): (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+) \rightarrow (\mathcal{B}_{G,T'}^+, \partial_{G,\Gamma,T'}^+).$$

We will prove functoriality with respect to these transformations, which will easily allow us to project to the version in which we have included the tree.

Lemma 7.2. *For any fixed garden Γ on G and any homotopy class of path γ from $*_T$ to $*_{T'}$ in Λ_G , we have*

$$(\tilde{\Phi}_T^{T'}(\gamma))^*: \text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F}) \rightarrow \text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F})$$

being just the identity.

Proof. Recall that $\tilde{\Phi}_T^{T'}(\gamma)$ acts on degree 0 generators as $\tilde{\mathcal{C}}_T^{T'}(\gamma)$, given by

$$(\tilde{\mathcal{C}}_T^{T'}(\gamma))(\lambda_{e,o}) = \beta_T^{T'}(f, \gamma, +) \cdot \lambda_{e,o} \cdot \beta_T^{T'}(g, \gamma, -)^{-1}.$$

Recall further that the words β were associated to certain loops of paths on Λ_G . In particular, the terms β appear as the diagonal terms in the matrix of words for these loops. Upon applying any representation ε , we find that $\varepsilon_\Pi(\beta) = I$. It follows now that if $\varepsilon \in \text{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F})$ and $\lambda_{e,o} \in \Pi$, then

$$\varepsilon_\Pi(\lambda_{e,o}) = \varepsilon_\Pi((\tilde{\mathcal{C}}_T^{T'}(\gamma))(\lambda_{e,o})).$$

Hence, $\varepsilon = (\tilde{\Phi}_T^{T'}(\gamma))^* \varepsilon$. ■

Corollary 7.3. *For any of the dg-isomorphisms $\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma)$, we have*

$$\Psi_{G,\Gamma,T} \circ (\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma))^* = \Psi_{G,\Gamma',T'}.$$

Proof. We have proved that $\Psi_{G,\Gamma}$ commutes with orientation changes and moves V–VII, as well as the morphisms $\tilde{\Phi}_T^{T'}(\gamma)$. Every $\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma)$ descends from a composition of these simple changes; i.e., $\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma)$ fits into a diagram of the form

$$\begin{array}{ccc} (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+) & \xrightarrow{\tilde{\Phi}_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma)} & (\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma'}^+) \\ \pi_T \downarrow & & \downarrow \pi_{T'} \\ (\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+) & \xrightarrow{\Phi_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma)} & (\mathcal{B}_{G,T'}^+, \partial_{G,\Gamma',T'}^+) \end{array}$$

with the property that $\Psi_{G,\Gamma} \circ (\tilde{\Phi}_{\Gamma,T}^{\Gamma',T'}(\zeta, \gamma))^* = \Psi_{G,\Gamma'}$. In a similar manner as in the composition property of Lemma 3.2, the desired equation is induced from this one. ■

7.3. The colors-to-representations map

In this final subsection, we prove the following proposition, which is enough to prove our desired theorem.

Proposition 7.4. *The maps*

$$\Psi_{G,\Gamma,T}: \mathcal{M}_r^{\text{Rep}}(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+; \mathbb{F}) \rightarrow \mathcal{M}_r^{\text{Col}}(G; \mathbb{F})$$

are bijections for each choice of Γ and T .

Proof of Theorem 1.12 (assuming Proposition 7.4). By Proposition 7.4, we have that the maps $\Psi_{G,\Gamma,T}$ are bijections. They were proved to satisfy the functoriality property in Corollary 7.3. ■

In the rest of this subsection, we prove Proposition 7.4, essentially by constructing inverse maps. Although we wish to find an inverse to $\Psi_{G,\Gamma,T}$, we will work with representations instead of moduli of representations, until we need to.

To begin, we wish to understand which elements of $(\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)}$ belong in the image of M^Γ . Instead of attacking this directly, note that we may encode more information in the construction of M^Γ as follows. For any given representation $\varepsilon \in \mathrm{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F})$, we may choose to encode

$$\varepsilon_\Pi(H(\tau_f(v))) \in \mathrm{GL}_r(\mathbb{F})$$

for each thread $\tau_f(v)$, as well as $\varepsilon_\Pi(\lambda_{e,\circ}) \in \mathrm{GL}_r(\mathbb{F})$ for each generator $\lambda_{e,\circ} \in \Pi$ corresponding to some oriented edge $(e, \circ) \in \tilde{E}$. In other words, we have a modified morphism

$$\tilde{M}^\Gamma: \mathrm{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F}) \rightarrow (\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\mathrm{Thr}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\tilde{E}},$$

where we have denoted by $\mathrm{Thr}(G)$ the collection of threads. We may write this map in terms of the three factors on the right as

$$\tilde{M}^\Gamma = (M^\Gamma, H^\Gamma, \Lambda^\Gamma),$$

where M^Γ is the map described previously, and

$$H^\Gamma(\varepsilon, \tau) := \varepsilon_\Pi(H(\tau)), \quad \Lambda^\Gamma(\varepsilon, \lambda_{e,\circ}) := \varepsilon_\Pi(\lambda_{e,\circ})$$

for all $\tau \in \mathrm{Thr}(G)$ and $(e, \circ) \in \tilde{E}$, respectively.

Lemma 7.5. *The map $\tilde{M}^\Gamma$ is injective.*

Proof. We have that $\Lambda^\Gamma(\varepsilon)$ is precisely encoding the images $\varepsilon_\Pi(\lambda_{e,\circ})$. But since the various $\lambda_{e,\circ}$ generate Π , and ε is only non-trivial on elements of Π , we have that $\Lambda^\Gamma(\varepsilon)$ alone is enough to recover ε . ■

An element in the image of $\tilde{M}^\Gamma$ satisfies a number of properties which are clear by construction:

- *Base-point fixing:* if $v_\Gamma \in \mathrm{Tri}(G)$ is the triangle containing the base point of the garden $M^\Gamma(v_\Gamma) = I$ (by which we mean this is the case for every ε).
- *Edge relation:* for every edge e , if $\circ$ is the orientation of e occurring in Γ , and v_1, v_2 are the adjacent triangles to the left and right of the edge, respectively, then

$$M^\Gamma(v_2) = M^\Gamma(v_1) \begin{pmatrix} 0 & \Lambda^\Gamma(\lambda_{e,\circ})^{-1} \\ -\Lambda^\Gamma(\lambda_{e,\bar{\circ}}) & 0 \end{pmatrix}$$

(again for every ε).

- *Thread relation:* for every thread τ , with its canonical orientation from center to vertex, if v_1 and v_2 are the triangles on the left and right of the thread, respectively, then

$$M^\Gamma(v_2) = M^\Gamma(v_1) \begin{pmatrix} I & H^\Gamma(\tau) \\ 0 & I \end{pmatrix}$$

(for every ε).

- *Compatibility:* for every vertex $v \in V$ and every thread τ terminating at v , let e_1 , e_2 , and e_3 be the edges around E appearing counter-clockwise around v starting at the face containing τ , and let $\mathfrak{o}_1$, $\mathfrak{o}_2$, and $\mathfrak{o}_3$ be the outward orientations. Then,

$$H^\Gamma(\tau) = (-1)^{r_v} \Lambda^\Gamma(\lambda_{e_1, \mathfrak{o}_1})^{-1} \Lambda^\Gamma(\lambda_{e_2, \mathfrak{o}_2}) \Lambda^\Gamma(\lambda_{e_3, \mathfrak{o}_3}^{-1}).$$

In fact, the compatibility condition is extraneous, as the following lemma demonstrates.

Lemma 7.6. *Suppose that we have a vertex $v \in V$. Suppose that for each of the six triangles around the vertex, each incident thread, and each incident edge with orientation, we have associated an element of $\mathrm{GL}_{2r}(\mathbb{F})$, $\mathrm{GL}_r(\mathbb{F})$, and $\mathrm{GL}_r(\mathbb{F})$, with the property that these associated elements satisfy the edge and thread relations around each incident thread and edge. Then, the compatibility condition is automatically satisfied.*

Proof. For simplicity, we assume that all of the edges are oriented inwards in Γ ; we may change the orientation of an edge by negating the two elements of $\mathrm{GL}_r(\mathbb{F})$ associated to that edge. Then, just as in the proof of Lemma 5.11, we have that for v a triangle adjacent to the vertex such that the counter-clockwise loop around v traverses the threads and edges $\tau_1, e_1, \tau_2, e_2, \tau_3, e_3$, in that order, we find

$$M(v) = M(v) \cdot Z,$$

where, letting $H_j \in \mathrm{GL}_r(\mathbb{F})$ be the component corresponding to τ_j and $\Lambda_j \in \mathrm{GL}_r(\mathbb{F})$ be the component corresponding to e_j inwardly oriented,

$$Z := \begin{pmatrix} 1 & H_1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & -\Lambda_1^{-1} \\ \bar{\Lambda}_1 & 0 \end{pmatrix} \begin{pmatrix} 1 & H_2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & \Lambda_2^{-1} \\ -\bar{\Lambda}_2 & 0 \end{pmatrix} \begin{pmatrix} 1 & H_3 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & -\Lambda_3^{-1} \\ \bar{\Lambda}_3 & 0 \end{pmatrix}.$$

Since $M^\Gamma(v) \in \mathrm{GL}_{2r}(\mathbb{F})$ is invertible, it follows that

$$Z = I.$$

Multiplying out and looking at the bottom right $r \times r$ matrix yields

$$I = \bar{\Lambda}_1 H_2 \bar{\Lambda}_2 \Lambda_3^{-1},$$

and hence,

$$H_2 = \bar{\Lambda}_1^{-1} \Lambda_3 \bar{\Lambda}_2,$$

which is one of the compatibility relations. The others follow by symmetry. $\blacksquare$

Let

$$V_r(G, \Gamma; \mathbb{F}) \subset (\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\mathrm{Thr}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\tilde{E}}$$

be the subset satisfying these relations simultaneously. As the lemma indicates, since the compatibility relation is extraneous, we will henceforth ignore it.

Proposition 7.7. *The image of $\tilde{M}^\Gamma$ is $V_r(G, \Gamma; \mathbb{F})$. Hence, there is an induced bijection*

$$\mathrm{Rep}_r(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+; \mathbb{F}) \cong V_r(G, \Gamma; \mathbb{F}).$$

Proof. As just described, the image clearly lands in this subset. Conversely, if we have an element of this subset, then we may simply read off the representation ε from Λ^Γ . It suffices to check that $\sum H^\Gamma(\tau) = 0$, where the sum is over threads around a center, but this is clear by using the thread relation around this center, and noting that for $H_1, \dots, H_m \in \mathrm{GL}_r(\mathbb{F})$, we have

$$\begin{pmatrix} I & H_1 \\ 0 & I \end{pmatrix} \cdots \begin{pmatrix} I & H_m \\ 0 & I \end{pmatrix} = \begin{pmatrix} I & \sum_{j=1}^m H_j \\ 0 & I \end{pmatrix}.$$

The last sentence follows because we have already proved that $\tilde{M}^\Gamma$ is injective. $\blacksquare$

Now, let us include the tree into the story. Notice that elements in the image of $\tilde{M}^\Gamma \circ (\pi_T)^*$ have the following further property:

- *Tree triviality:* if $e \in E \setminus T$, then $\Lambda^\Gamma(\lambda_{e,o}) = I$ for either orientation o of e .

Let $V_r(G, \Gamma, T; \mathbb{F}) \subset V_r(G, \Gamma; \mathbb{F})$ be the further subset of elements with the tree triviality property. We have

$$\begin{aligned} V_r(G, \Gamma, T; \mathbb{F}) &\subset (\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\mathrm{Thr}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\tilde{T}} \\ &\subset (\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\mathrm{Thr}(G)} \times (\mathrm{GL}_r(\mathbb{F}))^{\tilde{E}}, \end{aligned}$$

where the second inclusion is given by setting each coordinate of $(\mathrm{GL}_r(\mathbb{F}))^{\tilde{E} \setminus T}$ to the identity matrix $I \in \mathrm{GL}_r(\mathbb{F})$.

Proposition 7.8. *The image of $\tilde{M}^\Gamma \circ (\pi_T)^*$ is $V_r(G, \Gamma, T; \mathbb{F})$. Hence, there is an induced bijection*

$$\mathrm{Rep}_r(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+; \mathbb{F}) \cong V_r(G, \Gamma, T; \mathbb{F}).$$

Proof. The proof is essentially identical, noting that the only requirement of including the tree is that the edges not in the tree are sent to the identity. The last sentence follows because both $(\pi_T)^*$ and $\widetilde{M}^\Gamma$ are injective, and hence, their composition is injective. ■

Our description of $V_r(G, \Gamma, T; \mathbb{F})$ is a seemingly extraneous construction, in the sense that the coordinates H and Λ are simply disregarded in the map M^Γ ; we are interested in M^Γ instead of $\widetilde{M}^\Gamma$ after all. Our goal is to show that $V_r(G, \Gamma, T; \mathbb{F})$ is itself naturally in bijective correspondence with the image of $M^\Gamma \circ (\pi_T)^*$. To do this, we will essentially remove the tree and thread coordinates one by one. Suppose that

- T' a subtree of T (possibly including $T' = \emptyset$),
- $Y \subset \text{Thr}(G)$ a subset of threads.

Let us consider the set

$$V_r(G, \Gamma, T' \subset T, Y; \mathbb{F}) \subset (\text{GL}_{2r}(\mathbb{F}))^{\text{Tri}(G)} \times (\text{GL}_r(\mathbb{F}))^Y \times (\text{GL}_r(\mathbb{F}))^{\widetilde{T \setminus T'}}$$

satisfying essentially the same properties as before wherever they can, but where the edge and thread relations are replaced with the following weaker version:

- *Edge* relation:* we require the old edge relation to hold for edges in $E \setminus T'$ (as we do not have access to edge data for elements in T'). For edges in T' with adjacent triangles v_1 and v_2 , we require the weaker condition that the span of the first r columns of $M^\Gamma(v_1)$ be equal to the span of the last r columns of $M^\Gamma(v_2)$ and vice versa. Equivalently, there exist elements $\Lambda_1, \Lambda_2 \in \text{GL}_r(\mathbb{F})$ such that

$$M^\Gamma(v_2) = \begin{pmatrix} 0 & \Lambda_1^{-1} \\ -\Lambda_2 & 0 \end{pmatrix} M^\Gamma(v_1).$$

- *Thread* relation:* we require the old thread relation for elements of Y . For threads τ in $\text{Thr}(G) \setminus Y$, we require the weaker condition that there exists $H \in \text{GL}_r(\mathbb{F})$ such that, for the triangles v_1 and v_2 neighboring τ , we have

$$M^\Gamma(v_2) = \begin{pmatrix} I & H \\ 0 & I \end{pmatrix} M^\Gamma(v_1).$$

That is, instead of requiring that the data of the values of H and Λ be encoded, we require that, locally, such values exist (we do not require uniqueness; the point is essentially that uniqueness is automatic). Notice that when $T' = \emptyset$ and $Y = \text{Thr}(G)$, we have

$$V_r(G, \Gamma, T; \mathbb{F}) = V_r(G, \Gamma, \emptyset \subset T, \text{Thr}(G); \mathbb{F}).$$

We prove now that we may increase T' or decrease Y at will and still end up with an isomorphic set under the corresponding projection.

Lemma 7.9. *Suppose that T' , T'' are two subtrees of T with $T' \subset T''$, and suppose that $Y'' \subset Y' \subset \text{Thr}(G)$. Consider the projection*

$$\begin{array}{c} (\text{GL}_{2r}(\mathbb{F}))^{\text{Tri}(G)} \times (\text{GL}_r(\mathbb{F}))^{Y'} \times (\text{GL}_r(\mathbb{F}))^{\widetilde{T \setminus T'}} \\ \downarrow \Pi \\ (\text{GL}_{2r}(\mathbb{F}))^{\text{Tri}(G)} \times (\text{GL}_r(\mathbb{F}))^{Y''} \times (\text{GL}_r(\mathbb{F}))^{\widetilde{T \setminus T''}}. \end{array}$$

Then, it induces a bijection of subsets

$$V_r(G, \Gamma, T' \subset T, Y'; \mathbb{F}) \cong V_r(G, \Gamma, T'' \subset T, Y''; \mathbb{F}).$$

Proof. It is clear that an element of $V_r(G, \Gamma, T' \subset T, Y'; \mathbb{F})$ lands in $V_r(G, \Gamma, T'' \subset T, Y''; \mathbb{F})$, since we simply require weaker relations to hold. We prove that this is a bijection by constructing an inverse. It suffices to fill in the data on $Y' \setminus Y''$ and $T'' \setminus T'$ so that the edge and thread relations hold. But notice that, by the edge* and thread* relations, we have that there exist choices which make the edge and thread relation holds on the nose. It suffices to check that they are unique.

For the case of edges, notice that we have

$$\begin{pmatrix} 0 & \Lambda_1^{-1} \\ -\Lambda_2 & 0 \end{pmatrix} = M(v_2) \cdot M(v_1)^{-1},$$

so Λ_1 and Λ_2 are obviously determined by the values of $M(v_1)$ and $M(v_2)$. The same reasoning works for the threads. ■

In particular, we have that

$$\begin{aligned} \text{Rep}_r(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+; \mathbb{F}) &\cong V_r(G, \Gamma, T; \mathbb{F}) \\ &\cong V_r(G, \Gamma, \emptyset \subset T, \text{Thr}(G)) \\ &\cong V_r(G, \Gamma, T \subset T, \emptyset; \mathbb{F}). \end{aligned}$$

But this last set, $V_r(G, \Gamma, T \subset T, \emptyset; \mathbb{F})$, is just the image of $M^\Gamma \circ (\pi_T)^*$, simply because $V_r(G, \Gamma, T; \mathbb{F})$ was the image of $\widetilde{M}^\Gamma \circ (\pi_T)^*$, and the bijection between the two is just the projection down to $(\text{GL}_r(\mathbb{F}))^{\text{Tri}(G)}$, the same projection which when composed with $\widetilde{M}^\Gamma$ yields M^Γ .

Finally, we are in a position to build an inverse map to $\Psi_{G,\Gamma,T}$. Instead of constructing the inverse with image in $\text{Rep}_r(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^+; \mathbb{F})$, we will instead construct it to have image in our convenient model $V_r(G, \Gamma, T \subset T, \emptyset; \mathbb{F})$, where we only need to check the edge* and thread* relations.

By the functoriality of Corollary 7.3, we may assume that the triangle ν_Γ containing the base point is adjacent to the vertex $*_T$. The reason we do this is that

$\tilde{\Psi}_{G,\Gamma,T}(\varepsilon) \in \text{Col}_r(G, \mathbb{F})$ always has that for the three faces around $*_T$ given as f_Γ (containing v_Γ), f'_Γ , and f''_Γ , we always have that

$$\begin{aligned}\tilde{\Psi}_{G,\Gamma,T}(\varepsilon, f_\Gamma) &= \text{Span} \begin{pmatrix} I \\ 0 \end{pmatrix}, \\ \tilde{\Psi}_{G,\Gamma,T}(\varepsilon, f'_\Gamma) &= \text{Span} \begin{pmatrix} 0 \\ I \end{pmatrix}, \\ \tilde{\Psi}_{G,\Gamma,T}(\varepsilon, f''_\Gamma) &= \text{Span} \begin{pmatrix} I \\ I \end{pmatrix}.\end{aligned}$$

This is simply because, in our construction of M^Γ , we could have taken all of our paths to be located in a small disk around the vertex, so this is a local computation.

Suppose that we are given an element $\chi \in \mathcal{M}_r^{\text{Col}}(G; \mathbb{F})$. Let us choose a representative $\tilde{\chi} \in \text{Col}_r(G; \mathbb{F}) \subset (\text{Gr}(r, 2r; \mathbb{F}))^F$, by which we mean that χ is the equivalence class of $\tilde{\chi}$ under the $\text{PGL}_{2r}(\mathbb{F})$ action such that the following conditions are satisfied:

$$\begin{aligned}\tilde{\chi}(f_\Gamma) &= \text{Span} \begin{pmatrix} I \\ 0 \end{pmatrix}, \\ \tilde{\chi}(f'_\Gamma) &= \text{Span} \begin{pmatrix} 0 \\ I \end{pmatrix}, \\ \tilde{\chi}(f''_\Gamma) &= \text{Span} \begin{pmatrix} I \\ I \end{pmatrix}.\end{aligned}$$

Notice that such $\tilde{\chi}$ exists, because the three faces involved are pairwise adjacent to each other, and any three pairs of transverse Grassmannian r -planes in $\mathbb{F}^{2r}$ may be arranged in this way by the action of $\text{GL}_{2r}(\mathbb{F})$. Note also that the elements of $\text{GL}_{2r}(\mathbb{F})$ which fix these three planes are of the form $\begin{pmatrix} g & 0 \\ 0 & g \end{pmatrix}$ for $g \in \text{GL}_r(\mathbb{F})$. We see that our choice of $\tilde{\chi}$ is therefore well defined up to action by $\text{GL}_r(\mathbb{F})$ on $\text{Col}_r(\mathbb{F})$; we recall that $\tilde{\Psi}_{G,\Gamma,T}$ intertwined the conjugation action on representations with this action on colorings, which is a good sign that we have picked the right gauge fixing condition.

Suppose now that we have a curve $\gamma: [0, 1] \rightarrow S^2$ satisfying the following properties:

- $\gamma(0)$ is the base point of Γ .
- $\gamma(1)$ is located in the interior of a triangle.
- γ intersects none of the edges of the tree T , nor the vertex $*_T$, nor any of the centers of the garden.
- γ is otherwise transverse to the edges (of $E \setminus T$) and threads that it crosses, say, at a finite set of times $0 < t_1 < \dots < t_n < 1$.

Associated to γ and $\tilde{\chi}$ is a locally constant function

$$\Theta(\gamma, \tilde{\chi}): [0, 1] \setminus \{t_1, \dots, t_n\} \rightarrow \mathrm{GL}_{2r}(\mathbb{F})$$

determined by first setting $\Theta(\gamma, \tilde{\chi})(0) = I$ and providing the following rules for how this function changes as we pass through t_i as follows:

- $\Theta(\gamma, \tilde{\chi})(0) = I$ is the identity matrix.
- If γ crosses through an edge at time t_i , then, for $0 < \varepsilon < \max(t_i - t_{i-1}, t_{i+1} - t_i)$, we have

$$\Theta(\gamma, \tilde{\chi})(t_i + \varepsilon) = \pm \Theta(\gamma, \tilde{\chi})(t_i - \varepsilon) \cdot \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix},$$

where the sign is determined by whether γ crosses through the edge positively or negatively.

- If γ crosses through a thread at time t_i from v_1 to v_2 , then similarly

$$\Theta(\gamma, \tilde{\chi})(t_i + \varepsilon) = \pm \Theta(\gamma, \tilde{\chi})(t_i - \varepsilon) \cdot \begin{pmatrix} I & H \\ 0 & I \end{pmatrix},$$

where H is the unique element of $\mathrm{GL}_r(\mathbb{F})$ such that the resulting product

$$\Theta(\gamma, \tilde{\chi})(t_i - \varepsilon) \cdot \begin{pmatrix} I & H \\ 0 & I \end{pmatrix}$$

has that the span of the last r columns is the element of $\tilde{\chi}$ attached to the face adjacent to v_2 .

Notice that it is easy to check inductively that, for all $t \neq t_i$, we have that if $\gamma(t) \in v \subset f$, with f' the face opposite v via an edge, then the span of the first r columns of $\Theta(\gamma, \tilde{\chi})(t)$ is given by $\tilde{\chi}(f)$, whereas the span of the last r columns is $\tilde{\chi}(f')$. Indeed, this is certainly true for $t = 0$ and remains true when we pass through an edge since we essentially swap the first and last r columns up to a sign, and also remains true when we pass through a thread by the very definition of how $\Theta^\Gamma(\gamma, \tilde{\chi})$ changes.

Lemma 7.10. *Suppose that v_γ is the triangle containing $\gamma(1)$. We wish to claim that if $v_\gamma = v_{\gamma'}$, then*

$$\Theta(\gamma, \tilde{\chi})(1) = \Theta(\gamma', \tilde{\chi})(1).$$

Proof. Notice that $S^2 \setminus T$ is topologically a disk, so γ and γ' are homotopic, in fact via a homotopy, so that the ending points remain in

$$v_\gamma = v_{\gamma'}.$$

It follows that we may change γ into γ' via a small handful of moves, each of which

we check does not affect the resulting value of $\Theta(\cdot, \tilde{\chi})(1)$:

- We may isotope γ in such a way that it still intersects the same threads and edges in the same order. The resulting value clearly does not change.
- We may push γ through a tangency with an edge or thread (as in moves I and II, but where the line is replaced by γ). When we go back and forth through an edge, we see that we end up at the same place, because we multiply by

$$\begin{pmatrix} 0 & \pm I \\ \mp I & 0 \end{pmatrix} \cdot \begin{pmatrix} 0 & \mp I \\ \pm I & 0 \end{pmatrix} = I_{2r},$$

where the signs exactly cancel because we traverse the edge with opposite orientations. On the other hand, suppose that we go back and forth through a thread, from v_1 to v_2 back to v_1 . Let Θ be the matrix we have at the beginning, when we are in v_1 . Then, we have that the matrix goes from Θ to $\Theta \cdot \begin{pmatrix} I & H \\ 0 & I \end{pmatrix}$ to $\Theta \cdot \begin{pmatrix} I & H \\ 0 & I \end{pmatrix} \cdot \begin{pmatrix} I & H' \\ 0 & I \end{pmatrix}$ for some uniquely determined $H, H' \in \text{GL}_r(\mathbb{F})$. But these must satisfy that the end result has that the last r columns spanning the face opposite v_1 via an edge. But inductively, Θ already has this property, so by the uniqueness, we have that

$$H' = -H.$$

Hence, after going back and forth, we arrive back at Θ .

- We may push γ through a center. As we go around the center, starting at some triangle v and ending up back at v , because our curve goes through threads, if we started at Θ , we accumulate a product

$$\Theta \cdot \begin{pmatrix} I & H_1 \\ 0 & I \end{pmatrix} \cdots \begin{pmatrix} I & H_k \\ 0 & I \end{pmatrix},$$

where k is the number of threads around the center, again with the property that the last r columns must span the same as the last r columns of Θ . We see that

$$H_k = -H_1 - \cdots - H_{k-1}$$

seems to work, and by the uniqueness, it suffices to check that this is indeed an element of $\text{GL}_r(\mathbb{F})$. But this is true because the last r columns of

$$\Theta \cdot \begin{pmatrix} I & H_1 \\ 0 & I \end{pmatrix} \cdots \begin{pmatrix} I & H_{k-1} \\ 0 & I \end{pmatrix}$$

span a subspace transverse to the span of the last r columns of Θ , by the fact that the last r columns span the element of $\tilde{\chi}$ on the opposite face.

- We may push γ through the vertex $*_T$ (as in move III). But because everything is uniquely determined as we go around $*_T$, everything matches a standard local model up to multiplication on the left by the original Θ . ■

Therefore, given $\tilde{\chi}$, we obtain an element in $\Theta(\tilde{\chi}) \in (\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)}$, taking the value $\Theta(\gamma, \tilde{\chi})(1)$ on v_γ for any path γ from the base point to v_γ . In fact, by construction, we have that $\Theta(\tilde{\chi})$ satisfies the edge* relations with respect to the full tree T and the thread* relations with respect to the empty set of threads. That is, we have

$$\Theta(\tilde{\chi}) \in V_r(G, \Gamma, T \subset T, \emptyset; \mathbb{F}).$$

Although $\tilde{\chi}$ was only determined from χ up to the action of $\mathrm{GL}_r(\mathbb{F})$, this action intertwined with the conjugation action on representations, or in turn the conjugation action on $(\mathrm{GL}_{2r}(\mathbb{F}))^{\mathrm{Tri}(G)}$ by the corresponding element $\begin{pmatrix} g & 0 \\ 0 & g \end{pmatrix}$. Hence, the map $\tilde{\chi} \mapsto \Theta(\tilde{\chi})$ descends to a map

$$\Theta_{G,\Gamma,T}: \mathcal{M}_r^{\mathrm{Col}}(G; \mathbb{F}) \rightarrow V_r(G, \Gamma, T \subset T, \emptyset; \mathbb{F}) / \mathrm{GL}_r(\mathbb{F}) \cong \mathcal{M}_r^{\mathrm{Rep}}(\mathcal{B}_{G,T}^+, \partial_{G,\Gamma,T}^{\mathcal{B}^+}; \mathbb{F}).$$

If we follow $\Theta(\tilde{\chi})$ back along the construction of $\tilde{\Phi}_{G,\Gamma,T}$, i.e., taking

$$\mathrm{Span}\left(\Theta(\tilde{\chi}) \cdot \begin{pmatrix} I \\ 0 \end{pmatrix}\right),$$

by construction, we simply arrive back at $\tilde{\chi}$, which implies that

$$\Psi_{G,\Gamma,T} \circ \Theta_{G,\Gamma,T} = 1.$$

Therefore, $\Psi_{G,\Gamma,T}$ is surjective. In fact, since $\Theta(\tilde{\chi})$ was essentially determined *uniquely* by requiring that

- the edge* relations hold,
- the thread* relations hold, and
- $\tilde{\Psi}_{G,\Gamma,T}(\Theta(\tilde{\chi}))$ matches $\tilde{\chi}$ on the three faces around $*_T$,

we have that actually $\Theta(\tilde{\chi})$ is the *unique* element of $V_r(G, \Gamma, T \subset T, \emptyset; \mathbb{F})$ such that

$$\tilde{\Psi}_{G,\Gamma,T}(\Theta(\tilde{\chi})) = \tilde{\chi}.$$

But the specific value of $\tilde{\chi}$ only depended up to the action of $\mathrm{GL}_r(\mathbb{F})$ (given χ), and hence, $\Theta(\tilde{\chi})$ only varies by the action of $\mathrm{GL}_r(\mathbb{F})$ as we change our choice of $\tilde{\chi}$. In particular, if we have

$$\Psi_{G,\Gamma,T}(\varepsilon) = \chi,$$

then that implies that

$$\tilde{\Psi}_{G,\Gamma,T}(\varepsilon) = \tilde{\chi}$$




Figure 18. The graph and garden for which the enlarged dg-algebra is computed in Appendix A.

for one of the lifts of χ , and in turn by this uniqueness, we have

$$\varepsilon = g \cdot \Theta(\tilde{\chi})$$

for some $g \in \mathrm{GL}_r(\mathbb{F})$. Hence, we have that $\Psi_{G,\Gamma,T}$ is injective. The following proof is a summary.

Proof of Proposition 7.4. We have proved that $\Psi_{G,\Gamma,T}$ is both injective and surjective, which is what we needed to show. ■

Remark 7.11. We have done slightly better than to prove Proposition 7.4; we have actually constructed an explicit inverse in the form of $\Theta_{G,\Gamma,T}$.

A. Computational example

We provide a sample computation of the enlarged dg-algebra $(\tilde{\mathcal{B}}_G^+, \tilde{\partial}_{G,\Gamma}^+)$ for the specific pair (G, Γ) of Figure 18 (with $g = 1$). We note that, with commutative coefficients, this computation was performed by Casals and Murphy [3, Section 5], so the reader is invited to compare our computation to theirs.

In the figure, the tines are labeled γ_j for $j = 0, 1, 2, 3$. We will write f_j for the generator corresponding to the center through which γ_j passes. We have also labeled

the edges and oriented all edges e_k for $k = 1, \dots, 6$. We will use the generators A_k and B_k as in Figure 4.

The differentials of the faces are as follows:

$$\begin{aligned}\tilde{\partial}f_0 &= B_1^{-1}A_3B_4^{-1} - A_6^{-1}B_2A_1^{-1} + A_4^{-1}B_5B_6^{-1}, \\ \tilde{\partial}f_1 &= B_3^{-1}A_4B_1^{-1} - A_1^{-1}B_6A_2^{-1} - B_2^{-1}A_5A_3^{-1}, \\ \tilde{\partial}f_2 &= -B_5^{-1}B_3B_2^{-1} - A_2^{-1}B_1A_6^{-1} + B_6^{-1}B_4A_5^{-1}, \\ \tilde{\partial}f_3 &= B_4^{-1}A_1B_3^{-1} - A_3^{-1}A_2B_5^{-1} + A_5^{-1}A_6A_4^{-1}.\end{aligned}$$

Meanwhile, we may write

$$\begin{pmatrix} \tilde{\partial}z & \tilde{\partial}w \\ \tilde{\partial}y & \tilde{\partial}x \end{pmatrix} = \sum_{j=0}^3 W(\gamma_j),$$

where we compute each $W(\gamma_j)$ by breaking each time γ_j into elementary pieces in which it either crosses an edge, thread, or face:

$$\begin{aligned}W(\gamma_0) &= \begin{pmatrix} 0 & f_0 \\ 0 & 0 \end{pmatrix}, \\ W(\gamma_1) &= \begin{pmatrix} 1 & A_4^{-1}B_5B_6^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & A_6^{-1} \\ -B_6 & 0 \end{pmatrix} \begin{pmatrix} 1 & A_2^{-1}B_1A_6^{-1} \\ 0 & 1 \end{pmatrix} \\ &\quad \cdot \begin{pmatrix} 0 & A_2^{-1} \\ -B_2 & 0 \end{pmatrix} \begin{pmatrix} 0 & f_1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & A_1^{-1} \\ -B_1 & 0 \end{pmatrix} \begin{pmatrix} 1 & B_3^{-1}A_4B_1^{-1} \\ 0 & 1 \end{pmatrix}, \\ W(\gamma_2) &= \begin{pmatrix} 1 & A_4^{-1}B_5B_6^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & A_6^{-1} \\ -B_6 & 0 \end{pmatrix} \begin{pmatrix} 0 & f_2 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & A_2^{-1} \\ -B_2 & 0 \end{pmatrix} \\ &\quad \cdot \begin{pmatrix} 1 & -B_2^{-1}A_5A_3^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & B_3^{-1}A_4B_1^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & A_1^{-1} \\ -B_1 & 0 \end{pmatrix} \begin{pmatrix} 1 & B_1^{-1}A_3B_4^{-1} \\ 0 & 1 \end{pmatrix}, \\ W(\gamma_3) &= \begin{pmatrix} 1 & A_4^{-1}B_5B_6^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & A_6^{-1} \\ -B_6 & 0 \end{pmatrix} \begin{pmatrix} 1 & B_6^{-1}B_4A_5^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & -B_5^{-1} \\ A_5 & 0 \end{pmatrix} \\ &\quad \cdot \begin{pmatrix} 0 & f_3 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & A_3^{-1} \\ -B_3 & 0 \end{pmatrix} \begin{pmatrix} 1 & B_3^{-1}A_4B_1^{-1} \\ 0 & 1 \end{pmatrix} \\ &\quad \cdot \begin{pmatrix} 0 & A_1^{-1} \\ -B_1 & 0 \end{pmatrix} \begin{pmatrix} 1 & B_1^{-1}A_3B_4^{-1} \\ 0 & 1 \end{pmatrix}.\end{aligned}$$

One may multiply these out and use the relations between geometric elements to simplify. For example, the expression for $W(\gamma_3)$ simplifies to

$$W(\gamma_3) = \begin{pmatrix} 0 & 0 \\ -B_4 f_3 A_4 & 0 \end{pmatrix}.$$

This was expected by Lemma 4.8, since we may homotope γ_3 relative to its endpoints at infinity and without passing through other faces so that it crosses e_4 , then crosses f_3 , then crosses e_4 again. We see that f_3 only appears in the differential of the generator y , and not in the differential of the generators w , x , and z . As a sanity check, this matches with the Casals–Murphy computation in the commutative coefficient setting: only the differential of y has an f_3 term.

Acknowledgments. The author is grateful to a number of people who influenced this paper. Emmy Murphy initially suggested this project and provided helpful explanations of the underlying contact-geometric picture for the commutative case, especially in the early stages of this work. Discussions about Legendrian weaves with Roger Casals and Honghao Gao have influenced the expositional structure of this paper. Furthermore, encouragement from Honghao Gao helped to prevent this paper from entering development limbo. Almost all simplifications, clarifications, and corrections made from the original version of this article, as well as the inclusion of Section 1.4, were precipitated by the comments of two anonymous referees, to whom the author is grateful as well.

Funding. This work was partially supported by the National Science Foundation under grant DMS-1547145.

References

- [1] K. Appel and W. Haken, *Every planar map is four colorable*. Contemp. Math. 98, American Mathematical Society, Providence, RI, 1989 Zbl 0681.05027 MR 1025335
- [2] J. Asplund and T. Ekholm, *Chekanov–Eliashberg dg-algebras for singular Legendrians*. *J. Symplectic Geom.* **20** (2022), no. 3, 509–559 Zbl 1512.53083 MR 4563001
- [3] R. Casals and E. Murphy, *Differential algebra of cubic planar graphs*. *Adv. Math.* **338** (2018), 401–446 Zbl 1397.05042 MR 3861709
- [4] R. Casals and E. Zaslow, *Legendrian weaves: N -graph calculus, flag moduli and applications*. *Geom. Topol.* **26** (2022), no. 8, 3589–3745 Zbl 1521.53061 MR 4562568
- [5] B. Chantraine, G. Dimitroglou Rizell, P. Ghiggini, and R. Golovko, *Noncommutative augmentation categories*. In *Proceedings of the Gökova Geometry-Topology Conference 2015*, pp. 116–150, Gökova Geometry/Topology Conference (GGT), Gökova, 2016 Zbl 1358.16009 MR 3526840

- [6] B. Chantraine, L. Ng, and S. Sivek, [Representations, sheaves and Legendrian \$\(2, m\)\$ torus links](#). *J. Lond. Math. Soc. (2)* **100** (2019), no. 1, 41–82 Zbl [1428.57011](#) MR [3999682](#)
- [7] Y. Chekanov, [Differential algebra of Legendrian links](#). *Invent. Math.* **150** (2002), no. 3, 441–483 Zbl [1029.57011](#) MR [1946550](#)
- [8] T. Ekholm, [Morse flow trees and Legendrian contact homology in 1-jet spaces](#). *Geom. Topol.* **11** (2007), 1083–1224 Zbl [1162.53064](#) MR [2326943](#)
- [9] T. Ekholm, J. Etnyre, and M. Sullivan, [The contact homology of Legendrian submanifolds in \$\mathbb{R}^{2n+1}\$](#) . *J. Differential Geom.* **71** (2005), no. 2, 177–305 Zbl [1103.53048](#) MR [2197142](#)
- [10] T. Ekholm, J. Etnyre, and M. Sullivan, [Non-isotopic Legendrian submanifolds in \$\mathbb{R}^{2n+1}\$](#) . *J. Differential Geom.* **71** (2005), no. 1, 85–128 Zbl [1098.57013](#) MR [2191769](#)
- [11] T. Ekholm, J. Etnyre, and M. Sullivan, [Legendrian contact homology in \$P \times \mathbb{R}\$](#) . *Trans. Amer. Math. Soc.* **359** (2007), no. 7, 3301–3335 Zbl [1119.53051](#) MR [2299457](#)
- [12] T. Ekholm and Y. Lekili, [Duality between Lagrangian and Legendrian invariants](#). *Geom. Topol.* **27** (2023), no. 6, 2049–2179 Zbl [07781510](#) MR [4634745](#)
- [13] Y. Eliashberg, [Invariants in contact topology](#). In *Proceedings of the International Congress of Mathematicians, Vol. II (Berlin, 1998)*, pp. 327–338, Extra Vol. II, 1998 Zbl [0913.53010](#) MR [1648083](#)
- [14] Y. Eliashberg and L. Polterovich, [Local Lagrangian 2-knots are trivial](#). *Ann. of Math. (2)* **144** (1996), no. 1, 61–76 Zbl [0872.57030](#) MR [1405943](#)
- [15] J. B. Etnyre, L. L. Ng, and J. M. Sabloff, [Invariants of Legendrian knots and coherent orientations](#). *J. Symplectic Geom.* **1** (2002), no. 2, 321–367 Zbl [1024.57014](#) MR [1959585](#)
- [16] H. Federer and B. Jónsson, [Some properties of free groups](#). *Trans. Amer. Math. Soc.* **68** (1950), 1–27 Zbl [0035.29402](#) MR [32643](#)
- [17] S. Ganatra, J. Pardon, and V. Shende, [Microlocal Morse theory of wrapped Fukaya categories](#). 2020, arXiv:[1809.08807v2](#)
- [18] H. Gao, [Simple sheaves for knot conormals](#). *J. Symplectic Geom.* **18** (2020), no. 4, 1027–1070 Zbl [1469.57007](#) MR [4174293](#)
- [19] H. Gao, [Augmentations and sheaves for links](#). 2021, arXiv:[2109.01285](#)
- [20] T. Kálmán, [Contact homology and one parameter families of Legendrian knots](#). *Geom. Topol.* **9** (2005), 2013–2078 Zbl [1095.53059](#) MR [2209366](#)
- [21] D. Nadler, [Microlocal branes are constructible sheaves](#). *Selecta Math. (N.S.)* **15** (2009), no. 4, 563–619 Zbl [1197.53116](#) MR [2565051](#)
- [22] D. Nadler and E. Zaslow, [Constructible sheaves and the Fukaya category](#). *J. Amer. Math. Soc.* **22** (2009), no. 1, 233–286 Zbl [1227.32019](#) MR [2449059](#)
- [23] L. Ng, D. Rutherford, V. Shende, S. Sivek, and E. Zaslow, [Augmentations are sheaves](#). *Geom. Topol.* **24** (2020), no. 5, 2149–2286 Zbl [1457.53064](#) MR [4194293](#)
- [24] J. Nielsen, Om Regning med ikke-kommutative Faktorer og dens Anvendelse i Gruppeteorien. *Mat. Tidsskrift B* **1921** (1921), 77–94 Zbl [48.0123.03](#)
- [25] J. Nielsen, On calculation with noncommutative factors and its application to group theory. *Math. Sci.* **6** (1981), no. 2, 73–85 Zbl [0489.62020](#) MR [621096](#)
- [26] D. Rutherford and M. Sullivan, [Sheaves via augmentations of Legendrian surfaces](#). *J. Homotopy Relat. Struct.* **16** (2021), no. 4, 703–752 Zbl [1496.57028](#) MR [4343078](#)

- [27] V. Shende, D. Treumann, H. Williams, and E. Zaslow, [Cluster varieties from Legendrian knots](#). *Duke Math. J.* **168** (2019), no. 15, 2801–2871 Zbl [1475.53094](#) MR [4017516](#)
- [28] D. Treumann and E. Zaslow, [Cubic planar graphs and Legendrian surface theory](#). *Adv. Theor. Math. Phys.* **22** (2018), no. 5, 1289–1345 Zbl [07430949](#) MR [3952351](#)

Received 26 August 2022.

Kevin Sackel

Department of Mathematics and Statistics, University of Massachusetts Amherst,
Lederle Graduate Research Tower, 1623D, 710 N. Pleasant Street, Amherst, MA 01003-9305,
USA; ksackel@umass.edu