Advances in Mathematics 412 (2023) 108796

Contents lists available at ScienceDirect

Advances in Mathematics

journal homepage: www.elsevier.com/locate/aim

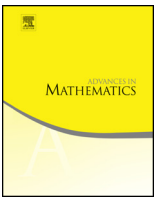

Tensor diagrams and cluster combinatorics at punctures ☆

Chris Fraser *, Pavlo Pylyavskyy

School of Mathematics, University of Minnesota, Minneapolis, MN, USA

ARTICLE INFO

Article history:
Received 21 October 2021
Received in revised form 31 August 2022
Accepted 31 October 2022
Available online 6 December 2022
Communicated by B. Keller

MSC:
primary 13F60
secondary 05E99, 57M50

Keywords:
Cluster algebras
Skein algebra
Web
Tensor diagram

ABSTRACT

Fock and Goncharov introduced a family of cluster algebras associated with the moduli of SL_k-local systems on a marked surface with extra decorations at marked points. We study this family from an algebraic and combinatorial perspective, emphasizing the structures which arise when the surface has punctures. When $k = 2$, these structures are the tagged arcs and tagged triangulations of Fomin, Shapiro, and Thurston. For higher k, the tagging of arcs is replaced by a Weyl group action at punctures discovered by Goncharov and Shen. We pursue a higher analogue of a tagged triangulation in the language of tensor diagrams, extending work of Fomin and the second author, and we formulate skein-algebraic tools for calculating in these cluster algebras. We analyze the finite mutation type examples in detail.

1. Introduction

The aim of *cluster combinatorics* is to identify, in a particular cluster algebra of interest, a labeling set for all of the cluster variables and a rule which specifies when two

☆ P. P. was partially supported by NSF grants DMS-1148634, DMS-1351590, and Sloan Fellowship. C. F. was supported by NSF grant DMS-1745638 and a Simons Travel Fellowship.

* Corresponding author.
E-mail addresses: cfraser@umn.edu (C. Fraser), ppylyavs@umn.edu (P. Pylyavskyy).

https://doi.org/10.1016/j.aim.2022.108796
0001-8708/

such cluster variables are *compatible* in the sense of residing in the same cluster. With these rules uncovered, one knows the *cluster monomials*, a certain subset of a "canonical" basis for the cluster algebra. Contrary to the name, there are several constructions of canonical bases valid in various settings, e.g. the dual canonical, dual semicanonical, generic, Satake, and theta bases. These bases do not all coincide with one another, but it is believed, and in some instances proved, that each of these bases contains the cluster monomials.

The *cluster algebras from surfaces* [9] have been very widely studied, perhaps especially because their cluster combinatorics is completely understood. The story is very clean when the underlying marked surface $\mathbb{S}$ has no punctures. In this case, cluster variables in the associated cluster algebra $\mathscr{A}(\mathbb{S})$ are labeled by isotopy classes of arcs connecting marked points in $\mathbb{S}$. Two cluster variables are compatible exactly when their corresponding arcs are noncrossing. The clusters are the maximal sets of noncrossing arcs, i.e. the ideal triangulations of $\mathbb{S}$.

When $\mathbb{S}$ has punctures, the above description captures certain clusters in $\mathscr{A}(\mathbb{S})$, but not all of them. Fomin, Shapiro and Thurston identified the combinatorial wrinkle closing this gap. Each end of an arc which is at a puncture, i.e. not at a marked point on the boundary, needs to be *tagged* as either a *plain* or *notched* end. There is a compatibility notion for tagged arcs – the arcs themselves should be noncrossing and their taggings must satisfy some extra conditions. A *tagged triangulation* is a maximal set of pairwise compatible tagged arcs. With these notions in place, one has that the cluster variables in $\mathscr{A}(\mathbb{S})$ are exactly the tagged arcs and the clusters are exactly the tagged triangulations. (We exclude here the exceptional case of once-punctured closed surfaces.)

The above cluster algebras $\mathscr{A}(\mathbb{S})$ fit into a broader family of cluster algebras $\mathscr{A}(G, \mathbb{S})$ indexed by a choice of surface $\mathbb{S}$ and simple Lie group G. Clusters in $\mathscr{A}(G, \mathbb{S})$ provide rational coordinate charts on the moduli space of *decorated* G-local systems on $\mathbb{S}$ and are part of an approach to Higher Teichmüller theory of $\mathbb{S}$. These cluster algebras were introduced by Fock and Goncharov in type A [6] and extended to other Lie types in [24,16]. They include the aforementioned $\mathscr{A}(\mathbb{S})$ in the special case $\mathscr{A}(\mathbb{S}) = \mathscr{A}(\mathrm{SL}_2, \mathbb{S})$.

The cluster algebras $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ (and their dual "$\mathcal{X}$-space") recover important structures in representation theory, for example the Knutson-Tao hive polytopes, the Schützenberger involution on semistandard Young tableaux, and the quantum group with its braid group action, see e.g. [15,14,16,20,31,26]. The cluster combinatorics of these cluster algebras is much more wild once $k > 2$ and is not yet well understood.

This paper presents conjectures and theorems concerning cluster combinatorics of $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$. We focus on the structures which arise when $\mathbb{S}$ has punctures, i.e. on the higher rank generalizations of tagged arcs and tagged triangulations. We now summarize our results and conjectures in a different order than they appear in the body of the paper. We use the following standard notations: T, P, and W for a maximal torus in G with its weight lattice and Weyl group; $[k] := \{1, \ldots, k\}$ and $\binom{[k]}{a}$ for the collection of a-subsets drawn from $[k]$; f^λ for the number of standard Young tableaux of shape λ; $S_{g,h}$ for the closed genus g surface with h punctures.

One expects a connection between the cluster algebra $\mathscr{A}(G, \mathbb{S})$ and a related commutative algebra $\mathrm{Sk}(G, \mathbb{S})$, the *$G$-skein algebra on* $\mathbb{S}$. For example, when $k = 2$ and $\mathbb{S}$ has no punctures, one has $\mathscr{A}(\mathbb{S}) \subseteq \mathrm{Sk}(\mathrm{SL}_2, \mathbb{S})$ and this containment is typically an equality [28]. As a vector space, $\mathrm{Sk}(G, \mathbb{S})$ consists of formal linear combinations of *G-tensor diagrams* modulo the G-skein relations. A tensor diagram is a certain graph drawn on $\mathbb{S}$ with edges labeled by fundamental G-weights. A contractible piece of a G-tensor diagram encodes a map between tensor products of fundamental G-representations and G-skein relations capture the linear relations between such maps. Multiplication in $\mathrm{Sk}(G, \mathbb{S})$ is induced by superposition of diagrams. The *invariant* $[T]$ encoded by a diagram T is its class in the skein algebra.

Tensor diagram calculus when $G = \mathrm{SL}_2$ is classical and involves arcs and Ptolemy relations. The calculus for rank two groups G was pioneered in [23] and is very rich while still being tractable. For $\mathbb{S}$ with no punctures, Fomin and the second author put forth precise and as yet unproved conjectures formulating cluster combinatorics of $\mathscr{A}(\mathrm{SL}_3, \mathbb{S})$ in terms of SL_3 tensor diagram calculus [8,7]. See [30,11,13,18,4,21] for subsequent work furthering this viewpoint.

Any element of the skein algebra is a linear combination of planar diagram invariants. The conjectures from [8,7] predict more strongly that each cluster monomial is an invariant of a planar tensor diagram. This way of encoding a cluster variable is more compact than a Laurent polynomial expansion with respect to a particular seed, and moreover is independent of such a choice of seed. The conjectures also predict which planar invariants are cluster monomials. Tensor diagram calculus becomes more cumbersome when $k > 3$, but we expect that variations on these recipes still hold. At a first pass, one can think that the higher rank generalization of an "arc" is a planar tensor diagram.

Goncharov and Shen [14] identified the higher rank generalization of tagging cluster variables at punctures: if $\mathbb{S}$ has h many punctures, then there is an action of the h-fold product of Weyl groups $\prod_{i=1}^{h} W$ on the cluster algebra $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ by cluster automorphisms. We define a *tagged tensor diagram invariant* as the pullback of a diagram invariant $[T]$ along this action. These serve as the higher rank generalization of tagged arcs. When $k = 2$, one has $W = \mathbb{Z}/2\mathbb{Z}$ and the action interchanges plain and notched tags at punctures. Not every tagged tensor diagram invariant is a cluster variable although we expect that the converse of this statement is true. In this paper, we address the question of when two tagged diagram invariants, each of which is a cluster variable in $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$, are compatible.

We introduce the class of pseudotagged tensor diagrams (T, φ), a pair consisting of an SL_k tensor diagram T and a pseudotagging function φ. The edges e in a tensor diagram are labeled by fundamental weights, or equivalently by integers $\mathrm{wt}(e)$ drawn from $[k-1]$. The function φ is a choice of subset $\varphi(e) \in \binom{[k]}{\mathrm{wt}(e)}$ (that is, of a point in the W-orbit of the fundamental weight $\omega_{\mathrm{wt}(e)}$) for every edge e which is incident to a puncture. When $k = 2$ every edge has weight 1 and the two possible values of φ correspond to plain or notched tagging.

By appropriately incorporating the $\prod_i W$ action at punctures, a pseudotagged diagram (T, φ) encodes an *invariant* $[(T, \varphi)]$ in the fraction field of the skein algebra. A pseudotagged diagram is *tagged* provided it satisfies the following additional condition: for any edges e and e' ending at the same puncture, one has $\varphi(e) \subseteq \varphi(e')$ or $\varphi(e') \subseteq \varphi(e)$. This extra condition is equivalent to requiring that $[(T, \varphi)]$ is the pullback of $[T]$ along the Weyl group action at punctures. When $k = 2$, an example of a diagram which is pseudotagged but not tagged is given by a loop based at a puncture with one of its ends plain and the other end notched. Another example is a union of two arcs whose taggings disagree at a puncture.

The seemingly exotic class of pseudotagged tensor diagram invariants is natural for the following reason. Multiplication of invariants is implemented by superposition of diagrams in the sense that $[T_1][T_2] = [T_1 \cup T_2]$, and such a superposition is typically pseudotagged but not tagged.

Our first result (the "flattening theorem" Theorem 9.1) clarifies the relationship between these two classes of diagram invariants: any pseudotagged diagram invariant is a $\mathbb{Q}$-linear combination of tagged diagram invariants. In particular, the vector space spanned by tagged diagram invariants is an algebra. We expect that this algebra coincides with the (upper) cluster algebra $\mathscr{A}^{\mathrm{up}}(\mathrm{SL}_k, \mathbb{S})$ in most cases.

The proof of the flattening theorem is constructive: we identify an algebraic relation between pseudotagged diagram invariants and show that repeated application of this algebraic relation "flattens" any pseudotagged diagram invariant into a linear combination of tagged diagram invariants. The $\mathbb{Q}$-coefficients appearing in the flattening relation have binomial coefficients in their numerator and the f^λ's in their denominator. Certain exchange relations in $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ are instances of the flattening relation, and we expect that *every* exchange relation is a consequence of the skein relations together with the flattening relations. Fig. 8 illustrates a typical application of the flattening relation in terms of tensor diagrams.

Our first main cluster combinatorics conjecture (Conjecture 8.18) asserts that every cluster monomial in $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ is the invariant of a planar tagged diagram and also of a tagged diagram with no cycles on interior vertices. This conjecture extends those from [8,7] to higher rank, and more novelly to surfaces with punctures.

Our second main cluster combinatorics conjecture (Conjecture 8.20) posits a necessary and sufficient condition for a product of tagged tensor diagram invariants, each of which is a cluster variable, to determine a cluster monomial. For example, consider a set of cluster variables given by taggings (T, φ_i) of the same underlying tensor diagram T and whose tagging functions φ_i only disagree on a single edge e of T incident to some puncture p. Then we conjecture that the product of these variables is a cluster monomial. This conjecture is a higher rank version of the Fomin-Shapiro-Thurston compatibility rule for tagged arcs.

Our second result (the "spiral theorem" Theorem 9.6) is a skein-algebraic sense in which our two conjectures are consistent with one another. Consider a set of variables $[(T, \varphi_i)]$ as in the preceding paragraph. The union of these diagrams is pseudotagged

and not tagged. Conjecture 8.20 predicts that the corresponding pseudotagged invariant $\prod_i[(T, \varphi_i)]$ is a cluster monomial. Hence, by Conjecture 8.18, it should be the invariant of a tagged diagram with no interior cycles and also of a tagged planar diagram. The spiral theorem confirms the first of these expectations, and establishes the second expectation in a weak sense: we show that $\prod_i[(T, \varphi_i)]$ can be given by a tagged diagram which is "planar nearby the puncture p" in a certain sense. We give a stronger planarity statement when $k = 3$ in Proposition 9.11. We note that while the spiral theorem is motivated by our cluster combinatorics conjectures, the theorem itself is purely skein-theoretic. Fig. 9 illustrates the pictures underlying the spiral theorem.

Heading in a different direction, we turn our attention from the delicate setting of tensor diagram calculus towards a coarser setting which nonetheless captures some essential features of $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$. When $k = 2$, this means forgetting the topological information encoded in a tagged arc and remembering only how many ends it has at each puncture and how the ends are tagged. This data is the weight of the corresponding cluster variable with respect to a natural algebraic torus action on $\mathscr{A}(\mathbb{S})$. One can parse clusters in a similar way. At an even coarser level, each cluster comes in one of three flavors at any puncture p: either all arcs are plain at p, all are notched at p, or there is exactly one plain and one notched arc at p.

One might think that the tagging formalism is an unfortunate departure of cluster algebra theory from more aesthetically pleasing notions such as arcs and triangulations. To the contrary, tagging is the key ingredient in three notable phenomena in cluster algebra theory [9, Sections 9 and 10]. Namely:

- Each algebra $\mathscr{A}(S_{g,1})$ does not admit a reddening sequence; the g-vector fan is contained in a half-space.
- The exchange graph of each $\mathscr{A}(S_{g,2})$ admits a "long cycle," i.e. an element of its fundamental group which is not a product of the 4- and 5-cycles coming from finite type rank 2 cluster subalgebras.
- The cluster complex of each $\mathscr{A}(S_{g,h})$ is homotopy equivalent to an $(h-1)$-sphere, providing an example of an infinite cluster type whose cluster complex is not contractible.

These three phenomena rely heavily on the coarse structures alluded to above. The first relies on the fact that in order to pass from a tagged triangulation which is plain at p to one that is notched at p by mutations, we must pass through one that is both plain and notched at p as an intermediate step. The second relies on this, and also on the fact that when $\mathbb{S} = S_{g,2}$ (or more generally when $\mathbb{S} = S_{g,h}$), a tagged triangulation cannot have a plain and notched arc at every puncture. For the third, consider a sign vector $\varepsilon = \{\pm 1\}^h$ indicating a choice of plain or notched at each puncture. It determines a subcomplex X_ε of the cluster complex, namely the subcomplex on the tagged arcs whose tagging at every puncture weakly agrees with ε. Then X_ε, moreover any intersection of these, is

contractible, so that the cluster complex is homotopy equivalent to the h-cube by the nerve lemma from topology.

We propose higher analogues of these coarse structures with a view towards understanding the topology of the cluster complex and exchange graph of $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$. We introduce *decorated* ordered set partitions of $[k]$, namely an ordered set partition with a choice of sign for each block of cardinality at least three. We call them dosp's for short. We propose that a dosp at p is the higher rank generalization of the three flavors of tagged triangulation at p from above.

We also study the possible weights of clusters in $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ with respect to the natural T-action at each puncture. We call such a set of weights a *P-cluster*. The mutation rule for P-clusters is linear algebra with integer vectors and can easily be experimented with by hand or on a computer.

Our third main conjecture (Conjecture 5.4) proposes a necessary condition which any P-cluster must satisfy. It is a combinatorial shadow of our second cluster combinatorics conjecture from above. By the *good part of the exchange graph*, we mean those clusters which can be reached from the initial seed by mutations which never violate Conjecture 5.4. Our expectation is that this good part is in fact the entire exchange graph.

Our third main result (developed in Sections 6 and 7) associates a dosp at each puncture to any cluster in the good part of the exchange graph. The construction depends only on the P-cluster, i.e. not on the cluster variables themselves. We identify a mutation operation on dosp's which models mutation of clusters: performing a mutation in the good part of the exchange graph changes at most one of the dosps at the punctures and does so by a dosp mutation. When $k = 2$, this mutation operation captures the first fact from above, i.e. that we cannot pass from plain at p to notched at p by a mutation. We show moreover that *every* dosp mutation "comes from" a mutation of clusters when $\partial\mathbb{S}$ carries at least two boundary points.

To broaden our family of examples, we introduce a "Grassmannian version" of the Fock-Goncharov moduli space. In this version, marked points on boundary components are decorated by vectors rather than by affine flags. We use primes $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S})$ to denote the Grassmannian version of the cluster algebra. All results and conjectures in this paper work equally well for both versions $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ and $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S})$ of the cluster algebra.

As a fourth result, we explore our conjectures for those $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ and $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S})$ of finite mutation type. Building on prior work and some coincidences, there are only three cluster algebras which need to be understood. Each of these is an $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S})$, either when $\mathbb{S}$ is a once-punctured digon and $k = 3$ or 4 or when $\mathbb{S}$ is a once-punctured triangle and $k = 3$. In each case, we prove that any cluster can be moved to a finite set of clusters modulo the action of some explicit quasi cluster automorphisms using the ideas from [11, Section 10]. We carry out the needed finite check in the two $k = 3$ examples, verifying all of our conjectures in these cases. We view the success of these calculations in the $k = 3$ case as further evidence in support of the conjectures from [8,7]. The finite check required in the $k = 4$ case is lengthy and we did not carry it out.

Fifth and finally, we prove that the Weyl group action on $\mathscr{A}(\mathrm{SL}_k, S_{g,1})$ is cluster when $k > 2$. This case was left open in [14,16]. Modulo some missing extra details, one can conclude that the Donaldson-Thomas transformation on $\mathscr{A}(\mathrm{SL}_k, S_{g,1})$ is cluster when $k > 2$.

The paper is organized as follows.

Sections 2 through 4 provide background with some new definitions and results as needed to fill in gaps. Section 2 collects standard background on type A combinatorics, cluster algebras, and cluster algebras from surfaces. It ends with a discussion of P-clusters, which is not a new notion although our terminology is. Section 3 recalls the Fock-Goncharov moduli space, introduces its Grassmannian cousin, and discusses cluster algebras associated to both moduli spaces. Section 4 recalls the Weyl group action at punctures and proves that this action is cluster in the case of $\mathscr{A}'$ and also when $\mathbb{S} = S_{g,1}$ and $k > 2$.

Sections 5 through 7 concern the possible weights of cluster variables at punctures. Section 5 introduces a compatibility notion for weight vectors and states our main conjecture concerning P-clusters. Section 6 explains that when this P-cluster conjecture holds, then one can associate a dosp at every puncture to a cluster. It also shows that mutation of clusters induces a mutation of dosps. Section 7 studies the extent to which *every* dosp mutation is induced by a mutation of clusters. The answer is sensitive to the chosen surface $\mathbb{S}$.

Sections 8 and 9 concern tensor diagram calculus and its pseudotagged version. Section 8 introduces tensor diagrams and skein relations, pseudotaggings and taggings, and states our two main cluster combinatorics conjectures. Section 9 states and proves the flattening and spiral theorems and discusses the special case $k = 3$.

Section 10 discusses the finite mutation type examples.

Acknowledgements. We thank Ian Le, Greg Muller, Linhui Shen, and David Speyer for conversations which inspired parts of this work. We thank the referees for helpful suggestions and for spotting errors in our original proofs of Lemma 3.7 and Proposition 5.6.

2. Background

We collect background on cluster algebras, aspects of type A combinatorics, cluster algebras from surfaces, and P-clusters. All but the last of these is standard.

Convention. Throughout the paper we use the notation $\mathcal{M}$ when we have fixed in mind one of the two possible versions of moduli space of decorated G-local systems. As we will explain, the symbol $\mathcal{M}$ implicitly requires that one has made a choice of marked surface $\mathbb{S}$, type A complex simple Lie group $G = \mathrm{SL}_k$ for some k, and style of decoration at boundary points, either by affine flags (the "Fock − Goncharov version" of the moduli space) or by vectors ("the Grassmannian version").

With such an $\mathcal{M}$ fixed, we use the notation $\mathscr{A}(\mathcal{M})$ for the associated cluster algebra defined below. We use primes $\mathscr{A}(G, \mathbb{S})$ versus $\mathscr{A}'(G, \mathbb{S})$ to distinguish between the Fock-Goncharov and Grassmannian versions of the cluster algebra when this is needed.

2.1. Cluster algebras

We assume familiarity with the construction of a *cluster algebra* and *upper cluster algebra* from an initial seed inside an ambient field of rational functions, and with the terminology of mutable and frozen variables, clusters, cluster monomials, cluster complex, and exchange graph. Two cluster variables are *compatible* when they reside in the same cluster, i.e. when their product is a cluster monomial. The *cluster type* is the underlying mutation pattern of mutable subquivers.

A *cluster automorphism* is an algebra automorphism which separately permutes the mutable and frozen variables while preserving the clusters. The *cluster modular group* can be defined as the cluster automorphism group of the cluster algebra once frozen variables are specialized to 1. A *quasi cluster automorphism* is a looser notion of automorphism of cluster algebra which allows for "rescalings" of cluster and frozen variables by Laurent monomials in frozen variables [12].

We assume familiarity with the related constructs of reddening sequences, maximal green sequences, and cluster Donaldson-Thomas transformations. The first of these is a mutation sequence which reaches the cluster variables whose g-vectors are the negative of the g-vectors of the initial cluster variables, the second is a restricted class of such sequences in which each mutation is "positive" in a certain sense, and the third is a quasi cluster automorphism which can be computed by a reddening mutation sequence.

2.2. Marked surfaces and triangulations

Let $\mathbb{S} = (\mathbf{S}, \mathbb{M})$ be an oriented marked surface [9]. The set of *marked points* $\mathbb{M}$ decomposes into the set of *punctures* $\mathbb{M}_\circ := \mathbb{M} \cap \operatorname{int} \mathbf{S}$ and the set of *boundary points* $\mathbb{M}_\partial := \mathbb{M} \cap \partial\mathbf{S}$. Denote by $S_{g,h}$ the oriented closed genus g surface with h punctures and by $D_{n,h}$ the n-gon with h punctures.

We henceforth exclude the following surfaces which have no ideal triangulations: the sphere with 2 or fewer punctures and the unpunctured monogon or digon. We also exclude the once-punctured monogon $D_{1,1}$. This last exclusion is not essential – it is possible to assign a cluster algebra to the Fock-Goncharov version of $D_{1,1}$, at least when $k > 3$. We thank L. Shen for explaining this to us. However, explaining the details of this construction and accordingly modifying our proofs adds extra technicalities to our presentation which we have chosen to avoid.

An *arc* on $\mathbb{S}$ is an immersion $\gamma\colon [0,1] \to \mathbf{S}$, such that $\gamma(0), \gamma(1) \in \mathbb{M}$ and γ restricts to an injection $(0,1) \hookrightarrow \mathbf{S} \setminus (\mathbb{M} \cup \partial\mathbf{S})$. Such arcs are considered up to isotopy fixing $\mathbb{M}$. A *loop* is an arc whose endpoints coincide. The subset of $\partial\mathbf{S}$ which connects adjacent boundary points along a boundary component is a *boundary interval*. An arc which is isotopic to a union of (at least two) consecutive boundary intervals is a *boundary arch*. An arc which is not a boundary arch is a *spanning arc* – such an arc either has its two endpoints on two different boundary components or has at least one end at a puncture.

Two arcs are *noncrossing* if one can find representatives for their isotopy classes which do not intersect except perhaps at their endpoints. A *triangulation* Δ of $\mathbb{S}$ is a maximal set of noncrossing arcs. Such a Δ decomposes $\mathbb{S}$ into "triangles" whose sides are either arcs or boundary intervals and whose vertices are marked points.

A triangulation is *regular* if it contains no loop γ which encloses a single puncture p on either of its two sides. In a regular triangulation, each triangle has three distinct sides. Because we exclude once-punctured monogons, any $\mathbb{S}$ considered in this paper has a regular triangulation.

A regular triangulation is *taut* if each of its triangles has at most one side which is a boundary interval. Equivalently, a taut triangulation is one which has no boundary arches.

In the n-gon, *every* arc is a boundary arch, so there are no taut triangulations. On the other hand, provided $\mathbb{S}$ is not the n-gon, then it admits a taut triangulation. For example, the unique taut triangulation of a once-punctured n-gon is the one in which each boundary point is connected to the puncture by an arc.

2.3. Cluster algebras from surfaces

Fomin, Shapiro, and Thurston associated to any marked surface $\mathbb{S}$ a cluster algebra $\mathscr{A}(\mathbb{S})$ [9]. Initial seeds for $\mathscr{A}(\mathbb{S})$ are the $k = 2$ case of the construction presented in Section 3. In this section, we summarize the cluster combinatorics underlying $\mathscr{A}(\mathbb{S})$.

To *tag* an arc γ is to make a binary choice (either *plain* or *notched* tagging) at every endpoint of γ which is a puncture. There is no binary choice for endpoints which are boundary points. We indicate notched tagging using the bowtie symbol $\bowtie$. If γ is a loop based at p, we require that both ends of γ are tagged the same way. We also require that the underlying arc γ is not contractible to a puncture or a boundary interval and does not cut out a once-punctured monogon.

Distinct tagged arcs are *compatible* if the arcs which underlie them are noncrossing and the following additional conditions on their taggings are satisfied. If the underlying arcs do not coincide, then their taggings agree at any puncture which they have in common. If the underlying arcs do coincide, then the taggings disagree at exactly one end. For example, if γ is a loop based a puncture p and tagged plain at p, then any tagged arcs which are compatible with γ are also plain at p if they have an endpoint at p. A *tagged triangulation* of $\mathbb{S}$ is a maximal set of pairwise compatible tagged arcs.

Suppose that $\mathbb{S} \neq S_{g,1}$ for some $g \geq 1$. Then cluster variables (resp. clusters) in $\mathscr{A}(\mathbb{S})$ are in bijection with tagged arcs (resp. tagged triangulations) in $\mathbb{S}$. The case $\mathbb{S} = S_{g,1}$ is exceptional: in this case the cluster variables (resp. clusters) in $\mathscr{A}(\mathbb{S})$ are in bijection with the arcs (resp. triangulations) of $\mathbb{S}$.

In Section 10, we assume familiarity with the following nested sequence of subgroups associated to a given $\mathbb{S}$: the *pure mapping class group* $\mathrm{PMCG}(\mathbb{S})$, the mapping class group $\mathrm{MCG}(\mathbb{S})$, and the *tagged mapping class group* $\mathrm{MCG}^{\bowtie}(\mathbb{S})$. The first of these consists of mapping classes which fix each puncture pointwise, the second allows mapping classes

to permute the punctures, and the third allows compositions of mapping classes and tag-changing transformations (interchanging plain tags with notched tags) at punctures. These groups act on the cluster algebra $\mathscr{A}(\mathbb{S})$ by cluster automorphisms. In particular, the reader should be familiar with *half-twists* dragging one puncture over another puncture [5, Section 9]. The square of such a half-twist is a Dehn twist about the simple closed curve enclosing the two punctures.

2.4. Type A combinatorics

We use the following standard Lie theoretic terminology in the special case of $G = \mathrm{SL}(V)$ for a complex k-dimensional vector space V, with hopes that our constructions might also be valid in other Lie types.

The group G has Weyl group $W = W_k$ the symmetric group on k letters. It has a Coxeter presentation generated by the simple transpositions $s_i = (i, i+1)$ for $i \in [k-1]$ modulo the relations

$$s_i s_j = s_j s_i \text{ when } |i-j| > 1 \text{ and } s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}. \tag{1}$$

A *reduced expression* for an element $w \in W$ is a minimal-length factorization of w as a product of simple transpositions, i.e. and expression $w = s_{i_1} \cdots s_{i_\ell}$ with ℓ minimal.

A *weight* is an algebraic group homomorphism $T \to \mathbb{C}^*$, where $T \subset G$ is a choice of maximal torus. We can choose a basis $e_1, \ldots, e_k$ for V so that $T \subset G$ is identified with the set of unimodular diagonal matrices. Then a weight is a Laurent monomial map $\mathrm{diag}(t_1, \ldots, t_k) \mapsto t^\lambda$, for a *weight vector* $\lambda \in \mathbb{Z}^k/(1, 1, \ldots, 1)$. We denote by

$$P = \mathbb{Z}^k/\mathrm{span}(1, 1, \ldots, 1)$$

the *weight lattice*, inside the ambient real vector space $P_{\mathbb{R}} := P \otimes_{\mathbb{Z}} \mathbb{R}$. We have *fundamental weights* $\omega_i = e_1 + \cdots + e_i \in P$. For $S \subset [k]$, we have an *indicator vector* $\iota_S := \sum_{s \in S} e_s \in P$.

In examples, we often choose to encode weight vectors in "multiplicative notation," i.e. as (Laurent) monomials in the letters $a, b, c, \ldots$. For example, we would encode the weight vector $e_1 + 2e_3 + e_4$ as the monomial ac^2d.

An *ordered set partition* of $[k]$ (an *osp* for short) is a tuple $\Pi = B_1|B_2|\cdots|B_\ell$ of nonempty disjoint subsets with union $[k]$. The B_i are the *blocks* of the osp. Block of cardinality one and two are called *singletons* and *doubletons* respectively.

An osp $\Pi = B_1|\cdots|B_\ell$ encodes a *Weyl region* $C(\Pi) \subset P$ consisting of vectors whose coordinates satisfy the inequalities determined by the blocks of Π. That is, $(\lambda_1, \ldots, \lambda_k) \in C(\Pi)$ means that $\lambda_a = \lambda_b$ whenever a and b are in the same block of Π, and $\lambda_a > \lambda_b$ whenever $a \in B_i$, $b \in B_j$, and $i < j$. Given $\lambda \in P$, we write $\Pi(\lambda)$ for the osp which encodes its coordinate inequalities: $\lambda \in C_{\Pi(\lambda)}$.

Thinking instead of $C(\Pi)$ as a subset of $P_{\mathbb{R}}$, the Weyl regions are the faces of a geometric realization of a flag simplicial complex (the *Coxeter complex*). The partial

order which encodes the closure of Weyl regions is the *coarsening order* on ordered set partitions. That is, $C(\Pi')$ is in the closure of $C(\Pi)$ if each block of Π' is a union of several consecutive blocks of Π. We also say that Π *refines* Π' in this case. For example, $12|3|57|68|9$ is a refinement of $12|35678|9$. We use overlines to denote closures of Weyl regions, writing e.g. $\overline{C(\Pi)}$ for the union of the Weyl regions indexed by Π and those osp's which coarsen Π.

We write $\vee$ for the operation of taking the least upper bound of several ordered set partitions in the refinement partial order when this least upper bound exists. For example, $13|245|67 \vee 1|3|24567 \vee 13|24|567 = 1|3|24|5|67$. The Weyl region closure of the right hand side contains the three regions on the left hand side and is the minimal region with this property.

The group W acts on the weight lattice P by coordinate permutation:

$$w \cdot (\lambda_1, \ldots, \lambda_k) = (\lambda_{w(1)}, \ldots, \lambda_{w(k)}).$$

We say that distinct weight vectors λ and μ are *w-conjugate* if $w \cdot \lambda = \mu$. We typically use this notion when w is an involution, in which case the notion is symmetric in λ and μ. They are *conjugate* if they are *w-conjugate* for some w. For example, every indicator vector ι_S is conjugate to a fundamental weight $\omega_{|S|}$.

The *root system* Φ consists of the vectors (themselves known as *roots*) $e_a - e_b \in P$ for $a \neq b \in [k]$. Such a root determines a *root hyperplane* in $P_{\mathbb{R}}$ defined by the coordinate equality $\lambda_a = \lambda_b$. The action of the transposition $(ab) \in W$ on P is by reflection about this root hyperplane.

A subset $S \in \binom{[k]}{a}$ determines a *Grassmannian permutation* $w_S \in W$ whose first a symbols (resp. last $k-a$ symbols) in one-line notation are the elements of S (resp. $[k] \setminus S$) written in sorted order. We abbreviate $\ell(S) := \ell(w_S)$ where the latter is Coxeter length. The subset S also determines a lattice walk in the $a \times (k-a)$ rectangle consisting of unit steps in the $(0,1)$ direction at times S and in the $(1,0)$ direction at times $[k] \setminus S$. This lattice walk determines a Young diagram $\lambda(S)$, namely the boxes which are in the $a \times (k-a)$ rectangle and which lie weakly northwest of the lattice walk. It is known that the number of reduced expressions for the permutation w_S equals the number $f^{\lambda(S)}$ of standard Young Tableaux (SYT's) of shape $\lambda(S)$. Moreover, this set of reduced expressions is connected by the commutation moves $s_i s_j = s_j s_i$ only, see e.g. [2].

2.5. P-clusters

Cluster algebras often come with a grading by an abelian group such that each cluster variable is a homogeneous element. See [17] for a general discussion. The cluster algebras considered in this paper are graded by the direct sum $P^{\oplus \mathbb{M}_\circ}$ of copies of the G-weight lattice indexed by the punctures in $\mathbb{S}$.

Definition 2.1. A *P-seed* is a pair $(Q, \mathrm{wt}(v)_{v\in V(Q)})$ consisting of a quiver Q (with no frozen vertices) and a weight vector $\mathrm{wt}(v) \in P^{\oplus \mathbb{M}_\circ}$ for each vertex v of Q, subject to the *balancing constraint*

$$\sum_{u\to v \text{ in } Q} \mathrm{wt}(u) = \sum_{v\to u \text{ in } Q} \mathrm{wt}(u) \text{ for all vertices } v. \tag{2}$$

The *P-cluster* is the multiset of weights at the vertices $(\mathrm{wt}(v)_{v\in V(Q)})$. It consists of *P-cluster variables*.

To mutate a P-seed at a mutable vertex v, we perform quiver mutation at v and we modify the weight at v according to

$$\mathrm{wt}(v) \mapsto -\mathrm{wt}(v) + \sum_{u\to v \text{ in } Q} \mathrm{wt}(u)$$

leaving the other weights unchanged.

It is known that mutation of a P-seed again yields a P-seed [17]. The *P-exchange graph* is the graph whose vertices are the P-seeds and whose edges indicate mutations.

In what follows, the weight vector $\mathrm{wt}(v)$ represents the grading of a cluster variable $x(v)$ with respect to a natural $\prod_{\mathbb{M}_\circ} T$-action. The common weight of the two terms in (2) is the weight of the right hand side of the exchange relation describing mutation at the cluster variable $x(v)$ out of the current seed. Thus, there is a map from clusters to P-clusters which commutes with mutations.

3. Cluster nature of the moduli spaces

Fock and Goncharov associated a cluster algebra to the moduli space of decorated G-local systems. We review this construction and introduce its "Grassmannian cousin."

3.1. Fock-Goncharov version of the moduli space

Definition 3.1. Let V be a complex vector space of dimension k with a fixed volume form $\xi \in \bigwedge^k(V) \cong \mathbb{C}$. An *affine flag* F in V is a choice of tensors $F_{(i)} \in \bigwedge^i(V)$ subject to the normalization $F_{(k)} = \xi$ and also the *flag condition*: for each $i \in [k-1]$, there exists $v \in V$ such that $F_{(i)} \wedge v = F_{(i+1)}$.

Due to the normalization condition, these are sometimes called affine flags for SL_k.

The left G-action on V, hence on $\bigwedge^i(V)$, induces an action on affine flags. The stabilizer of any affine flag is a maximal unipotent subgroup in G. In particular, a matrix which stabilizes an affine flag is unipotent. (All of its eigenvalues equal 1.)

We have a right T-action on affine flags by rescaling the steps of the affine flag:

$$(F \cdot t)_{(i)} := t_1 \cdots t_i F_{(i)} \in \bigwedge^{(i)} V \text{ for } t = \operatorname{diag}(t_1, \ldots, t_k) \in T. \tag{3}$$

Remark 3.2. One can specify an affine flag F by setting $F_{(i)} := v_1 \wedge \cdots \wedge v_i$ where $v_1, \ldots, v_{k-1}$ is a linearly independent sequence of vectors. Each v_i is determined modulo $\operatorname{span}\{v_1, \ldots, v_{i-1}\}$. The sequence of subspaces $\operatorname{span} F_{(i)} := \operatorname{span}(v_1, \ldots, v_i)$ is a complete flag of subspaces depending only on F, not on the chosen v_i's. Two affine flags give rise to the same complete flag if and only if they are related by the right T-action on affine flags (3).

Suppose for the moment that k is odd. By a *G-local system* on $\mathbb{S}$ we will mean a copy V_p of the vector space V with its volume form ξ at every point $p \in \mathbb{S}$, and a *transport isomorphism* $\Xi_\alpha \colon V_p \to V_q$ for every path $\alpha \colon p \to q$ in $\mathbb{S}$, preserving the form ξ and depending only on the isotopy class of α.

Once we have chosen a choice of G-local system, every closed curve α based at a point $p \in \mathbb{S}$ determines a *monodromy matrix* $M_\alpha \in \operatorname{Aut}(V_p) \cong G$ given by the transport isomorphism along α. The local system is *unipotent* if the monodromy u_p of a simple closed curve contractible to p is a unipotent matrix.

A *decorated* G-local system on $\mathbb{S}$ is a G-local system on $\mathbb{S}$ together with a choice of affine flag $F_p \in V_p$ at every marked point $p \in \mathbb{M}$, subject to the extra constraint that the monodromy u_p around every puncture $p \in \mathbb{M}_\circ$ stabilizes the corresponding affine flag F_p at p. (The affine flags are "decorations.")

When k is even one makes essentially the same definitions, except that rather than considering local systems on $\mathbb{S}$, one considers local systems on the *punctured tangent bundle* of $\mathbb{S}$ whose monodromy in the fiber direction equals $-\mathrm{Id} \in G$. Fock and Goncharov call these *twisted local systems*. The space of twisted decorated G-local systems is isomorphic to the space of decorated G-local systems but the isomorphism involves making certain choices. Making a different choice of isomorphism will perhaps change certain cluster variables and tensor diagram invariants we discuss below by a sign. The details of these signs do not strongly affect the combinatorial results presented here. One can typically change the sign of a tensor diagram invariant as needed by "migrating hairs," i.e. the constructions below are typically flexible enough to capture any needed sign ambiguity.

We denote by $\mathcal{A}_{G,\mathbb{S}}$ the *moduli space* of (twisted) decorated G-local systems on $\mathbb{S}$, i.e. the set of decorated G-local systems considered up to simultaneous G-action.

Let g be the genus of the closed surface obtained from $\mathbf{S}$ by filling in all of the boundary components with disks, and let b be the number of such boundary components. Then

$$\dim \mathcal{A}_{G,\mathbb{S}} = (2g - 2 + b + |\mathbb{M}|)|G| - |\mathbb{M}_\partial| \dim U. \tag{4}$$

3.2. Grassmannian version of the moduli space

Elements in the moduli space just defined come with an affine flag at every marked point. We now introduce a variation in which boundary points carry a different type of decoration.

Definition 3.3. Denote by $\mathcal{A}'_{G,\mathbb{S}}$ the space of twisted G-local systems on $\mathbf{S}$ together with an affine flag F_p in V_p at every puncture p and a vector $v_p \in V_p$ at every boundary marked point p, considered up to simultaneous G-action.

When $k = 3$, this version of the moduli space matches the perspective of [8,7]. We introduce it here in hopes that it will appear in future contexts. More immediately, it provides us with a richer set of explicit examples in which our cluster combinatorics conjectures can be tested.

We have

$$\dim \mathcal{A}'_{G,\mathbb{S}} = (2g - 2 + b + |M_\circ|)|G| + k|M_\partial|. \tag{5}$$

Remark 3.4. The moduli space $\mathcal{A}'_{\mathrm{SL}_k, D_{n,0}}$ is the space of n-tuples of vectors in $\mathbb{C}^k$ modulo simultaneous SL_k action. Its algebra of regular functions is the homogeneous coordinate ring of the Grassmannian $\mathrm{Gr}(k,n)$ of k-subspaces in $\mathbb{C}^n$ in the Plücker embedding. More generally, for unpunctured surfaces $\mathbb{S}$, a point in $\mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$ consists of a tuple of vectors $v \in V$ together with a tuple of endomorphisms $M \in \mathrm{End}(V)$, with such tuples considered up to simultaneous G-action. This is a classical object of study in invariant theory, see [8] for some discussion.

Remark 3.5. We have the following relationships between the two flavors $\mathcal{A}$ versus $\mathcal{A}'$ of moduli space. First, the two moduli spaces coincide when $G = \mathrm{SL}_2$ or when $\mathbb{S}$ has no boundary.

Second, we will shortly associate a cluster algebra $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ (resp. $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S})$) to the moduli space $\mathcal{A}_{\mathrm{SL}_k,\mathbb{S}}$ (resp. $\mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$). There is a projection $\mathcal{A}_{\mathrm{SL}_k,\mathbb{S}} \twoheadrightarrow \mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$ induced by forgetting all but the first step of each affine flag. We conjecture that this map is a *cluster fibration* [25, Definition 3.2], i.e. that the inclusion of algebras $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S}) \hookrightarrow \mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ sends a cluster in the former to a partial cluster in the latter while moreover preserving the mutable part of the quiver. This holds when $\mathbb{S} = D_{n,0}$ is an n-gon [25, Proposition 3.4.(3)].

Third, we expect the following generalization of [11, Section 7] from n-gons to arbitrary surfaces. Suppose that every boundary component of $\mathbb{S}$ has an even number of marked points. Associate to it a new marked surface $\mathbb{S}'$ in which we change neither the underlying surface $\mathbf{S}$ nor the number of punctures but we increase the number of boundary points on each boundary component by a factor of $\frac{k}{2}$. We expect that the cluster algebras $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ and $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S}')$ are related by a quasi cluster isomorphism, namely the one

already proposed in [11]. One needs to verify that this candidate map indeed sends a seed in the former cluster algebra to one in the latter in the appropriate sense.

3.3. A family of quivers

Before introducing initial seeds for the moduli spaces $\mathcal{A}_{G,\mathbb{S}}$ and $\mathcal{A}'_{G,\mathbb{S}}$, we take a detour to introduce a family of quivers Q_k^s where $k \in \mathbb{Z}_{\geq 2}$ and $s \in [k-1]$. This wider family shows up in certain of our proofs.

We let $\hat{\mathbb{H}}_k = \{(a,b,c) \in \mathbb{N}^3 \colon a+b+c = k\}$ and $\mathbb{H}_k = \hat{\mathbb{H}}_k \setminus \{(k,0,0),(0,k,0),(0,0,k)\}$. We think of $\mathbb{H}_k$ as the vertex set of a quiver Q_k whose arrows are obtained by drawing small counterclockwise triangles as indicated in the examples of Q_4 and Q_5 below:

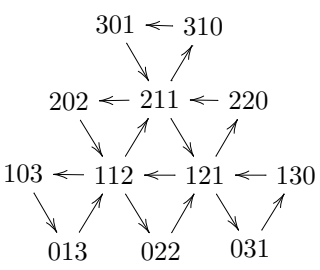

Q_4

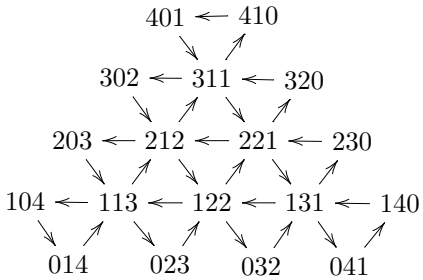

Q_5

The larger family of quivers Q_k^s for $s \in [k-1]$ is obtained from Q_k by identifying the vertices $(a,b,c) \sim (a,b+c,0) \in \mathbb{H}_k$ when $b = s$ and $a \neq 0$ and deleting the vertices (a,b,c) when $b \in [s+1, k-1]$. We draw below this family of quivers when $k = 5$:

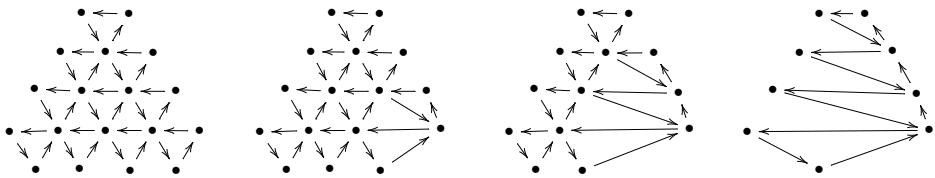

Q_5^4 Q_5^3 Q_5^2 Q_5^1

Note that $Q_k = Q_k^{k-1}$, and that each row of Q_k^s has no more than $s+1$ vertices.

3.4. Initial seeds for $\mathcal{A}_{G,\mathbb{S}}$

Fock and Goncharov described, for each choice of regular triangulation Δ of $\mathbb{S}$, a seed $\Sigma_k(\Delta) = (Q_k(\Delta), \mathbf{x}_k(\Delta))$ in the field of rational functions on $\mathcal{M} = \mathcal{A}_{G,\mathbb{S}}$. They proved that seeds from different triangulations are mutation-equivalent, so that the seeds $\Sigma_k(\Delta)$ together give rise to a cluster algebra $\mathscr{A}(G,\mathbb{S})$.

The initial quiver $Q_k(\Delta)$ is obtained as follows. To each triangle $\delta \in \Delta$ we associate a copy $Q_k(\delta)$ of the quiver Q_k from above. We identify the three vertices p_1, p_2, p_3 of the triangle δ with the three coordinates $(k,0,0)$, $(0,k,0)$, $(0,0,k)$ in $\hat{\mathbb{H}}_k$. We associate $k-1$ vertices of Q_k to each side of the triangle δ, associating e.g. the vertices $(a, k-a, 0)$ for $a \in [k-1]$ to the side $\overline{p_1 p_2}$ of δ.

If triangles δ and δ' have a common side $\gamma \in \Delta$ then we can glue the quiver fragments $Q_k(\delta)$ and $Q_k(\delta')$ by identifying the $k-1$ vertices associated to this shared side. The initial quiver $Q_k(\Delta)$ is the result of carrying out all such gluings:

$$Q_k(\Delta) = \coprod_{\delta \in \Delta} Q_k(\delta) / \sim, \tag{6}$$

where $\sim$ here indicates the gluing. The sides of triangles corresponding to boundary intervals are not glued during this process. We declare the $k-1$ vertices associated to such a boundary side as frozen vertices of $Q_k(\Delta)$, for a total of $(k-1)|\mathbb{M}_\partial|$ many frozen vertices. If any oriented two-cycles are created in the gluing process then they should be deleted from the quiver.

The initial cluster $\mathbf{x}_k(\Delta)$ is obtained as follows. A choice of point $z \in \mathcal{M}$ determines an affine flag F_p at every marked point $p \in \mathbb{M}$. Given a triangle $\delta \in \Delta$, the cluster variable $x_\delta(a,b,c)$ indexed by vertex $(a,b,c) \in Q_k(\delta)$ evaluates on z as follows. First, we parallel transport the tensors $F_{p_1,(a)}$, $F_{p_2,(b)}$, and $F_{p_3,(c)}$ to a common point p' in the interior of δ, and then we set $x_\delta(a,b,c)(z) := F_{p_1,(a)} \wedge F_{p_2,(b)} \wedge F_{p_3,(c)} \in \bigwedge^k (V_{p'}) \cong \mathbb{C}$.

3.5. Initial seeds for $\mathcal{A}'_{G,\mathbb{S}}$

We now carry out a similar procedure for the alternative version of the moduli space. As noted in Remark 3.4, when $\mathbb{S}$ is an n-gon, the moduli space is a Grassmannian and one should use the well-known initial seeds coming from *plabic graphs*. Excluding this case, we can assume that $\mathbb{S}$ admits a taut triangulation Δ.

Given such a taut Δ, we construct an initial seed $\Sigma'_k(\Delta) = (Q'_k(\Delta), \mathbf{x}'_k(\Delta))$. We subsequently check that seeds from different taut triangulations are related by mutations, so that we obtain in this way a cluster algebra $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S})$.

Recall that a choice of point $z \in \mathcal{A}'_{G,\mathbb{S}}$ provides us with an affine flag at each puncture and a vector at each boundary point. We associate a "proxy affine flag" F_p to boundary points p by the following construction. Let C be a component of $\partial\mathbb{S}$ carrying boundary points $p_1, \ldots, p_m$ in counterclockwise order. A choice of z yields vectors $v_i := v_{p_i}$. We define v_i for $i \in \mathbb{Z}$ using the monodromy matrix M_C around the component C: $v_{i+jm} := M_C^j(v_i)$ where $i \in [m]$ and $j \in \mathbb{Z}$. Then the proxy affine flag F_{p_i} which we associate to the boundary point p_i has steps $F_{p_i,(a)} := v_i \wedge v_{i+1} \wedge \cdots v_{i+a-1}$ for $a \in [k-1]$. (This recipe only works for sufficiently generic v_i's, but that is all we need to specify our initial cluster variables.)

Since Δ is a taut triangulation, each triangle δ has at most one side which is a boundary interval. If δ has no such sides, we assign to it the same quiver fragment

$Q_k(\delta)$ and initial cluster variables $x_\delta(a,b,c)$ as in the previous section, using the proxy flags F_p at boundary points p. If δ has a boundary side, we assign to it the quiver fragment Q_k^1 (cf. Section 3.3) with the bottom row of Q_k^1 identified with the boundary side and with the bottom left corner of Q_k^1 following the bottom right corner of Q_k^1 in the counterclockwise order around $\partial\mathbf{S}$. The cluster variables on the right and left sides are of the form $x_\delta(a,k-a,0)$ and $x_\delta(a,0,k-a)$. The vertex on the bottom side is frozen, and it carries the cluster variable $\det(v_i,\dots,v_{i+k-1})$ in the notation of the preceding paragraph. There are $|\mathbb{M}_\partial|$ many frozen vertices in total. The quiver $Q'_k(\Delta)$ is obtained by gluing the quiver fragments associated to triangles along their shared edges.

This completes our description of initial seeds. To complete our definition of the corresponding cluster algebra, our next three lemmas will establish that seeds from different taut triangulations are related by mutations.

We denote by $d(\Delta,\Delta')$ the distance between triangulations in the flip graph, i.e. the minimal number of flips needed to transform Δ into Δ'.

We let $b(\Delta)$ denote the number of boundary arches present in the triangulation Δ. We write $\Delta\sim\Delta'$ to indicate that Δ' can be obtained from Δ by a flip and use the notation $\gamma\mapsto\gamma'$ to indicate that a given flip move replaces the arc $\gamma\in\Delta$ with the arc $\gamma'\in\Delta'$. We classify the flip $\Delta\sim\Delta'$ as either a 0-move, a 1-move, or a -1-move according to the value of $b(\Delta)-b(\Delta')$.

Clearly, if $d(\Delta,\Delta')=1$ and both Δ and Δ' are taut, then the flip between them is a 0-move.

Lemma 3.6. *Let Δ and Δ' be taut triangulations satisfying $d(\Delta,\Delta')\geq 2$. Then any minimal-length sequence of flips connecting Δ to Δ' has at least two 0-moves.*

Proof. We will assume that there are either no 0-moves in the flip sequence, or that there is exactly one 0-move in the flip sequence, and eventually contradict the length-minimality of the flip sequence.

When we perform a 1-move, we flip a spanning arc γ to obtain a boundary arch γ'. Any boundary arches which γ' nests will remain boundary arches until we perform another flip at γ'. That is, boundary arches are removed in a "last in first out" order.

Let us focus on the *last* 1-move performed in our minimal flip sequence. The local picture before such move is drawn in Fig. 1. Vertices $L,M,R\subset\mathbb{M}$ reside on some boundary component C while $p\in\mathbb{M}\setminus C$. We write $\gamma=pM$ as a shorthand to say that γ has endpoints p and M. The below calculations take place in a polygon so we lose nothing by labeling arcs by endpoints in this way. We have $\gamma'=LR$, have arcs $\beta_1=LM$ and $\beta_2=MR$ nested by γ' (either boundary arches or boundary intervals, depending on whether the vertices L, M, and R are adjacent vertices along C), and have spanning arcs $\alpha_1=pL$ and $\alpha_2=pR$ forming part of the quadrilateral containing γ.

If there are no 0-moves in our flip sequence, then after the flip $\gamma\mapsto\gamma'$, we do not flip at either of the arcs α_i because these would not be -1-moves. And by the "last in first out" property, we do not flip at either of the arcs β_i until we have first performed the inverse

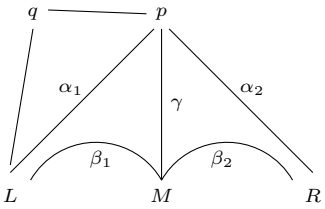

Fig. 1. The local picture before the last 1-move in the proof of Lemma 3.6. Here L, M, R, q, p indicate marked points whereas α_i, β_i, and γ indicate arcs. The flip at γ is a 1-move.

flip $\gamma' \mapsto \gamma$. At some point we *must* perform the -1-move at γ', but after appropriately commuting terms in the flip sequence, we see that the given flip sequence contains the subsequence $\gamma \mapsto \gamma' \mapsto \gamma$ contradicting the length-minimality of the sequence.

If there *is* a 0-move in the flip sequence, then the preceding argument shows that this 0-move comes *after* the last 1-move, and moreover that this 0-move must be a flip at one of the α_i's or one of the β_i's.

In each of these four cases, one can use the flip identity in a pentagon (together with commutation moves) to obtain a shorter sequence of flips, contradicting the supposed minimality of the sequence. For example, consider the case that the 0-move is performed at α_1 to obtain a new arc $\alpha_1' = qR$. Since this is a 0-move, we have $q \notin C$. In the given flip sequence, we see the flips $\gamma \mapsto \gamma'$, followed by the flip $\alpha_1 \mapsto \alpha_1'$ followed by $\gamma' \mapsto \gamma'' = qM$. But up to commutation moves, we can replace these three flips by the shorter sequence $\alpha_1 \mapsto \gamma''$ followed by $\gamma \mapsto \alpha'$ contradicting minimality. (Note that the flip at qL is forbidden because it is not a -1-move. Flips at qp are possible in the case that q and p are on the same boundary component, but the resulting arc will never be flipped again in the sequence. The contradiction from the preceding sentences still works in this case.) The case that we flip at a β_i is similar. □

Lemma 3.7. *The set of taut triangulations of* $\mathbb{S}$ *is connected by flips.*

By a *taut* flip sequence we will mean a set of sequence of flips which never leaves the set of taut triangulations.

Proof. Observe that any triangulation Δ is exactly $b(\Delta)$ many flips away from a taut triangulation in the flip graph: one can reach such a triangulation by performing -1-moves at boundary arches in "last in first out" order.

We now prove by induction on $d(\Delta, \Delta')$ that any two taut triangulations Δ, Δ' are connected by a taut flip sequence.

Let $\underline{\Delta} := \Delta = \Delta_0 \sim \Delta_1 \sim \cdots \sim \Delta_d = \Delta'$ be a shortest walk from Δ to Δ' in the flip graph. Abbreviate $b_i := b(\Delta_i)$.

Define $i \geq 1$ by the property that the first $i-1$ many moves in the flip sequence are 1-moves and the ith move is either a 0-move or a -1-move. Thus, $b_{i-1} = i-1$ and $b_i < i$.

Altogether, the flip sequence has at least $i-1$ many 1-moves, thus at least $i-1$ many -1-moves, and also at two many 0-moves by Lemma 3.6. Thus, $2i \leq d$.

By the observation in the first paragraph, we can choose a taut triangulation Δ'' whose distance from Δ_i in the flip graph equals b_i. Then

$$d(\Delta'', \Delta_d) \leq d(\Delta'', \Delta_i) + d(\Delta_i, \Delta_d) = b_i + (d-i) < d$$

and similarly

$$d(\Delta_0, \Delta'') \leq d(\Delta_0, \Delta_i) + d(\Delta_i, \Delta'') = i + b_i < 2i \leq d.$$

By induction, Δ'' is connected to both Δ_d and Δ_0 by a taut flip sequence, so that Δ_d and Δ_0 are connected to each other by a taut flip sequence as required. □

Lemma 3.8. *If Δ and Δ' are taut triangulations related by a flip, then $\Sigma'_k(\Delta)$ and $\Sigma'_k(\Delta')$ are mutation-equivalent seeds.*

Proof. If neither of the two triangles involved in a flip has a boundary side, then the statement is a specialization of Fock and Goncharov's [6, Proposition 10.1].

We next consider the case that one of the two triangles involved in the flip has a boundary side. We indicate the mutation sequence which implements the flip in the case $G = \mathrm{SL}_4$ in the below picture. The triangle pb_1b_2 has an exposed side b_1b_2. The quiver pqb_1 has no exposed sides, and we wish to perform the flip at the arc pb_1 obtaining the non-exposed triangle pqb_2 and the exposed triangle qb_1b_2. The mutation sequence which simulates this flip of triangulations is to mutate at the vertices labeled 1,2,3,4,5,6 (in that order) in the left picture below. The result is the right picture, which is the quiver we expect after performing the flip $pb_1 \mapsto qb_2$.

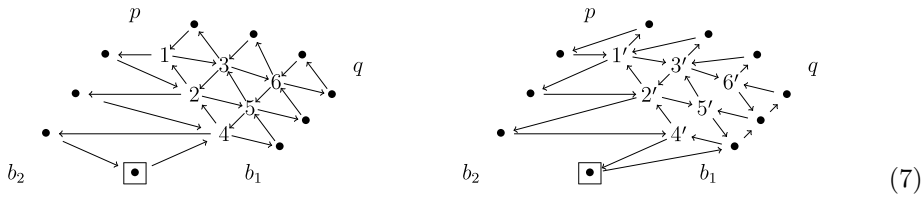

(7)

We have used primes to denote the result of performing the mutation, so that e.g. $1'$ is the result of mutating at vertex 1. The asserted mutation sequence readily extends to the SL_k case.

One should check that the cluster variables transform as expected. To see this, observe that each vertex in the mutation sequence is 4-valent at the moment it is mutated. The corresponding exchange relation is an instance of the three-term Plücker relations. Suppose the proxy affine flag at b_1 is represented by vectors $v_1, \ldots, v_k$, the proxy affine flag at b_2 is represented by $v_2, \ldots, v_{k+1}$, the affine flag at p is represented by $u_1, \ldots, u_k$,

and the affine flag at q is represented by $w_1, \ldots, w_k$. (See Remark 3.2 for the notion of vector representatives of an affine flag.) Let $(\alpha, \beta, \gamma, \eta)$ denote the function

$$\det(u_1, \ldots, u_\alpha, w_1, \ldots, w_\beta, v_1, \ldots, v_\gamma, v_2, \ldots, v_{1+\eta}),$$

then the mutation sequence sends the variable $(\alpha, \beta, \gamma, 0)$ to the variable $(\alpha-1, \beta+1, 0, \gamma)$ via the Plücker relation

$$\begin{aligned}&(\alpha, \beta, \gamma, 0) \times (\alpha - 1, \beta + 1, 0, \gamma)\\ &\quad = (\alpha, \beta, 0, \gamma) \times (\alpha - 1, \beta + 1, \gamma, 0) + (\alpha, \beta + 1, 0, \gamma - 1) \times (\alpha - 1, \beta, \gamma + 1, 0).\end{aligned}$$

For example, in the above example, mutation at vertex 1 (whose variable is $\det(u_1u_2u_3v_1)$) yields the new variable given by $\det(u_1u_2w_1v_2)$. The right hand side of the Plücker relation is

$$\det(u_1u_2v_1v_2)\det(u_1u_2u_3w_1) + \det(u_1u_2w_1v_1)\det(u_1u_2u_3v_2).$$

Note that $1'$ is a mutable variable in the seed fragment $\Sigma_r(\delta)$ where δ is the triangle pqb_2. This completes our discussion of the case in which we have one boundary side.

Finally we consider the case that both triangles have a boundary side. Since the flip results in a taut triangulation, these two sides must be opposite each other in the quadilateral in which the flip is performed. There are $k-1$ vertices of $Q_k(\delta)$ on the edge which needs to be flipped. These vertices have no arrows between them and each vertex is four-valent in the quiver. Mutating at each of them (in any order, since these mutations commute) realizes the flip. As above, each exchange relation is a three-term Plücker relation, and the variable that results is the one which we would expect to see when we perform the flip. □

Remark 3.9. We do not prove here that our initial clusters provide rational coordinate systems on $\mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$ although we believe this should be true. One might be able to prove these statements by mimicking the proofs of these statements given in [6], or by establishing the cluster fibration property from Remark 3.5.

We content ourselves here by showing that the size of our initial cluster matches the dimension of the moduli space. Indeed, the difference in the dimensions of the two moduli spaces is $|\mathbb{M}_\partial|(\dim G/U - \dim V) = |\mathbb{M}_\partial|(\binom{k}{2} - 1)$ using (4) and (5). By inspection, the quiver Q_k^1 has $\binom{k}{2} - 1$ fewer vertices than Q_k. The cardinality $|\mathbb{M}_\partial|$ is the number of boundary intervals, which is the number of times Q_k^1 rather than Q_k is used when creating the quiver $Q'_k(\Delta)$. The claim follows.

Example 3.10. When $\mathbb{S}$ is an annulus with one point on each boundary component, the cluster algebra $\mathscr{A}'(\mathbb{S}, \mathrm{SL}_k)$ has the same cluster type as the Q-system of type A_{k-1} see e.g. [22].

Recall that $D_{n,1}$ has a unique taut triangulation, hence $\mathscr{A}'(\mathrm{SL}_k, D_{n,1})$ has a canonical choice of initial seed. The mutable part of this seed is the cylindrical triangulated grid quiver $Q_{m,n}$ considered in [20]. Thus Lemma 4.5 below provides a different proof of Theorem 3.2 therein.

Remark 3.11. We believe that it is possible to assign a seed in the cluster algebra $\mathscr{A}'_{G,\mathbb{S}}$ to any regular triangulation Δ, not only to the taut ones. We explain this recipe now without proof.

Suppose that $\delta \in \Delta$ has exactly one side γ which is a boundary arch contractible to s many consecutive boundary intervals. Then one should associate to this triangle the quiver $Q_k^{\min(s,k-1)}$. If $s \leq k-1$, then the rightmost vertex on the bottom side is frozen.

If δ has two such sides, one should apply this recipe to *both* of the exposed sides (deleting and gluing vertices as in the definition of $Q_k^s(\delta)$). See [7, Figure 21] for an illustration of this recipe when $G = \mathrm{SL}_3$.

3.6. Weights of cluster variables

Let $\mathcal{M}$ be one of the moduli spaces either $\mathcal{A}_{G,\mathbb{S}}$ or $\mathcal{A}'_{G,\mathbb{S}}$. Recall the right T-action (3) on affine flags by rescaling the $F_{(i)}$'s. It determines a right action $\mathcal{M} \curvearrowleft T^{\mathbb{M}_\circ}$ by separately rescaling the decorations at each puncture.

We let $P^{\mathbb{M}_\circ}$ denote the direct sum of weight lattices indexed by the punctures in $\mathbb{S}$, and denote by

$$\pi_p \colon P^{\mathbb{M}_\circ} \to P$$

the projection onto the copy of the weight lattice indexed by a given puncture $p \in \mathbb{M}_\circ$.

A rational function $f \in \mathbb{C}(\mathcal{M})$ is *homogeneous* if there exists a weight vector $\lambda \in P^{\oplus \mathbb{M}_\circ}$ such that

$$f(z \cdot t) = t^\lambda f(z), \text{ for } z \in \mathcal{M}, t \in T^{\mathbb{M}_\circ}.$$

For such an f, define

$$\mathrm{wt}(f) := \lambda \in P^{\oplus \mathbb{M}_\circ}, \text{ the } \textit{weight} \text{ of } f \qquad \mathrm{wt}_p(f) := \pi_p(\lambda) \in P, \text{ the } \textit{weight of } f \textit{ at } p.$$

Each of the initial cluster variables in $\mathscr{A}(\mathcal{M})$ is a homogeneous function with nonzero weight at three or fewer punctures. For example, let $\delta \in \Delta$ be a triangle with three distinct vertices and with "top vertex" p_1 a puncture rather than a boundary point. Then the initial cluster variable $x_\delta(a,b,c)$ in this triangle has weight ω_a at p_1. One can see that the two cluster monomials on the right hand side of its exchange are homogeneous functions of the same weight, i.e. they satisfy the balancing condition (2) necessary to determine an initial P-seed. The frozen variables have weight zero, so one can delete them with no effect on the P-cluster variables.

Henceforth when we refer to P-clusters, P-cluster variables, etc., we always have in mind that we have fixed a choice of moduli space $\mathcal{M}$ and are discussing those P-clusters which are reachable from our initial P-seeds by mutations.

For a P-cluster $\mathcal{C} \subset P^{\mathbb{M}_\circ}$ and puncture p, the projection $\pi_p(\mathcal{C}) \subset P$ is the *P-cluster at p*.

Example 3.12. When $k = 2$, a plain (resp. notched) end of a tagged arc ending at a puncture p contributes a copy of the standard basis vector e_1 (resp. e_2) to the weight of the corresponding cluster variable at p. Thus, the weight of any cluster variable at p lies in $\{0, e_1, e_2, 2e_1, 2e_2\}$. The P-cluster at p is either a multiset drawn from $\{0, e_1, 2e_1\}$ (all arcs plain at p), a multiset drawn from $\{0, e_2, 2e_2\}$ (all arcs notched at p), or it has exactly one P-cluster variable of weight e_1, one of weight e_2, and all others of weight zero at p.

4. Weyl group action at punctures

For any puncture p, Goncharov and Shen identified a birational action of the Weyl group of G on the space of decorated G-local systems by changing the affine flag at p. They proved, in most cases, that the resulting Weyl group action on $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$ is by cluster automorphisms. These Weyl group symmetries play an important role in the cluster combinatorics of $\mathscr{A}(\mathcal{M})$. We review the construction and show that the action is cluster in some new cases.

4.1. Weyl group action on affine flags

Definition 4.1. Let u be a unipotent matrix and F be an affine flag stabilized by u. Choose a vector v with the property that $F_{(i+1)} = F_i \wedge v$. Define an affine flag $s_i \cdot F$ whose ith step is given by the formula $(s_i \cdot F)_{(i)} := F_{(i-1)} \wedge (u(v) - v)$ and whose other steps agree with F.

This recipe only yields an affine flag when $u(v) - v \notin \operatorname{span} F_{(i-1)}$, in which case it is independent of the chosen v. The resulting affine flag is again stabilized by u. The operations s_i are involutions satisfying the relations (1), thus they determine a (rational) W-action on the set of affine flags stabilized by u.

Remark 4.2. Representing F by vectors $v_1, \ldots, v_k$ satisfying $F_{(a)} = v_1 \wedge \cdots \wedge v_a$, the flag $s_i \cdot F$ can be represented by vectors

$$v_1, \ldots, v_{i-1}, \mathcal{W}_i v_i, \frac{1}{\mathcal{W}_i} v_{i+1}, v_{i+2}, \ldots, v_k \text{ where } \mathcal{W}_i := \frac{F_{(i-1)} \wedge (u(v_{i+1}) - v_{i+1})}{F_{(i)}} \in \mathbb{C}. \tag{8}$$

Thus, applying s_i changes the vector representatives in two positions i and $i+1$ even though it only changes the affine flag in step i. Expressing the linear transformation u

in the ordered basis $v_1, \ldots, v_k$, the quantity $\mathcal{W}_i$ is the subdiagonal matrix entry in row i and column $i+1$. Fixing $v_1, \ldots, v_k$ and letting u vary, the quantity $\mathcal{W}_i$ is logarithmic in u: multiplication of matrices u corresponds to addition of the quantity $\mathcal{W}_i$.

The following lemma follows from the definition and is used later in the paper.

Lemma 4.3. *Let F be an affine flag with a choice of $u \in \mathrm{stab}(F)$. Then the tensor $(w \cdot F)_{(a)} \in \bigwedge^a(V)$ obtained by Weyl group action along w depends only on the first a symbols of w, i.e. on $w|_{[a]}$.*

Now, following Goncharov and Shen, we relate the Weyl group action on affine flags with the cluster algebra $\mathscr{A}(\mathcal{M})$ for one of the moduli spaces $\mathcal{M}$ considered in this paper. We denote by $\mathrm{Bir}(\mathcal{M})$ the group of birational automorphisms of the moduli space $\mathcal{M}$.

Fix $p \in \mathbb{M}_\circ$. A choice of point in $\mathcal{M}$ determines an affine flag F_p at p, stabilized by the monodromy u_p of the local system around p. Applying Definition 4.1 when u is taken to be this u_p, we have a map

$$\psi_{i,p} \in \mathrm{Bir}(\mathcal{M}) \text{ sending } F_p \mapsto s_i \cdot F_p, \tag{9}$$

changing the decoration at the puncture p while changing neither the underlying local system nor the decorations at $\mathbb{M} \setminus \{p\}$.

The actions at various punctures commute, yielding a group homomorphism $\psi\colon W^{\mathbb{M}_\circ} \to \mathrm{Bir}(\mathcal{M})$. This action is compatible with the W-action on the weight lattice: for homogeneous f,

$$\mathrm{wt}(f \circ \psi(\underline{w})) = \underline{w} \cdot \mathrm{wt}(f) \in P^{\oplus \mathbb{M}_\circ}, \text{ for } \underline{w} \in W^{\mathbb{M}_\circ}. \tag{10}$$

In particular the homomorphism ψ is injective.

Theorem 4.4 *([14, Theorem 1.2]). Let $\mathcal{M} = \mathcal{A}_{\mathrm{SL}_k,\mathbb{S}}$ where $\mathbb{S}$ is neither an $S_{g,1}$ (for any k) when $k = 2$. Then $\psi(W^{\mathbb{M}_\circ}) \subset \mathrm{Bir}(\mathcal{M})$ acts by cluster automorphisms of the cluster algebra $\mathscr{A}(\mathcal{M})$.*

For an extension of this theorem to other Lie types see [19].

In the remainder of this section we extend Theorem 4.4 to two new settings, namely to $\mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$ and also the cases $(k, S_{g,1})$ when $k > 2$. The statement is false for $S_{g,1}$ and for $S_{0,3}$ when $k = 2$.

The Weyl group action at punctures is a crucial ingredient to the cluster combinatorics of $\mathscr{A}(\mathcal{M})$, so we felt it was important to treat these cases here. Our proof is a modification of the ideas already present in Goncharov and Shen's proof of the above theorem. The reader who is not interested in these two special cases can safely skip to Section 5.

4.2. Reminders on the proof of Theorem 4.4

We fix a puncture $p \in \mathbb{M}_\circ$ and a regular triangulation Δ of $\mathbb{S}$ with the property that no arc in Δ is a loop based at p. (The existence of such triangulations when $\mathbb{S} \neq S_{g,1}$ is asserted in [14].) In the special case that $\mathbb{S} = S_{g,1}$, *every* arc is a loop based at the puncture, which is why this case is excluded from Theorem 4.4.

For an initial cluster variable $x \in \mathbf{x}_k(\Delta)$, we will compute $\psi^*_{i,p}(x)$ as a Laurent polynomial in the initial cluster variables $\mathbf{x}_k(\Delta)$. Then we show that we can reach this set of Laurent polynomials $\{\psi^*_{i,p}(x) \colon x \in \mathbf{x}_k(\Delta)\}$ by mutations, proving in this way that $\psi^*_{i,p}$ sends clusters to clusters.

Consider a triangle $\delta \in \Delta$ with p as one of its vertices. Identify p with the "top vertex" $(k, 0, 0) \in \hat{\mathbb{H}}_k(\delta)$ and recall that the initial cluster variable $x_\delta(a, b, c)$ has $\mathrm{wt}_p(x_\delta(a, b, c)) = \omega_a \in P$. From this and the definition (9), we get a formula for the Weyl group action:

$$\psi^*_{i,p}(x_\delta(a, b, c)) = x_\delta(a, b, c) \text{ if } a \neq i, \text{ and} \qquad \psi^*_{i,p}(x_\delta(i, b, c)) = \mathcal{W}_i x_\delta(i, b, c), \tag{11}$$

where $\mathcal{W}_i$ is the quantity (8) in the case that $u = u_p$ is the monodromy around p. That is, $\psi^*_{i,p}$ only affects the variables in row i around p and it rescales these variables by the factor $\mathcal{W}_i$.

An edge $e = (i, b, c) - (i, b + 1, c - 1)$ in a triangle δ with p as one of its vertices determines a *rhombus monomial* $\mathrm{RM}(e) = \frac{x_\delta(i+1,b,c)x_\delta(i-1,b+1,c)}{x_\delta(i,b,c)x_\delta(i,b+1,c)}$. Goncharov and Shen show that the quantity $\mathcal{W}_i$ is a Laurent polynomial in initial cluster variables:

$$\mathcal{W}_i = \sum_{e \text{ in row } i \text{ of } \delta,\ \delta \text{ incident to } p} \mathrm{RM}(e). \tag{12}$$

Combining formulas (11) and (12) gives us a Laurent polynomial formula for the action of $\psi^*_{i,p}$.

It remains to show that this formula amounts to a sequence of mutations. As a preliminary step, consider a quiver which is an oriented m-cycle on the vertex set $[m]$ with arrows $j \to j + 1 \mod m$. Add to this quiver $2m$ frozen vertices $j^\pm$ as well as arrows $j + 1 \to j^\pm \to j$. We denote the mutable initial cluster variable at vertex j by x_j and the frozen variables at $j^\pm$ by $x_j^\pm$. Note that the mutable edge $j \to j + 1$ is part of two three-cycles $j \to j+1 \to j^+ \to j$ and $j \to j+1 \to j^- \to j$ together forming a "rhombus."

Consider the permutation-mutation sequence

$$\mu_1 \circ \mu_2 \circ \cdots \mu_{m-1} \circ (m - 1, m) \circ \mu_{m-1} \circ \mu_{m-2} \circ \cdots \mu_1 \tag{13}$$

where the middle term is a transposition written in cycle notation and μ_i denotes mutation at vertex i. It is known that this permutation-mutation sequence determines an involutive automorphism of the cluster algebra which is independent of the choice of

which vertex is called 1. The automorphism rescales initial mutable variables by the Laurent polynomial

$$\sum_{j\in[m]} \frac{x_j^+ x_j^-}{x_j x_{j+1}}, \tag{14}$$

cf. [20, Theorem 3.1] or [14, Theorem 7.7].

Returning to the setting of Theorem 4.4, note that when we restrict the quiver $Q_k(\Delta)$ to the vertices which are in rows i in the triangles surrounding p, we get an oriented m-cycle for some m. The subquiver in rows $i-1, i, i+1$ is a specialization of the quiver from (14): the frozen variables $x_j^{\pm}$ from (14) should be specialized to the variables $x_\delta(i+1, b, c)$ and $x_\delta(i-1, b+1, c)$ from (12). After such specialization, the Laurent polynomials (14) and (12) coincide, completing the proof.

Finally, for the purposes of our generalization below, let us recall how the formula (12) is proved. By a well-known recipe a generic pair (F, G), with F an affine flag and G a complete flag, determines an ordered basis in V. (Intersecting span F_i with span G_{k-i+1} determines an ordered basis up to rescaling each of the basis vectors; the scalars are fixed using the extra information provided by the affine flag F.)

Let δ be a triangle with a puncture p as its top vertex and with p_2 and p_3 its bottom left and bottom right vertices. A choice of point in the moduli space $\mathcal{M}$ determines affine flags F_p, F_{p_2}, and F_{p_3}. The preceding paragraph gives ordered bases associated to the pairs (F_p, F_{p_2}) and also to (F_p, F_{p_3}), i.e. to the left and right sides of δ. There is a unique matrix $u(\delta)$ which sends the left ordered basis to the right ordered basis, and this matrix stabilizes F_p. By [15, Section 3.1], the quantity $\mathcal{W}_i$ (8) when calculated on the matrix $u(\delta)$ is the Laurent polynomial $\sum_{e \text{ in row } i \text{ of } \delta} \mathrm{RM}(e)$. If $\delta_1, \ldots, \delta_s$ are the triangles surrounding p in cyclic order, it follows that $u_p = \prod_{i=1}^{s} u(\delta_i)$. The formula (12) follows because the quantity $\mathcal{W}_i$ is logarithmic in the matrix u.

4.3. The action on $\mathcal{A}'_{G,\mathbb{S}}$ is cluster

Now we give the analogue of Theorem 4.4 in the case that boundary points carry vector decorations.

Lemma 4.5. *The $W^{\mathbb{M}_\circ}$-action on $\mathscr{A}'(\mathrm{SL}_k, \mathbb{S})$ is by cluster automorphisms.*

Proof. We can assume that $\mathbb{S}$ has boundary because otherwise we are dealing with the Fock-Goncharov version of the moduli space. Fix a puncture p and consider a taut triangulation Δ which has the property that no arc is a loop based at p. (To construct such, start with any regular triangulation with no loops and mutate at boundary arches in last in first out fashion as in the proof of Lemma 3.6.) It still makes sense to consider vertices which are in row i around p. These will reside in $Q_k(\delta)$ or in $Q_k^1(\delta)$ according to

whether or not δ has a boundary side and these again have weight ω_i at p so that the formula (11) holds.

We can mimic the argument from [15, Section 3.1] factoring the matrix u_p into matrices $u(\delta)$ indexed by triangles δ surrounding p and get the corresponding formula (12). When δ has a boundary side, there is only one rhombus which contributes to the sum (12) because there are fewer steps required to move the ordered basis associated to the left side of δ to that for the right side of δ.

The induced subquiver on the vertices in row i around p is again an oriented m-cycle for some m, and each edge in this cycle is part of a "rhombus." The formula (14) again specializes to the rhombus Laurent monomial (12), proving the assertion. □

4.4. The action on $\mathscr{A}(\mathrm{SL}_k, S_{g,1})$ is cluster

When $k = 2$, the Weyl group action at punctures is the tag-changing transformation and is not cluster in the special case that $\mathbb{S} = S_{g,1}$ for some $g \geq 1$. See Example 8.17 for a detailed comparison between the s_1 action at punctures and plain versus notched tagging of arcs.

Goncharov and Shen do not discuss whether or not the W-action on $\mathscr{A}(\mathrm{SL}_k, S_{g,1})$ is cluster when $k > 2$. We resolve this affirmatively:

Proposition 4.6. *When $k > 2$, the W-action on $\mathscr{A}(\mathrm{SL}_k, S_{g,1})$ is by cluster automorphisms.*

As an ingredient of our proof, we recall that the cluster algebra $\mathscr{A}_{\mathrm{SL}_k,\mathbb{S}}$ admits a cluster automorphism $*$, the *duality map*, which implements the outer automorphism of $i \mapsto k - i$ of the type A_{k-1} Dynkin diagram [14, Section 9]. In particular, we have

$$* \circ \psi_{i,p} \circ * = \psi_{k-i,p}, \tag{15}$$

where $\psi_{i,p} \in \mathrm{Bir}(\mathcal{M})$ denotes Weyl group action at p.

In terms of points in the moduli space $\mathcal{M}$, the automorphism $*$ amounts to replacing the vector space V and its volume form with their duals. Importantly for us, Goncharov and Shen show that $*$ is a cluster automorphism even for $\mathcal{M} = \mathcal{A}_{\mathrm{SL}_k,S_{g,1}}$.

Proof. We denote by p the puncture and abbreviate $\psi_{i,p} = \psi_i$. We will explain that ψ_i is a cluster automorphism when $i \geq \frac{k}{2}$. This implies the statement for all $i \in [k-1]$ since $*$ is a cluster automorphism, via (15).

Consider a regular triangulation Δ of $S_{g,1}$ (cf. Fig. 2 for the case $g = 1$). All three vertices of any triangle $\delta \in \Delta$ are the puncture p. The cluster variable $x(a, b, c)$ having coordinates $(a, b, c) \in \mathbb{H}_k$ in this triangle has weight $\omega_a + \omega_b + \omega_c$ at p, regardless of how we identify the three corners of δ with the vertices $(k, 0, 0)$, $(0, k, 0)$ and $(0, 0, k)$ in $\hat{\mathbb{H}}_k$.

It still makes sense to consider the vertices in row i, except that these vertices will now be a union of three line segments parallel to the three sides of the triangle δ. Each of the

three segments is contractible to one of the three corners of δ. The union of these three segments, taken over all of the various triangles δ, is therefore a closed curve contractible to p. If we temporarily view one of the three corners as the top vertex of a triangle, it makes sense to move the "left side" of δ to the "right side" as in the proof of Theorem 4.4. Doing this for each of the three corners, we again factorize the monodromy u_p around p into monodromies associated to each of the corners of the various triangles δ. Thus the rhombus monomial formula (12) still holds.

When $i > \frac{k}{2}$ the three segments are disjoint and the aforementioned closed curve has no selfintersections. For an example see the blue cycle in Fig. 2. In this case, the cluster variable $x(i,b,c)$ has weight $\omega_i + \omega_b + \omega_c$ with $\omega_b, \omega_c \neq \omega_i$, hence (11) holds. The restriction of the quiver $Q_k(\Delta)$ to the vertices which are in row i (with respect to any corner of any triangle) is an oriented m-cycle for some m. We can apply the mutation sequence (13) and the formula (14) still holds. This completes the proof when $i > \frac{k}{2}$.

What remains is to argue the case $i = \frac{k}{2}$ when k is even and $k > 2$. In this case the three segments in row i intersect each other at the points $(i,i,0)$, $(i,0,i)$ and $(0,i,i)$ in each triangle. See the red segments in Fig. 2. We call the corresponding vertices of $Q_k(\Delta)$ the *midpoint vertices* and call the remaining vertices in row i the *nonmidpoint vertices*. The midpoint vertices have weight $2\omega_i$, so that the correct version of (11) rescales the cluster variable $x(i,i,0)$ by $\mathcal{W}_i^2$. On the other hand, (11) holds for the nonmidpoint vertices.

Provided $k > 2$, there are no edges between the various midpoint vertices. The mutation μ_{midpoint} at all the midpoint vertices (in any order) is well-defined. We use primes $x'(i,i,0)$ to denote the result of such mutation. Note that each midpoint vertex has degree 4 with 2 incoming and two outgoing neighbors. Each neighbor has weight $\omega_i + \omega_{i-1} + \omega_1$, so that $\text{wt}_p(x'(i,i,0)) = 2(\omega_1 + \omega_{i-1})$. In particular, $\psi_{i,p}^*$ fixes $x'(i,i,0)$.

The induced subquiver of $\mu_{\text{midpoint}}(Q_k(\Delta))$ on the nonmidpoint vertices is an oriented m-cycle for some m. We claim that performing the sequence (13) to the seed $\mu_{\text{midpoint}}(\Sigma_k(\Delta))$ is the Weyl group automorphism ψ_i^*.

The action of $\psi_{i,p}^*$ on $\mu_{\text{midpoint}}(\Sigma_k(\Delta))$ is to rescale each nonmidpoint cluster variable by the quantity $\mathcal{W}_i$ (fixing all other cluster variables). The expression (12) for $\mathcal{W}_i$ as a Laurent polynomial in the initial cluster variables is still valid. What needs to be checked is that when we rewrite the quantity (12) in terms of the variables in $\mu_{\text{midpoint}}(\Sigma_k(\Delta))$, it matches the formula (14) for the mutation sequence (13) applied to the cycle on the nonmidpoint variables. This is a straightforward check which we illustrate in Example 4.7. □

Example 4.7. We illustrate the ideas in the proof of Proposition 4.6 in the case that $G = \text{SL}_4$ and $\mathbb{S} = S_{1,1}$. The initial quiver is in Fig. 2 with rows 3 and 2 drawn in blue and red. The induced subquiver on the red vertices is not an oriented cycle. To rectify this, one should mutate at vertices 2, 4 and 10 to obtain the quiver on the right. In this quiver, the red form an oriented 6-cycle. The formula (12) takes the form

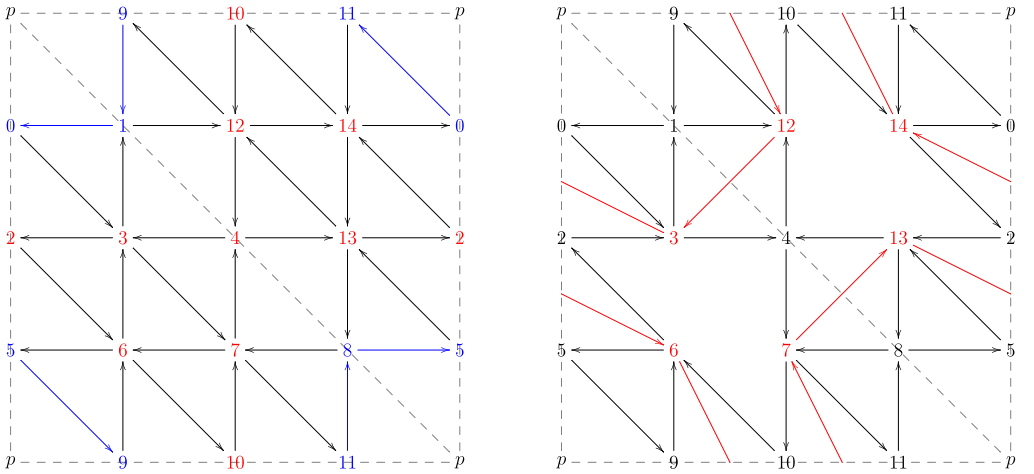

Fig. 2. From the proof of Proposition 4.6. The gray dashed lines describe a triangulation of the once-punctured torus $S_{1,1}$ (with top/bottom and left/right sides identified). The left picture is the initial quiver when $G = \mathrm{SL}_4$. The blue vertices comprise row 3; the induced subquiver on these vertices is an oriented 6-cycle contractible to the puncture p. The red vertices comprise row 2. Mutation at the midpoint vertices 2, 4, and 10 yields the quiver on the right. The nonmidpoint vertices (drawn red in the right figure) form an oriented 6-cycle contractible to the puncture. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

$$\begin{aligned}
&\frac{x_0x_6}{x_2x_3}+\frac{x_1x_7}{x_3x_4}+\frac{x_1x_{13}}{x_4x_{12}}+\frac{x_9x_{14}}{x_{10}x_{12}}+\frac{x_7x_9}{x_6x_{10}}+\frac{x_3x_5}{x_2x_6}+\frac{x_5x_{14}}{x_2x_{13}}+\frac{x_8x_{12}}{x_4x_{13}}\\
&\quad+\frac{x_3x_8}{x_4x_7}+\frac{x_6x_{11}}{x_7x_{10}}+\frac{x_{11}x_{12}}{x_{10}x_{14}}+\frac{x_0x_{13}}{x_2x_{14}}\\
&=\frac{x_4'x_1}{x_3x_{12}}+\frac{x_2'x_0}{x_3x_{14}}+\frac{x_{10}'x_{11}}{x_7x_{14}}+\frac{x_4'x_8}{x_7x_{13}}+\frac{x_2'x_5}{x_6x_{13}}+\frac{x_{10}'x_9}{x_6x_{12}}.
\end{aligned}$$

In passing from the first to the second line, we e.g. grouped the terms $\frac{x_1}{x_4}(\frac{x_7}{x_3}+\frac{x_{13}}{x_{12}})$ and then used the mutation formula $x_4' = \frac{x_7x_{12}+x_3x_{13}}{x_4}$. The result is the formula (14) applied to the red oriented 6-cycle on the right of Fig. 2 with the frozen variables appropriately specialized.

Remark 4.8. In [14], Goncharov and Shen give a formula for the DT-transformation of $\mathscr{A}(\mathrm{SL}_k, \mathbb{S})$. It is a composition of several commuting maps, one of which is the action of the longest element $w_0 \in W$ at each puncture. We expect that this formula works for $S_{g,1}$'s, so that our Theorem 4.4 shows that the DT transformation of $\mathcal{A}_{\mathrm{SL}_k,S_{g,1}}$ is cluster once $k > 2$. We believe this would follow from Section 8.2 in [14] provided one checks the analogue of Theorem 4.4 on the $\mathcal{X}$-space. When $k = 3$, we have checked that the candidate map is indeed the DT transformation and that the standard sequence of mutations which realizes this map is a maximal green sequence.

5. Compatibility of weight vectors

With the above preliminaries established, the next several sections investigate properties of P-clusters in the cluster algebras $\mathscr{A}(\mathcal{M})$. The present section introduces a compatibility notion for elements of the weight lattice which we expect models compatibility of cluster variables. We use freely here the Coxeter-theoretic terminology of roots, weights, ordered set partitions, etc. from Section 2.4.

The maximal faces of the Coxeter complex are indexed by permutations. The closure $\overline{C}_w$ of the facet C_w indexed by the permutation w consists of weight vectors λ whose coordinates satisfy $\lambda_{w(1)} \geq \cdots \geq \lambda_{w(k)}$.

Definition 5.1. A pair of weight vectors $\lambda, \mu \in P$ is *w-sortable* if both these vectors lie in $\overline{C_w}$. The pair is W-sortable if it is w-sortable for some $w \in W$.

The pair is *root-conjugate* if the difference $\lambda - \mu$ lies in the root system Φ and moreover, the vectors λ and μ are conjugate by reflection in this root vector.

The pair is *compatible* if it is either W-sortable or root-conjugate.

A compatible pair of vectors is either W-sortable or root-conjugate, never both.

Whenever weight vectors λ and μ are conjugate by reflection in the root $e_a - e_b$, one has $\lambda - \mu \in \mathrm{span}_{\mathbb{Z}}(e_a - e_b)$. Root-conjugacy requires more strongly that $\lambda - \mu = \pm(e_a - e_b)$. This condition is equivalent to requiring that there exists a weight vector $\nu \in P$ whose coordinates satisfy $\nu_a = \nu_b$, and a choice of sign $\varepsilon \in \{\pm 1\}$, such that $\lambda = \nu + \varepsilon e_a$ and $\mu = \nu + \varepsilon e_b$. In fact, one readily sees that if this property holds for ε, then it also holds for $-\varepsilon$, so the choice of sign is immaterial. This sign ambiguity disappears once we have three vectors which are pairwise root-conjugate, see Lemma 5.3 below.

Example 5.2. The vectors $(7, 3, 2)$ and $(3, 7, 2)$ are s_1-conjugate but not root-conjugate. The vectors $(7, 3, 2)$ and $(6, 4, 2)$ differ by the root $e_1 - e_2$ but they are not root-conjugate. The vectors $(7, 3, 2)$ and $(7, 2, 3)$ are root-conjugate in the direction of the root $e_2 - e_3$. Applying the previous paragraph to this pair, we could take either $\nu = (7, 3, 3)$ and $\varepsilon = -1$ or take $\nu = (7, 2, 2)$ and $\varepsilon = +1$.

If $\mathcal{C}$ is a set of pairwise root-conjugate vectors, then every difference of two vectors in $\mathcal{C}$ takes the form $e_a - e_b$ for appropriate indices $a, b \in [k]$. Define the *ground set* of such a $\mathcal{C}$ as

$$\{a \in [k] : \text{ there exists } \lambda, \mu \in \mathcal{C} \text{ and } b \in [k] \text{ such that } \lambda - \mu = e_a - e_b\}. \tag{16}$$

As an example, the collection $\mathcal{C} = \{(6, 7, 6, 7), (7, 6, 6, 7), (7, 7, 6, 6)\}$ has ground set $\{1, 2, 4\}$.

Lemma 5.3. *Let $\mathcal{C} \subset P$ be a collection of weight vectors with $|\mathcal{C}| > 2$.*

If the elements of $\mathcal{C}$ are pairwise W-sortable, then there is a unique osp Π such that $\mathcal{C} \subset \overline{C_\Pi}$ and $\overline{C_\Pi}$ is minimal by inclusion with this property.

If instead the elements of $\mathcal{C}$ are pairwise root-conjugate, let $B \subset [k]$ be their corresponding ground set (16). Then there exists a unique weight vector ν with constant B-coordinates (i.e., $\nu_a = \nu_b$ when $a, b \in B$), and a unique sign $\varepsilon \in \{\pm 1\}$, such that $\mathcal{C} = \{\nu + \varepsilon e_a : a \in B\}$.

The ordered set partition Π alluded to in the first part of this lemma is the least upper bound of the ordered set partitions $\Pi(\lambda) : \lambda \in \mathcal{C}$. To illustrate the second part of the lemma, we would have uniquely that $\nu = (7, 7, 6, 7)$ and $\varepsilon = -1$ for the collection $\mathcal{C}$ introduced just before the lemma statement.

Proof. The first statement is an instance of the classical statement that the Coxeter complex is a flag complex. Given λ, the corresponding closed face $\overline{C_{\Pi(\lambda)}}$ of the Coxeter complex is a simplex on certain vertices $v_i(\lambda)$. If $\mu \in \mathcal{C}$, then the w-sortability hypothesis says that $\overline{C_{\Pi(\lambda)}} \cup \overline{C_{\Pi(\mu)}} \subset \overline{C_{\Pi(w)}}$. Thus, any pair of $v_i(\lambda)$ and $v_j(\mu)$ form a 2-simplex in the Coxeter complex, hence the union of the $v_i(\lambda)$'s and $v_j(\mu)$'s determines a simplex in the Coxeter complex (because the latter is a flag complex). Continuing in this way, the minimal simplex which contains all of the vectors $\lambda \in \mathcal{C}$ is the one spanned by the various vertices $v_i(\lambda)$. This minimal by inclusion simplex is the desired $\overline{C_\Pi}$.

For the second statement, let λ, λ', and λ'' be three elements of $\mathcal{C}$. Observe that if $a \in [k]$ and λ, μ are root-conjugate, then $\lambda_a - \mu_a \in \{0, \pm 1\}$. Suppose that $\lambda - \lambda' = e_a - e_b$. From the definition of root-conjugacy we can write $\lambda_a = \lambda'_b = x + 1$ and $\lambda_b = \lambda'_a = x$ for some $x \in \mathbb{Z}$.

From the observation two sentences previous, we have $\lambda''_a \in \{x, x+1\}$ and $\lambda''_b \in \{x, x+1\}$.

In the first case, we have $\lambda''_a = x$. Then $\lambda - \lambda'' = e_a - e_c$ for some $c \neq b$. It follows that $\lambda'_b = \lambda_b = x$ and also follows that $\lambda''_c = x + 1$ and $\lambda_c = x$. Letting $\nu = \lambda - e_a$, we see that λ, λ', and λ'' take the form $\nu + e_a$, $\nu + e_b$, and $\nu + e_c$ where $\nu_a = \nu_b = \nu_c = x$. It is clear that the vector $\nu \in P$ is unique and the sign $\varepsilon = +1$ is determined.

The second case is that $\lambda''_a = x + 1$. Then we have $\lambda''_b - \lambda'_b = e_a - e_c$ for some $c \neq b$. It follows that $\lambda''_b = x + 1$ and that $\lambda'_c = \lambda_c = x + 1$. Letting $\nu = \lambda + e_b$, we can write our three vectors as $\nu - e_b$, $\nu - e_a$, and $\nu - e_c$. As above, everything (including the sign $\varepsilon = -1$) is determined.

To complete the proof we suppose there is a fourth vector λ''' and repeat the argument with λ, λ', and λ'''. If we were in the first case with λ, λ', and λ'', one can check that must also be in the first case with λ'''_a and so on. So the argument persists and the claims about the existence and uniqueness of ν and ε follow. □

Let $\mathcal{C} \subset P$ be a multiset of pairwise compatible weight vectors. We call $\lambda \in \mathcal{C}$ a *root-conjugate vector in $\mathcal{C}$* if it is root-conjugate to some $\mu \in \mathcal{C}$. Otherwise we say that λ

is a W-sortable vector in $\mathcal{C}$. We call the collection $\mathcal{C}$ *basic* if every root-conjugate vector in $\mathcal{C}$ appears with multiplicity one in the multiset.

The following is our main conjecture concerning weights of cluster variables in the cluster algebra $\mathscr{A}(\mathcal{M})$.

Conjecture 5.4. Let $\mathcal{C} \subset P^{\mathbb{M}_\circ}$ be a P-cluster for $\mathscr{A}(\mathcal{M})$ and $p \in \mathbb{M}_\circ$ be a puncture. Then the projection $\pi_p(\mathcal{C}) \subset P$ is a basic multiset of pairwise compatible vectors.

This conjecture is purely about P-clusters at p, so one might be able to prove it by a careful analysis of P-seeds and their mutations.

We denote by $\mathbb{E} = \mathbb{E}(\mathcal{M})$ the exchange graph of the cluster algebra $\mathscr{A}(\mathcal{M})$.

Definition 5.5. We denote by $\mathbb{E}_{\text{good}} \subset \mathbb{E}$ the maximal connected subgraph of the exchange graph containing the initial seed and consisting of seeds whose P-cluster at any puncture p is a basic multiset of pairwise compatible vectors. We call $\mathbb{E}_{\text{good}}$ the *good part* of the exchange graph.

Thus, a cluster is in the good part of the exchange graph if it is connected to the initial seed by mutations which only pass through clusters whose underlying P-cluster satisfy Conjecture 5.4. Conjecture 5.4 is the assertion that $\mathbb{E}_{\text{good}} = \mathbb{E}$. Since we do not have a proof of this conjecture, our results in the next two sections address clusters which are in $\mathbb{E}_{\text{good}}$.

Our first result along these lines is as follows. We say that vertices v and v' are *symmetrical* in a quiver Q if the vertex transposition (vv') induces an isomorphism of Q.

Proposition 5.6. *Let $\mathcal{C} = (Q, \mathrm{wt}(v)_{v \in V(Q)})$ be a P-seed in the good part of the exchange graph. Let v, v' be vertices of its underlying quiver Q whose weights $\pi_p(\mathrm{wt}(v))$ and $\pi_p(\mathrm{wt}(v'))$ are root-conjugate at some puncture p. Then v and v' are symmetrical vertices in Q. Moreover, $\pi_{p'}(\mathrm{wt}(v)) = \pi_{p'}(\mathrm{wt}(v'))$ for punctures $p' \neq p$.*

Proof. Put $\lambda = \pi_p(\mathrm{wt}(v))$ and $\lambda' = \pi_p(\mathrm{wt}(v'))$. By assumption we have $\lambda = \nu + e_a$ and $\lambda' = \nu + e_b$ where $\nu_a = \nu_b$. By the goodness of the P-cluster, all other vertices $u \neq v, v'$ in Q have the property that $\pi_p(\mathrm{wt}(u))_a = \pi_p(\mathrm{wt}(u))_b$.

Consider a vertex $u \in Q$. Let $\mu \in P$ be the result of applying π_p to the weight vector (2) encoding the weight of the right hand side of the exchange relation at u. By considering the left and right hand sides of (2) in turn and using the observations from the previous paragraph, we see that

$$\begin{aligned} \mu_a - \mu_b &= \sum_{u \to v} 1 - \sum_{u \to v'} 1 \\ &= \sum_{v \to u} 1 - \sum_{v' \to u} 1. \end{aligned}$$

On the other hand, only one of the cases $u \to v$ or $v \to u$ is possible and likewise for $u \to v'$ and $v' \to u$. If, for example, the case $u \to v$ occurs, we conclude that the number of arrows $u \to v$ must equal the number of arrows $u \to v'$ and also conclude that there are no arrows $v \to u$ or $v' \to u$. That is, the vertices v and v' are symmetrical in the quiver.

For the second claim about weights at punctures $p' \neq p$, note by weight considerations that the automorphism $\psi_{(a,b),p}$ must send $\mathrm{wt}(v)$ to $\mathrm{wt}(v')$, and it follows that these two weight vectors coincide away from p. □

Example 5.7. The cluster algebra $\mathscr{A}'(\mathrm{SL}_3, D_{2,1})$ has finite cluster type D_4. We indicate the initial mutable P-seed indicated below (drawn inside the digon with puncture p). There are finitely many P-clusters which we computed exhaustively on a computer (code available at the first author's web site). The first column in the table below lists a set of representatives for the W-equivalence classes of P-clusters:

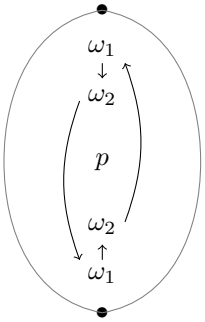

P-cluster	class sums	dosp
$e_1 \times 2, e_1 + e_2 \times 2$	$e_1, e_1 + e_2$	$1\vert 2\vert 3$
$e_1, e_2, e_1 + e_2 \times 2$	$e_1 + e_2$	$12\vert 3$
$e_1 \times 2, e_1 + e_2, e_1 + e_3$	e_1	$1\vert 23$
$e_1, e_2, e_1 + e_2, 0$	$0, e_1 + e_2$	$12\vert 3$
$e_1, e_1 + e_2, e_1 + e_3, 0$	$0, e_1$	$1\vert 23$
$e_1, e_2, e_3, 0$	0	123^+
$e_1 + e_2, e_1 + e_3, e_2 + e_3, 0$	0	123^-

(17)

We write $e_1 + e_2 \times 2$ to indicate that the vector $e_1 + e_2$ appears twice in the multiset, and so on. The initial cluster occupies the first row of the table. The second and third columns of the table are discussed in the next section. We give similar tables for $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$ and $\mathscr{A}'(\mathrm{SL}_4, D_{2,1})$, each of which also has only finitely many P-clusters, in (56) and (58) as part of our study of finite mutation types. By inspecting these tables, one can see that Conjecture 5.4 holds in all three of these examples.

6. From compatible collections to dosps

We explain a way of associating a certain combinatorial object (a "dosp" at every puncture) to P-clusters in the good part of the exchange graph. We identify the operation on these objects which is induced by mutations in the cluster algebra.

6.1. An osp at every puncture

Let $\mathcal{C} \subset P$ be a basic multiset of pairwise compatible weight vectors. We get an equivalence relation on the ground set $\mathcal{C}$ by declaring two vectors $\lambda \neq \mu \in \mathcal{C}$ to be

equivalent when they are root-conjugate. This relation is transitive by Lemma 5.3. Each W-sortable vector in $\mathcal{C}$ lies in its own equivalence class.

Each equivalence class $\mathcal{O} \subset \mathcal{C}$ gives rise to a *class sum*, defined as $\operatorname{sum} \mathcal{O} := \sum_{\lambda \in \mathcal{O}} \lambda \in P$. This construction is indicated in the passage from the first to second column in (17). For example, the fifth row has an equivalence class $\mathcal{O} = \{e_1 + e_2, e_1 + e_3\}$ with $\operatorname{sum} \mathcal{O} = 2e_1 + e_2 + e_3 = e_1 \in P$.

Lemma 6.1. *If $\mathcal{C}$ is a basic multiset of pairwise compatible weight vectors, then its collection of class sums consists of pairwise W-sortable vectors.*

We have listed the P-clusters up to W-action in (17) in such a way that the class sums in each row of the table are sortable by the identity permutation.

Proof. If $\mathcal{C}$ is a pairwise W-sortable collection then the claim is clear. Otherwise let $\mathcal{O}$ be a class of size at least two. We will show that the collection $\mathcal{O}' := (\mathcal{C} \setminus \mathcal{O}) \cup \{\operatorname{sum} \mathcal{O}\}$ is again a basic multiset of pairwise compatible weight vectors and that $\operatorname{sum} \mathcal{O}$ is a W-sortable vector in $\mathcal{O}'$. Repeating this process proves the claim.

Let $B \subset [k]$ be the ground set of the class $\mathcal{O}$ (16). By Lemma 5.3, there is a vector $\nu \in P$ with constant B-coordinates and a sign ε such that $\mathcal{O} = \{\nu + \varepsilon e_a : a \in B\}$. Thus $\operatorname{sum} \mathcal{O} = |B|\nu + \varepsilon \iota_B$ has constant B-coordinates where ι_B is the indicator vector of B. If $\mu \in \mathcal{C} \setminus \mathcal{O}$, then since μ is not root-conjugate to the vectors in $\mathcal{O}$, the definition of compatibility requires that μ is W-sortable with each element of $\mathcal{O}$. In particular, μ must have constant B-coordinates. It follows that μ is pairwise W-sortable with $\operatorname{sum} \mathcal{O}$. □

By this lemma and Lemma 5.3, any basic multiset $\mathcal{C}$ of pairwise compatible weight vectors gives rise to an ordered set partition $\Pi(\mathcal{C})$ via

$$\Pi(\mathcal{C}) := \bigvee_{i=1}^{s} \Pi(\operatorname{sum} \mathcal{O}_i), \tag{18}$$

where we recall that $\bigvee$ denotes the least upper bound operation on osp's. This osp records the global coordinate equalities and strict inequalities which hold for all class sums in $\mathcal{C}$.

This construction is illustrated in the passage from the first column to the third column in (17) provided the reader ignores for the moment the $+, -$ signs appearing as exponents in the third column. (These signs are the extra information of a *decorated* osp and will be explained shortly.) For the collection $\mathcal{C}$ appearing in the fifth row of (17), we have class sums e_1 and 0, with corresponding osp's $\Pi(e_1) = 1|23$ and $\Pi(0) = |123|$ so that $\Pi(\mathcal{C}) = 1|23$.

In this example, note that $\{2, 3\}$ is a ground set of the root conjugacy class $\{e_1 + e_2, e_1 + e_3\} \in \mathcal{C}$ and appears as a block of $\Pi(\mathcal{C}) = 1|23$. This is in fact always the case:

Lemma 6.2. *Continuing the setup of the previous lemma, let B_j be the ground set of an equivalence class $\mathcal{O}_j \subset \mathcal{C}$. Then B_j is a block of $\Pi(\mathcal{C})$.*

Proof. Let $\mathcal{O}_1, \ldots, \mathcal{O}_s$ be a listing of the classes in $\mathcal{C}$. Let us focus on a particular B_j. It is easy to see that B_j is a subset of a block of $\Pi(\mathcal{C})$, i.e. that each of the vectors in the set $S := \{\operatorname{sum} \mathcal{O}_i : i \in [s]\}$ has constant B_j-coordinates. We need to argue that no other coordinate is globally equal to these coordinates on this set S. As in the proof of the previous lemma, we have $\operatorname{sum}(\mathcal{O}_j) = |B_j|\nu + \varepsilon\iota_{B_j}$. We see that the ath coordinate $\operatorname{sum}(\mathcal{O}_j)_a = 0 \mod |B_j|$ when $a \notin B_j$ while $\operatorname{sum}(\mathcal{O}_j)_a = 1 \mod |B_j|$ when $a \in B_j$, hence these coordinates are not equal. □

So far the discussion has focused on a particular copy of the weight lattice P. If we are instead given a P-cluster $\mathcal{C} \subset P^{\oplus \mathbb{M}_\circ}$ appearing in the good part of the exchange graph, we use the projection $\pi_p \colon P^{\oplus \mathbb{M}_\circ} \to P$ to assign to this P-cluster an ordered set partition at each puncture:

$$\Pi_p(\mathcal{C}) := \Pi(\pi_p(\mathcal{C})). \tag{19}$$

We have the following relationship between this recipe and the mutation rule (2) for P-cluster variables. Given a P-seed and mutable vertex v in the underlying quiver, we denote by $\kappa(v) := \sum_{u \to v} \operatorname{wt}(u) \in P^{\mathbb{M}_\circ}$ the weight of the right hand side of the exchange relation describing mutation in direction v out of the current seed.

Proposition 6.3. *Let $\mathcal{C}$ be the P-cluster of a P-seed in the good part of the exchange graph determining an osp $\Pi_p(\mathcal{C})$ at each puncture p. Then each of the weight vectors $\pi_p(\kappa(v))$, as v varies of the vertices of the P-seed, lies in the closed Weyl region $\overline{C_{\Pi_p(\mathcal{C})}}$.*

That is, while the coordinate equalities described by $\Pi_p(\mathcal{C})$ need not hold for the P-cluster variables at p, they do hold for the weights of each exchange relation out of this P-cluster. In particular, the exchange relations for each P-cluster have pairwise W-sortable weights. This is a nontrivial assertion which might possibly be proved inductively by a careful analysis of the mutation rule (2) and might be easier to prove than Conjecture 5.4 itself.

Proof. By Proposition 5.6, vertices in each equivalence class $\mathcal{O}_i \subset \pi_p(\mathcal{C})$ are symmetrical in the corresponding quiver, so their contribution to (2) is symmetrical. So the weight vector $\pi_p(\kappa)$ is a nonnegative linear combination of the vectors $\operatorname{sum} \mathcal{O}_i$ as i ranges over the root-conjugacy classes in $\pi_p(\mathcal{C})$. Each vector $\operatorname{sum} \mathcal{O}_i$ lies in the closure $\overline{C_{\Pi_p(\mathcal{C})}}$ by Lemma 6.1 (this was the definition of $\Pi_p(\mathcal{C})$), hence so does $\pi_p(\kappa)$ since the closed Weyl region is a cone. □

6.2. A dosp at every puncture

The previous section associated an osp at each puncture to a P-cluster in the good part of the exchange graph. We now give a more refined version of this construction in which we associate a choice of sign to certain blocks of such an osp.

Our starting point is the following lemma which says informally that the P-cluster at a puncture p cannot be "too small."

A choice of weight vector $\lambda \in P$ determines a Laurent monomial map $T \to \mathbb{C}^*$ sending $t \mapsto t^\lambda$. If $\mathcal{C} \subset P$ is a multiset of weight vectors of cardinality N, then it determines in this way a map of algebraic tori $T \to (\mathbb{C}^*)^{\mathcal{C}} \cong (\mathbb{C}^*)^N$. Our next lemma addresses what happens when we remove a vector $\lambda \in \mathcal{C}$, getting a map from T to a torus of rank $N-1$.

Lemma 6.4. *Let $\mathcal{C}$ be a P-cluster, $\lambda \in \mathcal{C}$ one of its P-cluster variables, and p a puncture. Then $T \to (\mathbb{C}^*)^{\pi_p(\mathcal{C})\setminus\pi_p(\lambda)}$ is injective.*

That is, we can recover $t \in T$ from the monomials $t^{\pi_p(\lambda)} : \lambda \in \mathcal{C}$, and more strongly we can recover t from any $N-1$ of these monomials. Note that here N is the number of mutable variables in the cluster algebra $\mathscr{A}(\mathcal{M})$, which is much larger than the rank $k-1$ of the algebraic torus T, so the lemma is not so surprising.

Proof. In an initial seed, there are least two vectors of weight ω_i at p for $i = 1, \ldots, k-1$ (because the quiver fragments $Q_k(\delta)$ and $Q_k^1(\delta)$ have at least two variables in each row.) In particular, regardless of which cluster variable we remove, we still have at least one copy of the monomial t^{ω_i} for $i = 1, \ldots, k-1$, and these monomials are enough to recover t.

To complete the proof, we should check that the claimed injectivity propagates under mutation. That is, we assume the injectivity statement for a P-cluster $\mathcal{C}$ and consider a mutation $\lambda \mapsto \lambda'$ resulting in a new P-cluster $\mathcal{C}'$. The monomial map indexed by the vectors in $\mathcal{C}' \setminus \lambda'$ is clearly injective since it coincides with the map indexed by $\mathcal{C} \setminus \lambda$. If $\mu \in \mathcal{C}' \setminus \lambda'$, we need to argue that the monomial map indexed by $\mathcal{C}' \setminus \mu$ is injective. Note that μ appears in at most one of the two monomials appearing in the exchange relation (2) between λ and λ'. Suppose that $t \in \ker(T \to (\mathbb{C}^*)^{\mathcal{C}'\setminus\mu})$. By considering the monomial in the exchange relation that does not involve μ, we see that $t^{\pi_p(\lambda+\lambda')} = 1$. And on the other hand $t^{\pi_p(\lambda')} = 1$ since $\lambda' \in \mathcal{C}' \setminus \mu$. Thus $t^{\pi_p(\lambda)} = 1$ and $t \in \ker(T \to (\mathbb{C}^*)^{\pi_p(\mathcal{C}\setminus\mu)})$. Thus $t = 1 \in T$ using the injectivity statement for the P-cluster $\mathcal{C}$. □

Corollary 6.5. *Let $\mathcal{C}$ be a P-cluster which is in the good part of the exchange graph and p be a puncture. For any coordinates $a \neq b \in [k]$, there is either a pair of vectors $\lambda, \mu \in \pi_p(\mathcal{C})$ which are root-conjugate in the direction $e_a - e_b$, or one of the following global directions of inequality holds:*

- *for all $\lambda \in \pi_p(\mathcal{C})$ one has $\lambda_a \geq \lambda_b$ with strict inequality attained at least twice, or*
- *the same as above but with the direction of inequality reversed.*

As an illustration of this corollary, we could not possibly see the P-cluster multiset $\{e_1 \times 3, e_1 + e_2\}$ in table (17) because the global coordinate inequality between coordinates 2 and 3 is attained only once, not twice.

Proof. The argument only involves weights at p so we can abbreviate $\mathcal{C} = \pi_p(\mathcal{C})$. For $a \neq b \in [k]$ denote by $T_{ab} \subset T$ the rank one subtorus of matrices of the form $(1, \ldots, z, \ldots, z^{-1}, \ldots, 1)$ for $z \in \mathbb{C}^*$, with the z's in positions a and b.

Consider a pair of coordinates $a \neq b \in [k]$. Suppose $\mathcal{C}$ does not contain vectors which are root-conjugate by the root $e_a - e_b$. By the definition of pairwise compatibility, there is a global inequality $\lambda_a \geq \lambda_b$ or $\lambda_a \leq \lambda_b$ holding for all vectors λ. By considering the action of the torus T_{ab} and using Lemma 6.4, we see that the strict version of this inequality must be strict for at least two vectors $\lambda \in \mathcal{C}$. □

Definition 6.6. A *decorated ordered set partition* (a *dosp* for short) is a pair (Π, ε) consisting of an osp Π together with a choice of sign $\varepsilon(B)$ for each block B of Π of cardinality at least three.

Recall the construction (19) assigning an osp at each puncture to any P-cluster $\mathcal{C}$ in the good part of the exchange graph.

Lemma 6.7. *Let $\mathcal{C}$ be a P-cluster in the good part of the exchange graph, p be a puncture, and B be a non-singleton block of $\Pi_p(\mathcal{C})$. Then B is the ground set of a root-conjugacy class in $\pi_p(\mathcal{C})$.*

That is, every non-singleton block of $\Pi_p(\mathcal{C})$ "comes from" a root-conjugacy class in $\pi_p(\mathcal{C})$.

Proof. Suppose that a, b reside in a block of $\Pi_p(\mathcal{C})$ and are not in part of a ground set. By the corollary, there is a vector $\lambda \in \Pi_p(\mathcal{C})$ satisfying a strict inequality $\lambda_a > \lambda_b$ (or the opposite strict inequality, but this will not change the argument). On the other hand, passing from the collection $\Pi_p(\mathcal{C})$ to the collection of class sums, we only turn a strict inequality between coordinates into an equality of coordinates when the two coordinates are in the ground set of a root-conjugacy class. This contradicts the fact that the coordinates a and b are supposed to be equal on class sums. □

Using Lemma 6.7, we upgrade the recipe (19) to a "P-cluster to dosp" map

$$\mathcal{C} \mapsto ((\Pi_p(\mathcal{C}), \varepsilon_p(\mathcal{C}))_{p \in \mathbb{M}_\circ}, \tag{20}$$

defined for any cluster in the exchange graph whose P-cluster satisfies Conjecture 5.4. Namely, we can define the sign $\varepsilon(B)$ for any block of cardinality at least three to be the sign associated to its underlying root-conjugacy class by the construction in Lemma 5.3.

6.3. Dosp mutations

By the previous section, we can assign to any vertex of the good part of the exchange graph a tuple of dosps indexed by punctures. In this section, we show that every mutation of P-clusters changes these dosps in a predictable way.

Definition 6.8. A *mutation* of dosps is the following local move on decorated blocks:

$$(L \cup \{a\})^+ \,|\, R^- \leftrightarrow L^+ \,|\, (\{a\} \cup R)^-, \tag{21}$$

where $L, R \subset [k]$ are adjacent blocks, possibly the empty block.

When one of L or R is the empty block, the mutation move amounts to merging a singleton block with a neighboring block or the inverse of this move. If a block has cardinality one or two, then for the purposes of mutation one can regard the block as both signs $+$ and $-$ in (21). See the below example for further clarification.

Example 6.9. The following is a complete list of dosps which can be obtained from the dosp $12|345^+|6$ by a dosp mutation, with terms grouped if they are related by W-action: i) $12|3456^+$, ii) $1|2|345^+|6$ or $2|1|345^+|6$, iii) $12|34|56$ or $12|35|46$ or $12|45|36$, and finally iv) $12|34|5|6$ or $12|35|4|6$ or $12|45|3|6$.

We now justify our use of the word mutation.

Lemma 6.10. *Let $\mathcal{C}$ and $\mathcal{C}'$ be mutation-adjacent P-clusters in the good part of the exchange graph and consider a puncture $p \in \mathbb{M}_\circ$. Then the corresponding dosps $(\Pi_p(\mathcal{C}), \varepsilon(\mathcal{C}))$ and $(\Pi_p(\mathcal{C}'), \varepsilon(\mathcal{C}'))$ either coincide or are related by a dosp mutation.*

Proof. The argument is local in the puncture p so we abbreviate $\mathcal{C} = \pi_p(\mathcal{C})$, $\Pi = \Pi_p$, $\lambda = \pi_p(\lambda)$ etc.

Using Corollary 6.5, a given pair of coordinates a, b either satisfies a global inequality on $\mathcal{C}$ with strictness attained at least twice, or these coordinates are part of a ground set of a root-conjugacy class (in which case there is a strict inequality between these coordinates in both directions, attained exactly once). Using the data of the direction of these inequalities, and how many times they are attained, we clearly can recover the dosp.

Fixing a and b, let us analyze how the direction of inequality can change when we perform a P-cluster mutation removing λ and inserting some λ'. We use primes to denote the result of mutation, e.g. denoting the new P-cluster by $\mathcal{C}'$.

First we suppose there is no global inequality between coordinates a and b in $\mathcal{C}$. That is, these two coordinates reside in the same block $B \ni a, b$ of Π_1, or equivalently, in the ground set of a root-conjugacy class in $\mathcal{C}$. We can write the vectors in this root-conjugacy class as $\nu + \varepsilon(B)e_s$ as s varies over the block B, where ν has constant B-coordinates. The

inequality between a and b will only change if λ is either $\nu + \varepsilon(B)e_a$ or $\nu + \varepsilon(B)e_b$, and up to symmetry we can assume $\lambda = \nu + \varepsilon(B)e_a$. Removing λ removes a from the block B. The coordinates a and b must reside in different blocks of Π', i.e. there is a global inequality between these coordinates which holds on $\mathcal{C}'$ and is attained at least twice. Let $\mu := (ab)\lambda$; this vector lies in $\mathcal{C}$ hence also in $\mathcal{C}'$. When the block B has cardinality at least three, the sign $\varepsilon(B)$ determines the inequality between the coordinates μ_a and μ_b, hence determines the direction of inequality between coordinates a and b in Π'. (When $\varepsilon(B) = +1$, a must reside to the right of the coordinate b; when $\varepsilon(B) = -1$, a must reside to the left of the coordinate b.)

If there *is* a global inequality between coordinates a and b in $\mathcal{C}$, then because the strict inequality is attained at least twice, $\mathcal{C}'$ either has the same global inequality, or a and b are in the same block of $\mathcal{C}'$ and we are in the case of the previous paragraph.

After possibly interchanging λ and λ', then, we can assume that the removed vector $\lambda = \nu + \varepsilon(B)e_a$ is a root-conjugate vector residing in a root-conjugacy class $\mathcal{O}$ with ground set B, where ν has constant B-coordinates. Removing λ only changes the coordinate inequalities involving a. (All other coordinate inequalities are unchanged since they are witnessed by the vectors in $\mathcal{O}\backslash\lambda$.) If $\varepsilon(B) = +1$ (resp. $= -1$), we know that the coordinate a moves right (resp. left), and since no other coordinate inequalities are affected, a must move to the block immediately right of the block $B \setminus \{a\}$ in Π'. (If $\{a\}$ is a singleton block, then this block will be immediately right of $B \setminus \{a\}$.) To summarize: only the a-coordinate "moves" in passing from Π to Π'; we know the direction that this coordinate must move provided the block B has cardinality at least three, and we know that it must move into an immediately adjacent block. We get the dosp mutation rule (21). □

Our next lemma is a continuation of the previous one, but now considering all punctures at once.

Lemma 6.11. *Let $\mathcal{C}$ and $\mathcal{C}'$ be mutation-adjacent P-clusters in the good part of the exchange graph. Then the dosps $(\Pi_p(\mathcal{C}), \varepsilon_p(\mathcal{C}))$ and $(\Pi_p(\mathcal{C}'), \varepsilon_p(\mathcal{C}'))$ disagree at at most one puncture.*

Our proof relies in part on the assertion that the quiver for our moduli space has no isolated vertices if we exclude the special case of a once-punctured digon and $G = \mathrm{SL}_2$. This special case presents no counterexample to our lemma so we can ignore it. The assertion itself can be argued by a painful case analysis.

Proof. By Lemma 6.10, the dosps at each puncture either coincide or are related by a dosp mutation. We suppose that the mutation removes a vector λ and inserts a vector λ'. Our reasoning below relies on the observation at the construction of a dosp from a collection $\mathcal{C}$ depends only on $\mathcal{C}$ as a set, not as a multiset.

Suppose that a dosp mutation occurs at a puncture p. We will show it does not happen at punctures $q \in \mathbb{M}_\circ$ with $q \neq p$. As in the last paragraph of the previous proof, we can

assume (after possibly interchanging roles) that $\pi_p(\lambda) \in \pi_p(\mathcal{C})$ is a root-conjugate vector, root-conjugate to some $\mu \in \mathcal{C}$ at p. Using Proposition 5.6, we know that the vectors λ and μ coincide at q. Thus, in removing λ and inserting λ', we are removing a vector which is present at least twice in the multiset $\pi_q(\mathcal{C})$ and perhaps adding a new vector. In particular, $\pi_q(\lambda) \in \pi_q(\mathcal{C})$ is not a root-conjugate vector. By the same reasoning, if $\pi_p(\lambda') \in \pi_p(\mathcal{C}')$ is *also* a root-conjugate vector, then we have $\pi_q(\mathcal{C}) = \pi_q(\mathcal{C}')$ as sets (not as multisets). So in this case, the dosp at q does not change.

Suppose seeking contradiction that adding the vector λ' changes the dosp at q. Then $\pi_q(\lambda')$ must be a root-conjugate vector in $\pi_q(\mathcal{C}')$. (We argued in the previous proof that when a dosp mutation occurs, one of the two vectors involved in the mutation is a root-conjugate vector, and we know that $\pi_q(\lambda)$ is not a root-conjugate vector.) By the reasoning in the previous proof, there is a strict inequality between a pair of coordinates, say $\nu_c \geq \nu_d$, which holds on $\pi_q(\mathcal{C})$ and which is strictly attained exactly twice in the multiset $\pi_q(\mathcal{C})$. Moreover, $\pi_q(\lambda)$ must be one of the two vectors which attains this strict inequality because removing λ and inserting λ' changes the global inequality between c and d into a root-conjugacy. And since λ and μ coincide at q, the vector $\pi_q(\mu)$ is the other vector in $\pi_q(\mathcal{C})$ attaining the strict inequality. Finally, by Proposition 5.6, the P-cluster variables λ and μ are symmetrical in the quiver Q underlying the P-seed of $\mathcal{C}$. Let v be a vertex of the P-seed which is adjacent to the vertices λ and μ in Q; such vertex exists because Q has no isolated vertices. For simplicity assume that the arrows point from λ, μ to v in Q, so that λ and μ contribute to the first term in (2) describing mutation at vertex v in the P-seed. By our assumptions, π_q applied to this first sum $(\sum_{u \to v \text{ in } Q} \mathrm{wt}(u))$ has c-coordinate strictly greater than its d-coordinate because λ and μ both appear in this sum and all other summands have equal c and d coordinates. On the other hand, π_q applied to the second sum $(\sum_{v \to u \text{ in } Q} \mathrm{wt}(u))$ has equal c and d coordinates for the same reasons. But by the balancing condition for P-seeds, these two sums must equal one another, yielding the desired contradiction. □

7. Realizing dosp mutations via cluster mutations

We explore in this section the relation between the exchange graph of the cluster algebra $\mathscr{A}(\mathcal{M})$ and a simpler finite graph.

First we make some standard graph theoretical definitions. Let $G = (V, E)$ and $G' = (V', E')$ be simple undirected graphs without isolated vertices. A map $f\colon V \to V'$ on their vertex sets is *simplicial* if for each edge $(u, v) \in E$, one either has $f(u) = f(v)$ or has $(f(u), f(v)) \in E'$. A simplicial map is an *edge contraction map* if every $e' \in E'$ takes the form $f(e)$ for *some* $e \in E$. (This implies that f is surjective on vertices.) Viewing graphs as 1-dimensional CW complexes, an edge contraction map determines a homotopy from G to G' while a simplicial map induces a homotopy onto a subgraph of G'.

Definition 7.1. The *dosp mutation graph* $H_{\mathrm{dosp}} = H_{\mathrm{dosp}}(k)$ is the graph with vertex set consisting of the dosps on $[k]$ and with edges the dosp mutations.

When we have fixed in our minds a marked surface $\mathbb{S}$, we associate the graph $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$, the Cartesian product of copies of H_{dosp} indexed by the punctures in $\mathbb{S}$.

Thus, a vertex of $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$ corresponds to a choice of at each puncture, and an edge corresponds to a dosp mutation at one of the punctures.

Example 7.2. When $k = 2$, the dosp mutation graph is the path graph on three vertices

$$H_{\mathrm{dosp}}(2) = 1|2 \quad\text{——}\quad |12| \quad\text{——}\quad 2|1. \tag{22}$$

The left and right vertices correspond to tagged triangulations which are plain versus notched at a puncture p. The middle vertex corresponds to tagged triangulations containing one plain and one notched arc at p. When $|\mathbb{M}_\circ| = 2$, the Cartesian square $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$ of this graph is the 3×3 grid graph.

The graph $H_{\mathrm{dosp}}(3)$ is the right graph drawn in Fig. 3. Its relationship with the left graph in the figure, i.e. the exchange graph of a type D_4 cluster algebra, is explained in Section 10.

The graph $H_{\mathrm{dosp}}(4)$ has 84 vertices and seems unwieldy to draw. We draw instead its quotient by the W-action:

$H_{\mathrm{dosp}}(4)/W =$

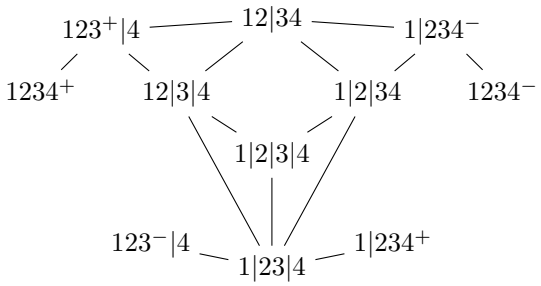

(23)

We now identify a relationship between the good part of the exchange graph and the graph $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$.

Consider a P-cluster in the good part of the exchange graph. By the recipe (20), it determines a dosp at each puncture, i.e. a vertex of $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$. By Lemmas 6.10 and 6.11, the map $\mathbb{E}_{\mathrm{good}} \to \square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$ is a simplicial map of graphs. Our next theorem addresses the image of this simplicial map. That is, we ask which dosps and which dosp mutations are "witnessed" by clusters and their mutations.

We call a dosp *oneblock* if its underlying osp has only one block. When $k > 2$ there are two oneblock dosps. We call a vertex of $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$ *everywhere oneblock* if its dosp at each puncture is oneblock; there are $2^{|\mathbb{M}_\circ|}$ many such vertices when $k > 2$.

Theorem 7.3. *Suppose that $\mathbb{S} \neq D_{1,h}$ for some $h \geq 1$. Then the simplicial map (20) has the following properties:*

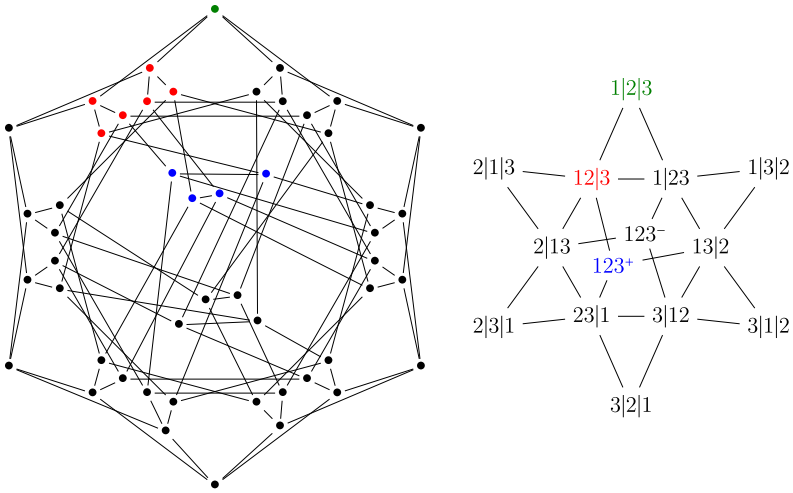

Fig. 3. The exchange graph for a D_4 cluster algebra and its edge contraction on to the dosp mutation graph, illustrating Theorem 7.3 when $\mathcal{M} = \mathrm{Gr}_{\mathrm{SL}_3, D_{2,1}}$. The green (resp. red, resp. blue) colors indicate the fibers of the edge contraction map.

(1) If $\mathbf{S}$ *has boundary, then* $\mathbb{E}_{\mathrm{good}} \to \square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$ *is an edge contraction map.*
(2) If $\mathbf{S}$ *has no boundary (hence when* $\mathbb{S} = S_{g,h}$*), then the map* $\mathbb{E}_{\mathrm{good}} \to \square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$ *misses the everywhere oneblock vertices.*

Conjecture 7.4. When $\mathbb{S} = S_{g,h}$ and $(k, \mathbb{S}) \neq (2, S_{0,3})$, the map (20) is an edge contraction onto the induced subgraph on the vertices which are not everywhere oneblock.

Let H' be the result of deleting the everywhere oneblock vertices from the graph $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$. We will prove momentarily that the image of (20) is contained in this smaller graph H' when $\mathbb{S}$ has no boundary. A verification of the above conjecture would involve showing that every edge in H' has a preimage under the map P-cluster to dosp map (20). This is a tractable explicit question which we do not address in this paper.

Remark 7.5. We expect that (1) from the above theorem holds for punctured monogons $D_{1,h}$ once $h \geq 2$, but there is a technical difficulty with the proof we present. We have excluded once-punctured monogons in this paper, but if one considers them, then the statement is false when $k = 2, 3$. It most likely holds when $k > 3$.

Remark 7.6. We illustrate (2) from Theorem 7.3 when $k = 2$. A tagged triangulation T is oneblock at p if there exists tagged arcs $\gamma, \gamma^{\bowtie} \in T$ with γ plain and $\gamma^{\bowtie}$ its notching at p. The other endpoint of γ is at a puncture q since there are no boundary points. The tagged arcs γ and $\gamma^{\bowtie}$ must be tagged the same way at q, hence T is not oneblock at q. That is, it is not possible to be oneblock at *every puncture* when $\mathbb{S}$ has no boundary.

By an explicit argument (constructing appropriate tagged triangulations and flips between them), one can check that Conjecture 7.4 holds in this case. That is, every

vertex which is not oneblock at every puncture has a preimage under (20), as does any edge between two such vertices.

We now prove part (2) of Theorem 7.3 based on the ideas of the preceding remark:

Proof of (2) from Theorem 7.3. For a weight vector $\lambda = (\lambda_1, \dots, \lambda_k) \in P$, $\operatorname{sum} \lambda := \sum_i \lambda_i \mod k$ is well defined. By induction on mutations, every P-cluster variable λ has the property that $\sum_p \operatorname{sum} \pi_p(\lambda) = 0 \mod k$. There are only two options for a collection $\mathcal{C} \subset P$ to have a oneblock dosp: either $\mathcal{C}$ consists of the vectors $e_1, \dots, e_k$ and multiple copies of the zero vector, or it consists of $-e_1, \dots, -e_k$ and multiple copies of the zero vector. By Proposition 5.6, if a P-cluster variable λ is a root-conjugate vector at p then it is not root-conjugate at any other puncture. If the P-cluster containing λ has everywhere oneblock dosps, then λ has zero weight at every puncture $q \neq p$. But then $\operatorname{sum} \lambda = \pm 1 \mod k$ which is a contradiction. □

Our proof of (1) from Theorem 7.3 is more involved. First, we prove the statement for the special cases of $\mathscr{A}'(\mathrm{SL}_k, D_{2,1})$ and then $\mathscr{A}'(\mathrm{SL}_k, D_{2,1})$ by providing explicit mutations. Then we reduce the assertion for an arbitrary version of the moduli space to this case by simple "cutting and gluing" type arguments.

The reader should consult the following example, which contains the heart of the whole argument, before reading the proof of Lemma 7.8.

Example 7.7. The cluster $\mathscr{A}'(\mathrm{SL}_7, D_{2,1})$ has a canonical initial seed since the once-punctured digon has a unique taut triangulation. Its P-cluster (with weight vectors written in multiplicative notation) is below:

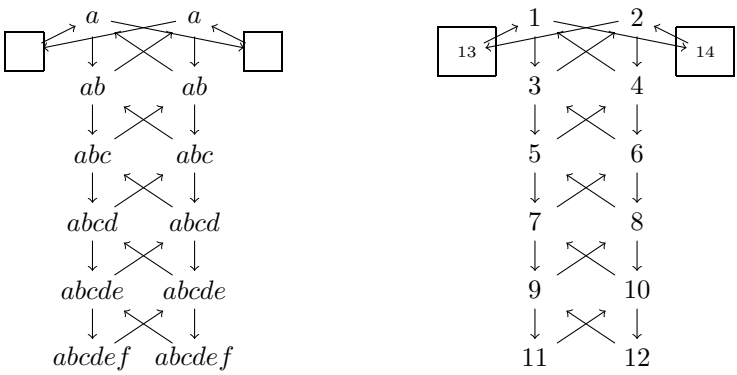

(24)

The two frozen vertices are boxed and have weight zero, hence can be deleted. The dosp associated to this P-seed is $1|2|3|4|5|6|7$. We now describe a sequence of P-seed mutations in the good part of the exchange graph which, after applying the "P-cluster to dosp map" (20), witnesses the sequence of dosp mutations

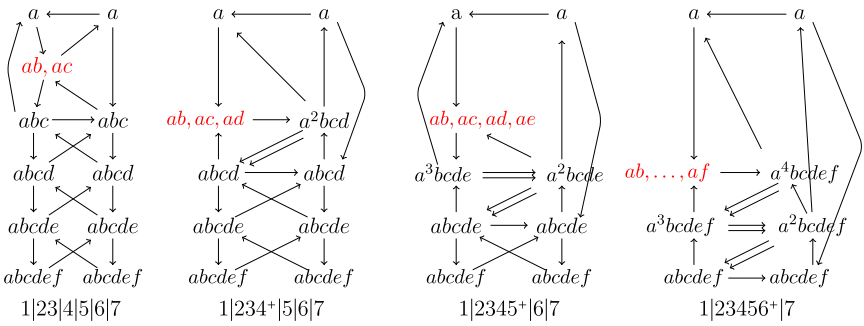

Fig. 4. A sequence of P-seeds appearing as intermediate steps of the mutation sequence from Example 7.7 and the proof of Theorem 7.3. Each P-seed is written above the dosp that its P-cluster exhausts. Red vertices indicate merging of several symmetrical vertices of the P-seed (thereby reducing the number of edges in the picture), e.g. ab, ac is obtained by gluing a vertex of weight ab with one of weight ac, creating a vertex of weight a^2bc.

$$1|2|3|4|5|6|7 \mapsto 1|23|4|5|6|7 \mapsto 1|234^+|5|6|7 \mapsto 1|2345^+|6|7 \mapsto 1|23456^+|7.$$

In the first step, perform mutation μ_4. The result is the leftmost P-seed in Fig. 4, drawn with vertices 3 and 4 merged and highlighted in red. The dosp of this P-seed is $1|23|4|5|6|7$. In the second step, mutate $\mu_5 \circ \mu_6$ to obtain the next P-seed in Fig. 4, with vertices 3,4,5 merged and in red. The dosp is $1|234^+|5|6|7$. Third, perform $\mu_6 \circ \mu_7 \circ \mu_8$ to obtain the next P-seed, with vertices 3,4,5,6 merged and with dosp $1|2345^+|6|7$. Finally mutate $\mu_7 \circ \mu_8 \circ \mu_9 \circ \mu_{10}$ to obtain the fourth seed, with dosp $1|23456^+|7$.

We can simulate the dosp mutations

$$1|23456^+|7 \mapsto 1|2345^+|6|7 \mapsto 1|234^+|56|7 \mapsto 1|23|456^-|7 \mapsto 1|2|3456^-|7 \mapsto 1|23456^-|7$$

by mutating at the red vertices in the rightmost P-seed one by one. That is, performing the mutation μ_7 we get $af \mapsto a^3bcde$ which witnesses the first of these dosp mutations. Performing μ_6 we get $ae \mapsto a^3bcdf$ and the resulting P-cluster exhausts $1|234^+|56|7$.

Lemma 7.8. *Theorem 7.3 holds for* $\mathscr{A}'(\mathrm{SL}_k, D_{2,1})$.

Proof. There is only puncture in this case. What we need to show is that every dosp is reachable inside the good part of the exchange graph by a sequence of mutations from the initial P-seed, and that every dosp mutation comes from a mutation of P-clusters inside the good part of the exchange graph.

The initial P-cluster when $G = \mathrm{SL}_k$ is the ladder quiver as in (24), but with $k-1$ rows. We number its vertices from left to right and top to bottom as in (24). Mutation at both vertices in the ith row is the Weyl group action s_i. Precompose with such mutations as needed, we can prove the desired surjectivity up to W-action.

The proof is based on the mutation sequences given in Example 7.7. If one wants to simulate a dosp mutation $(L \cup \{a\})^+|R^- \mapsto (L)^+|(R \cup \{a\})^-$, one should create the

positively decorated block $(L \cup R)^+$ using the first sequence of mutations given in that example. (Mutating down the rows of the ladder quiver, starting either in the right column or the left column depending on parity.) One can then "peel off" the coordinates in R by mutating at the corresponding merged "red" vertices, arriving at the dosp $(L \cup \{a\})^+|R^-$. Mutating at one more of these vertices implements the desired dosp mutation.

The mutation sequences involved are "local" in the coordinates $L \cup R$. That is, we can create "the rest of" the dosp on the coordinates $[k] \setminus (L \cup R)$ using the same mutation sequences, because the mutations in distinct blocks commute. To illustrate, in Example 7.7, we create the decorated block 23456^+ by mutating at rows $2, 3, 4, 5$ only. We would be able to merge 7 into a block (of any sign we like) by performing the needed mutations starting in row 7. □

For the next step in the proof of Theorem 7.3, we will show that various P-seeds are mutation-equivalent to each other provided we ignore vertices of weight zero. We will use the notation $\Sigma_1 \rightsquigarrow \Sigma_2$ to indicate that we can reach the P-seed Σ_2 from Σ_1 by a sequence of mutations which do not violate Conjecture 5.4 and deletion of weight zero vertices.

The once-punctured m-gon $D_{m,1}$ has a unique taut triangulation $\Delta = \Delta_m$. It gives rise to initial P-seeds $\Sigma_k(\Delta_m)$ and $\Sigma'_k(\Delta_m)$ for $\mathscr{A}(\mathrm{SL}_k, D_{m,1})$ and $\mathscr{A}'(\mathrm{SL}_k, D_{m,1})$ respectively.

We upgrade the intermediate quiver Q_k^s defined in Section 3.3 to a P-seed Σ_k^s by assigning to the vertex indexed by the triple $(a_1, a_2, a_3) \in \mathbb{H}_k$ the weight $\omega_{a_1} \in P$. We denote by $\Sigma_k^s \underset{\sim}{\coprod} \Sigma_k^t$ the result of gluing two such P-seeds along their left and right boundary sides. Thus $\Sigma_k^{k-1} \underset{\sim}{\coprod} \Sigma_k^{k-1} = \Sigma_k(\Delta_2)$ and $\Sigma_k^1 \underset{\sim}{\coprod} \Sigma_k^1 = \Sigma'_k(\Delta_2)$.

Lemma 7.9. *We have the following $\rightsquigarrow$-relationship between P-seeds:*

$$\Sigma_k(\Delta_2) \rightsquigarrow \Sigma'_k(\Delta_k) \tag{25}$$

$$\Sigma_k^s \underset{\sim}{\coprod} \Sigma_k^1 \rightsquigarrow \Sigma_k^{s-1} \underset{\sim}{\coprod} \Sigma_k^1 \text{ for } s \in [2, k-1] \text{ for any } k. \tag{26}$$

Proof. For the first statement (25), the frozen variables play no role since they have weight zero at the puncture. We claim that the mutable parts of the P-seeds $\Sigma_k(\Delta_2)$ and $\Sigma'_k(\Delta_k)$ are mutation-equivalent, establishing (25). This is a particular instance of the conjectured quasi cluster isomorphism asserted in Remark 3.5. We argue the statement we need in this case as follows.

There is a well known recipe that associates to a quiver to any reduced word for $w_0 \in W$. One can upgrade this quiver to a P-seed by assigning vertices in row i the weight ω_i. The seed Σ_k^{k-1} corresponds to the reduced word $1, \ldots, k-1, 1, \ldots, k-2, \ldots, 1, 2, 1$ and the seed $\Sigma_k(\Delta_2)$ is the gluing of two such seeds (corresponding to the two triangles in

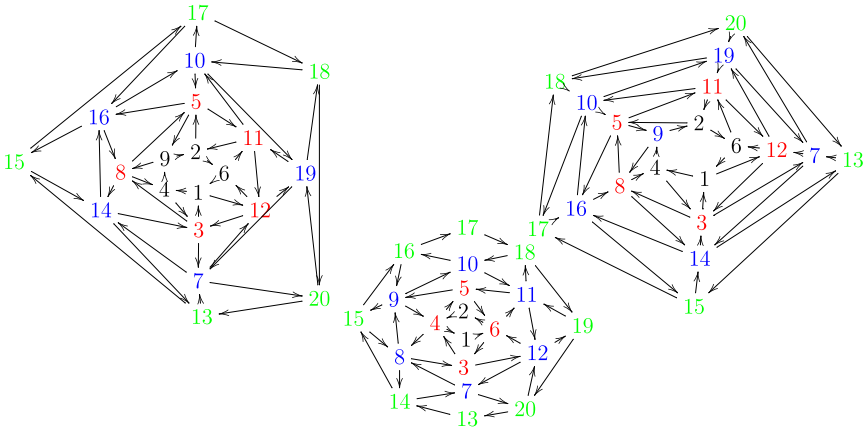

Fig. 5. From the proof of Lemma 7.9. The bottom quiver is that of $\Sigma_5(\Delta_2)$. The vertices $1, 3, 7, 13$ resp. $2, 5, 10, 17$ sit on the two internal edges of the triangulation. The upper left quiver is mutation-equivalent to the bottom quiver. It is obtained by gluing the quiver for the reduced word 1324132413 with that for 2413241324 along the edges of the triangulation. Rearranging the vertices of this quiver we get the upper right quiver, which is $\Sigma'_5(\Delta_5)$. The colors indicate weights of P-cluster variables.

the triangulation of the once-punctured digon) along their shared edges. One knows that we can modify our reduced word (in either triangle) and obtain a mutation-equivalent seed. Let $o = \prod_{i \text{ odd}} s_i$ and $e = \prod_{i \text{ even}} s_i$ be the product of odd (resp. even) Coxeter generators. Then $oeoe\cdots$ is reduced word for w_0 and so is $eoee\cdots$ (with k terms in either product). Depending on parity of k, one can either glue $oeoe\cdots$ to itself, or glue $oeoe\cdots$ with $eoeo\cdots$, so that the letters alternate as we read the two triangles in cyclic order.

Deleting the frozen variables, which have weight 0, gives the P-seed $\Sigma'_k(\Delta_k)$. We illustrate this argument when $k = 5$ in Fig. 5. The mutation needed to create the upper left quiver from the bottom quiver is $\mu_{19,12,11,6}$ (which creates the word $oeoeo$ in the right triangle) followed by $\mu_{14,16,9,4,8,9}$ (which creates $eoeo$ in the left triangle). Rearranging the vertices, we recognize that this P-seed is in fact that of $\Sigma'_k(\Delta_k)$, proving (25).

The proof of the second statement (26) is illustrated in Fig. 6 when $k = 8$. The topmost figure is $\Sigma_8^7 \underset{\sim}{\coprod} \Sigma_8^1$; moving rightwards and downwards across the page we see $\Sigma_8^6 \underset{\sim}{\coprod} \Sigma_8^1$ then $\Sigma_8^5 \underset{\sim}{\coprod} \Sigma_8^1$ etc. Each row of Σ_k^s has at most s many variables. To pass from $\Sigma_8^s \underset{\sim}{\coprod} \Sigma_8^1$ to $\Sigma_8^{s-1} \underset{\sim}{\coprod} \Sigma_8^1$, one should locate the highest row that has s variables in it. There will be a variable x that has no arrows to higher rows; one should mutate down the column containing x, starting in the row containing x. When a variable in row i is mutated, its weight changes $\omega_i \mapsto \omega_{i-1}$ and one should drag it down a row. One deletes a variable when it has weight zero. After mutating down a column, move one column right and mutate down this column starting one row lower than in the previous column. □

Lemma 7.10. *Theorem 7.3 holds when $\mathcal{M} = \mathcal{A}_{\mathrm{SL}_k, D_{m,1}}$ for any $m \geq 2$.*

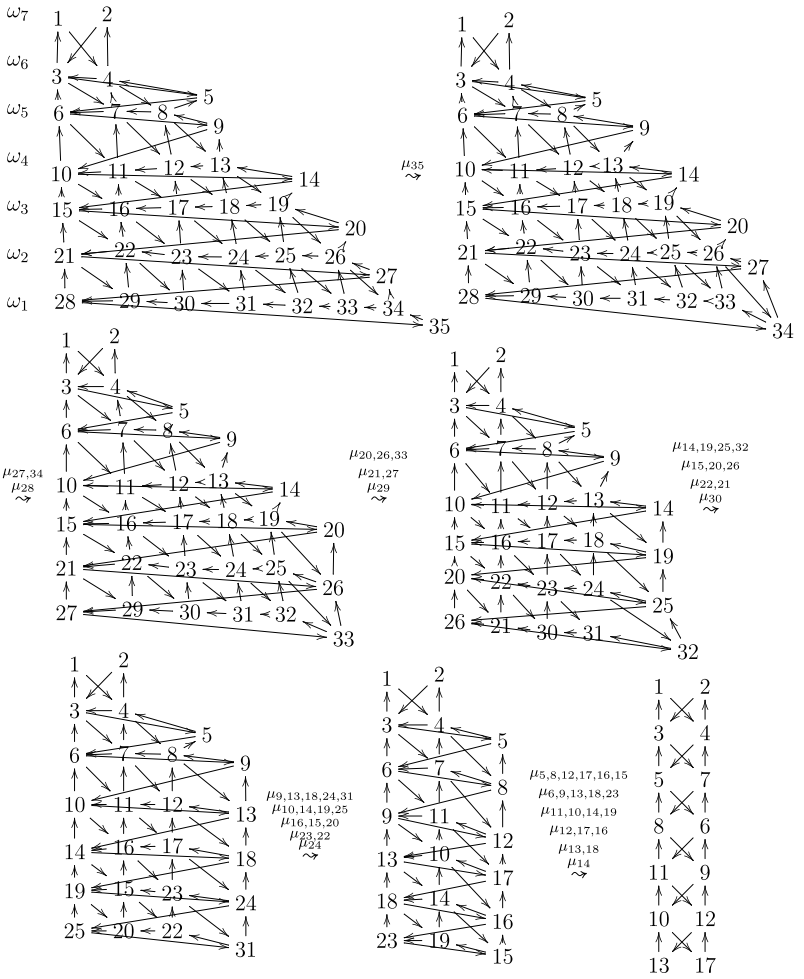

Fig. 6. From the proof of Lemma 7.9. The mutations and vertex deletions proving that $\Sigma_8^s \underset{\sim}{\coprod} \Sigma_8^1 \rightsquigarrow \Sigma_8^{s-1} \underset{\sim}{\coprod} \Sigma_8^1$ for $s \in [2, 7]$.

Proof. Cyclically number the arcs in the taut triangulation Δ_m of $D_{m,1}$ as $1, \ldots, m$. Flipping at arc 1 yields a new arc which is a boundary arch contractible to two boundary intervals. After performing this flip, we see a copy of Δ_{m-1} surrounding the puncture plus an extra "exterior triangle" not incident to the puncture. We will argue momentarily that we can simulate this flip of triangulations at the level of P-seeds in such a way that the P-cluster variables in the exterior triangle have zero weight and can be ignored. Deleting these variables of zero weight, the resulting P-seed is an $m-1$-cycle consisting of $m-2$ copies of Q_k^1 followed by one copy of Q_k^2 (corresponding to the triangle which has the flipped arc 1), with consecutive fragments glued along shared edges.

We next mutate at arc 2 to get $m-2$ triangles surrounding the puncture, with the new arc a boundary arch contractible to three boundary intervals. The corresponding P-seed

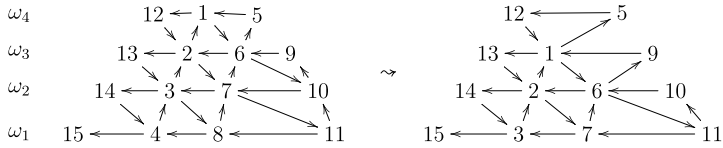

Fig. 7. From the proof of Lemma 7.10. Vertices $1, \dots, 11$ in the left figure form a copy of Q_5^2. We glue a copy of Q_5^1 to this quiver on the left with vertices $1, 2, 3, 4$ the glued vertices. Mutating at $1, 2, 3, 4, 6, 7, 8$ and deleting vertices $4, 8$ (which have zero weight) yields the right quiver, which is Q_5^3.

(after deleting vertices of zero weight) is an $m-2$-cycle consisting of $m-3$ copies of Q_k^1 followed by one copy of Q_k^3. We continue flipping in this way, flipping arc $3, \dots, m-2$, until the punctures sit in a once-punctured digon. One of the triangles surrounding the puncture has an exposed side (corresponding to the seed fragment Σ_k^1) and the other has a side which is contractible to $m-1$ boundary intervals (corresponding to the seed fragment $\Sigma_k^{\min(m-1,k-1)}$). In any case, the seed fragment is $\Sigma_k^s \underset{\sim}{\coprod} \Sigma_k^1$ for some s. Since $\Sigma_k^s \underset{\sim}{\coprod} \Sigma_k^1 \rightsquigarrow \Sigma_k^1 \underset{\sim}{\coprod} \Sigma_k^1$, the claim follows from Lemma 7.8.

What remains is to explain the mutations of P-cluster variables which realize the flips. We start with two adjacent Q_k^1's and claim that this $\rightsquigarrow Q_k^2$. Then we glue a copy of Q_k^1 to this Q_k^2 (simulating adding the next triangle), and claim that this $\rightsquigarrow Q_k^3$. The argument is indicated in Fig. 7 in the case of Q_5^1 glued to Q_5^2. In general, one should mutate down all of the "interior" diagonals from top to bottom, starting with the leftmost interior diagonal and ending with the rightmost diagonal. One should slide a variable one unit down its diagonal once it is mutated, and delete vertices of weight zero. The result in the above example is Q_5^3, as in the figure. □

Proof of part (1) of Theorem 7.3. By Lemmas 6.10 and 6.11, the map (20) is an edge contraction onto some subgraph of $\square_{\mathbb{M}_\circ} H_{\mathrm{dosp}}$. We need to argue the surjectivity on edges (hence on vertices). Since we are arguing part (1) of the theorem, $\mathbf{S}$ has boundary.

If there are at least two boundary components, then it is possible to choose a regular triangulation of $\mathbb{S}$ in which every puncture $p \in \mathbb{M}_\circ$ sits inside a once-punctured digon, with the two endpoints of the digon residing on different boundary components of $\mathbf{S}$. In particular, the two sides of such digon are not boundary arches, so that the seed near p is that of $\Sigma_k(\Delta_2)$ (in either version of the moduli space). Since $\Sigma_k(\Delta_2) \rightsquigarrow \Sigma'_k(\Delta_k)$, the claim follows from Lemma 7.10.

The final case to consider is that there is only one boundary component but more than one puncture. As in the theorem statement, we also assume that the boundary carries at least two boundary points. Then we can enclose all of the punctures inside once-punctured digons as before (this fails if there is only one boundary point, which is why we exclude it from Theorem 7.3). For the Fock-Goncharov case, the preceding argument works. For the Grassmannian case, the leftmost and rightmost punctures will be enclosed in digons which have an exposed triangle. The corresponding seed will be $\Sigma_k^{k-1} \coprod \Sigma_k^1 / \sim$ and the claim follows from Lemma 7.9 and the digon case Lemma 7.8. □

8. Tensor diagrams

We provide background on tensor diagrams on a surface, skein relations, and arborization moves following [3,8,7]. Then we introduce tagged and pseudotagged diagrams and state our cluster combinatorics conjectures.

8.1. Plain tensor diagrams

Let $\mathcal{M}$ be one of the moduli spaces considered in this paper, i.e. either $\mathcal{M} = \mathcal{A}_{\mathrm{SL}_k,\mathbb{S}}$ or $\mathcal{M} = \mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$.

Definition 8.1. An $\mathcal{M}$-*tensor diagram* is an immersed edge-labeled acyclic directed graph $T = (V(T), E(T)) \subset \mathbf{S}$ satisfying the following:

- Every marked point is a vertex of T and is a source vertex of T. The remaining vertices reside in $\operatorname{int}\mathbf{S} \setminus \mathbb{M}_\circ$, and each of these either has a unique outgoing or unique incoming edge.
- Edges have their endpoints at vertices but otherwise do not intersect $\mathbb{M}$ nor $\partial\mathbf{S}$. Any self-crossings in T consist of two edges crossing transversely and there are finitely many of these.
- Every edge $e \in E(T)$ carries an integer label $\mathrm{wt}(e) \in [k]$. When $\mathcal{M} = \mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$ we require that edges e incident to boundary points satisfy $\mathrm{wt}(e) = 1$.
- The local picture surrounding any vertex which is not a marked point is one of the following four pictures, each of which depicts such a vertex and the oriented and labeled half-edges incident to it:

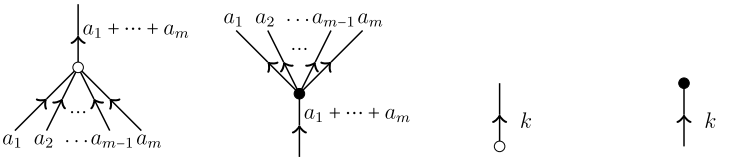

(27)

We consider such tensor diagrams T up to $\mathbb{S}$-isotopy fixing the marked points.

We use the adjectives *plain* or *underlying* to refer to the diagrams just defined when contrasting them with the *tagged* and *pseudotagged* diagrams to be defined below. The extra words wedge product, etc, appearing in (27) are explained in the next section.

We use the notations $s(e)$ and $t(e)$ for the source and target of an oriented edge e in a tensor diagram. We color vertices white or black according to whether they have a unique outgoing or incoming edge. Every coevaluation leaf has an outgoing edge whose endpoint is a black vertex and every evaluation leaf has an incoming edge whose source is a white vertex. We will depict marked points as black vertices (and as a black-white

vertex pair when $k = 3$); these vertices have no incoming edges and may have any number of outgoing edges. We refer to the first two types of vertices in (27) merely as *interior vertices*, using always the special word "leaf" for the third and fourth types of vertices. At interior vertices, the sum of incoming labels equals the sum of outgoing labels, but this property fails at leaves. This latter failure is a combinatorial encoding of the identification $\bigwedge^k(V) \cong \mathbb{C}$ determined by the chosen volume form ξ, as we explain in more detail in the next section.

In examples, we omit the integer label for edges e satisfying $\mathrm{wt}(e) = 1$. We depict evaluation leaves and coevaluation leaves, with their incident edge, as tiny "hairs" emanating from an interior vertex. See (33) for an example of a hair which is drawn as a standin for an evaluation leaf. These hairs are called tags in [3] but we avoid this technology as it clashes with the meaning of tagging arcs in the sense of Fomin, Shapiro, and Thurston.

In our tensor diagrammatical formulas below (e.g. in (34)), edges labeled by 0 should be deleted and any tensor diagram with at least one label in $\mathbb{Z} \setminus [0, k]$ should be discarded or treated as zero in our formulas. An edge of weight k can be chopped in thirds with the middle third discarded, creating an evaluation leaf at the source of the edge and a coevaluation leaf at the target of the edge.

Definition 8.2. The *legs* of a tensor diagram T are its edges incident to punctures. We denote these by $\mathrm{Legs}(T)$ and denote by $\mathrm{Legs}_p(T)$ those legs incident to a given puncture p. We define

$$\mathrm{wt}_p(T) := \sum_{e \in \mathrm{Legs}_p(T)} \omega_{\mathrm{wt}(e)} \in P, \tag{28}$$

and define similarly $\mathrm{wt}(T) \in P^{\mathbb{M}_\circ}$.

We give no special name to edges emanating from boundary points.

The weight vector $\lambda := \mathrm{wt}_p(T)$ always has weakly decreasing coordinates. We can translate it in the $(1, \ldots, 1)$ direction until $\lambda_k = 0$. Once this is done, we have the following useful translation between legs at p and the weight of $[T]$ at p:

$$(\text{number of legs of weight } a \text{ of } T \text{ at } p) = \lambda_a - \lambda_{a+1}. \tag{29}$$

Example 8.3. An arc $\gamma\colon [0, 1] \to \mathbb{S}$ with endpoints $p, q \in \mathbb{M}$ determines an SL_2 tensor diagram $T(\gamma)$ described as follows. The diagram has an oriented edge $\gamma(0) \to \gamma(\frac{1}{2})$ and an edge $\gamma(1) \to \gamma(\frac{1}{2})$, both traveling along γ. These edges meet up at a white vertex at which has an outgoing edge of weight 2 ending at an evaluation leaf. There are two directions in which this outgoing edge can point, but these only change our recipes below by a sign. We suppress this sign throughout.

In a similar way, any simple closed curve α gives rise to a tensor diagram $T(\alpha)$ consisting of a coevaluation leaf edge whose target is a black vertex at $\alpha(0)$. This black vertex has two outgoing edges, one oriented along α from time 0 to time $\frac{1}{2}$ and the other

from time 1 to time $\frac{1}{2}$. These two oriented edges meet at a white vertex at $\alpha(\frac{1}{2})$ whose outgoing edge is an evaluation leaf edge.

8.2. Plain tensor diagrams as functions

The four types of interior vertices in tensor diagrams – namely white and black interior vertices and the two types of leaves – encode certain $\mathrm{SL}(V)$-equivariant morphisms between tensor products of exterior powers of V. Tensor diagrams are a graphical way of encoding compositions of these maps.

We elaborate on this point now. White interior vertices encode the exterior product morphism $\wedge\colon \bigotimes_{j=1}^{m} \bigwedge^{a_j} V \to \bigwedge^{\sum_j a_j} V$. An evaluation leaf encodes the isomorphism $\bigwedge^k \mathbb{C}^k \xrightarrow{\cong} \mathbb{C}$ and a coevaluation leaf encodes its inverse $\mathbb{C} \xrightarrow{\cong} \bigwedge^k \mathbb{C}^k$. Black vertices encode the dual exterior product morphism

$$\bigwedge^{\sum_j a_j} V \to \bigwedge^{a_1} V \otimes \cdots \otimes \bigwedge^{a_m} V \tag{30}$$

$$x_1 \wedge \cdots \wedge x_b \mapsto \sum \pm (x_{i_1} \wedge x_{i_2} \wedge \cdots \wedge x_{i_{a_m}}) \otimes \cdots \otimes (x_{i_{b-a_1+1}} \wedge x_{i_{b-a_1+2}} \wedge \cdots \wedge x_{i_b}) \tag{31}$$

where the summation is over permutations $(i_1, i_2, \ldots, i_b)$ of $b := a_1 + \cdots + a_m$ which are increasing in each block, i.e. which satisfy $i_1 < \cdots < i_{a_1}$, $i_{a_1+1} < \cdots < i_{a_1+a_2}$, etc. The sign $\pm$ is the sign of the permutation $(i_1, i_2, \ldots, i_b)$, multiplied by the "global sign" of the permutation $(b - a_1 + 1, \ldots, b, \ldots, 1, \ldots, a_m)$. Thus, there is "no sign" associated with shuffling the first a_m vectors to the last tensor factor, shuffling the next a_{m-1} vectors to the penultimate factor, and so on.

Given a tensor diagram T, we have sets of source and sink edges

$$\vec{E}(T) = \{e\colon s(e) \in \mathbb{M} \text{ or is a coevaluation leaf}\} \text{ and}$$
$$\overleftarrow{E}(T) = \{e\colon t(e) \text{ is an evaluation leaf}\}.$$

They determine spaces

$$\vec{W}(T) := \otimes_{e \in \vec{E}(T)} \bigwedge^{\mathrm{wt}(e)} V_{s(e)} \text{ and } \overleftarrow{W}(T) := \otimes_{e \in \overleftarrow{E}(T)} \bigwedge^{\mathrm{wt}(e)} V_{t(e)}.$$

The second of these spaces has a canonical isomorphism with $\mathbb{C}$ given by the tensor product of the evaluation maps.

A choice of decorated local system $z \in \mathcal{M}$ provides us with a tensor $v_e(z) \in \bigwedge^{\mathrm{wt}(e)} V_{s(e)}$ for each $e \in \vec{E}(T)$: one should take the volume form χ at each coevaluation leaf and take the decoration (either the vector assigned to boundary points or the

wt(e)th step of the chosen affine flag) which the chosen z assigns at marked points. These give rise to a tensor $v_{initial}(z) := \otimes_{e \in \vec{E}(T)} v_e(z) \in \vec{W}(T)$.

The local system also provides us with parallel transport isomorphisms $\Xi_e \colon V_{s(e)} \to V_{t(e)}$, hence an isomorphism $\bigwedge^{\mathrm{wt}(e)} V_{s(e)} \to \bigwedge^{\mathrm{wt}(e)} V_{t(e)}$. Combining these isomorphisms with the exterior product, dual exterior product maps, and evaluation maps, the chosen z determines a map $\vec{W}(T) \to \mathbb{C}$.

Definition 8.4. The invariant $[T]$ encoded by an $\mathcal{M}$-tensor diagram T is the function on $\mathcal{M}$ whose value on a point $z \in \mathcal{M}$ is the image of the vector $v_{\mathrm{initial}}(z) \in \vec{W}(T)$ under the morphism $\vec{W}(T) \to \mathbb{C}$ encoded by z.

It can be seen that any invariant $[T]$ is a regular function on $\mathcal{M}$. The word invariant is used because regular functions on the moduli space are the G-invariant regular functions on the space of decorated local systems.

The weight of a tensor diagram, defined combinatorially in (28), matches the weight of its corresponding invariant $[T]$ with respect to the right T-action at punctures (3). In particular, if we know the weight of the invariant $[T]$ then we can recover how many legs of each weight the tensor diagram T has at each puncture using the translation (29).

In our algebraic calculations in Section 9, we typically forego the dual exterior product symbol and work with a related algebraic construct. Namely, combining the dual exterior product map with an evaluation map, we have a map $\bigwedge^{k-a} V \otimes \bigwedge^{k-b} V \to \bigwedge^{k-a-b} V \otimes \bigwedge^{b} V \otimes \bigwedge^{k-b} V \to \bigwedge^{k-a-b} V \otimes \bigwedge^{k} V \cong \bigwedge^{k-a-b} V$. We denote this composition by

$$\cap \colon \bigwedge^{k-a} V \otimes \bigwedge^{k-b} V \to \bigwedge^{k-a-b} V. \tag{32}$$

One computes $\cap$ by shuffling b vectors from the first tensor factor over to the second factor (for a total of b plus $k-b$ vectors), taking the determinant of the k vectors in the second factor after such shuffle, and using it to rescale the $k-a-b$ leftover vectors in the first factor. The value of $\cap$ is the signed sum of such terms. The operation $\cap$ is associative and is commutative up to a predictable sign; see [33, Section 3.3] for an exposition.

Remark 8.5. One can make Definition 8.4 more explicit: the value of $[T]$ on $z \in \mathcal{M}$ can be computed as a signed sum of complex numbers. Each term in this sum corresponds to a process of "flowing" the decorations along edges using the parallel transport isomorphisms, applying the exterior product map at white vertices and splitting up the incoming tensor at black vertices using the right hand side of (30). The terms in this sum correspond to a choice of term from the right hand side of (30) at each black vertex. The sign of a term is its product of signs from the various terms at the black vertices, and the complex number associated with a term is the product of the evaluation maps

at the evaluation leaves. For more details on this perspective when $k = 3$ see [8, Section 4].

Example 8.6. As a simplest example of the invariant by a tensor diagram, consider the case that $\mathbb{S} = D_{n,0}$ is an n-gon with boundary points $p_1, \ldots, p_n$ labeled in counterclockwise order. Consider the moduli space $\mathcal{M} = \mathcal{A}'_{\mathrm{SL}_k,\mathbb{S}}$. Any two paths joining a pair of points in $p, q \in \mathbb{S}$ are isotopic, so the fibers V_p, V_q of the local system are canonically identified by the transport isomorphisms. A choice of $z \in \mathcal{M}$ determines a vector $v_i \in P$ at the boundary point p_i and the vector v_{initial} is the tensor product of these vectors. Consider the invariant $[T]$ of the following diagram:

$T =$ 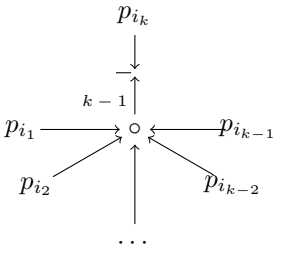

(33)

Applying the parallel transport isomorphisms we get a copy of $v_{i_1} \wedge \cdots \wedge v_{i_{k-1}}$ exiting the white vertex, and we get $v_{i_1} \wedge \cdots \wedge v_{i_{k-1}} \otimes v_{i_k}$ entering the "hair" (which is a visual shortcut for an evaluation leaf, see above). When we exit the hair, we get the tensor $v_{i_1} \wedge \cdots \wedge v_{i_k}$. Applying the evaluation map, we get the number $\det(v_{i_1} \wedge \cdots \wedge v_{i_k}) \in \mathbb{C}$.

When a similar tensor diagram is drawn on a surface, one transports the tensors to the evaluation leaf using the parallel transport isomorphisms and then takes a similar determinant.

Definition 8.7. A *web* is an embedded (i.e. planar) tensor diagram. A *forest diagram* is a tensor diagram whose underlying undirected graph has no cycles on interior vertices. A *tree diagram* is a forest diagram which is not a superposition of two forest diagrams.

A function on the moduli space $\mathcal{M}$ is a *diagram invariant* if it can be expressed as $[T]$ for some tensor diagram T. One has similarly *web invariants*, *forest invariants*, and *tree invariants*.

The product of diagram invariants $[T] \cdot [T']$ is again a diagram invariant encoded by the superposition of diagrams $T \cup T'$.

Definition 8.8. The *skein algebra* $\mathrm{Sk}(\mathcal{M})$ is the algebra generated by $\mathcal{M}$-tensor diagram invariants, thought of as a subalgebra of the algebra of regular functions on $\mathcal{M}$.

We stress that we do not define the skein algebra as the algebra generated by certain diagrams and modulo certain linear relations, which is the typical usage of this terminology. We discuss partial steps towards such a definition in the next section.

Remark 8.9. We expect that in many cases, the skein algebra $\mathrm{Sk}(\mathcal{M})$ coincides with the algebra of regular functions on $\mathcal{M}$, but we do not have a sense of precisely when this holds.

Example 8.10. Recall the tensor diagram $T(\gamma)$ associated to an arc γ as in Example 8.3. As in Example 8.6, the invariant $[T(\gamma)] \in \mathrm{Sk}(\mathrm{SL}_2, \mathbb{S})$ evaluates on a point in the moduli space by first transporting the chosen vectors at the endpoints of γ to the midpoint of γ and then taking their determinant. The two possible signs for such a determinant correspond to the two possible directions in which the evaluation edge can point. As we have said, we gloss over this sign in what follows.

The invariant $[T(\alpha)] \in \mathrm{Sk}(\mathrm{SL}_2, \mathbb{S})$ encoded by a simple closed curve α is the trace of the monodromy matrix M_α assigned to α. Indeed, the volume form $\xi = e_1 \wedge e_2$ at the coevaluation stub is sent to $e_1 \otimes e_2 - e_2 \otimes e_1$ by the coevaluation map. The two ways of proceeding from the coevaluation leaf to the evaluation leaf differ by M_α, so that the parallel transport of the above tensor is $e_1 \otimes M_\alpha(e_2) - e_2 \otimes M_\alpha(e_1)$. Applying the exterior product map we get

$$[T(\alpha)] = e_1 \wedge M_\alpha(e_2) - e_2 \wedge M_\alpha(e_1) = \mathrm{tr}(M_\alpha) \in \mathbb{C}.$$

8.3. Skein relations between diagram invariants

There are interesting linear relations between diagram invariants known as *skein relations*. Each skein relation $\sum a_i[T_i] = 0$ is local in the sense that one can choose an open neighborhood $U \subset \mathrm{int}\,\mathbf{S}$ such that the involved diagrams T_i coincide on $\mathbf{S} \setminus U$.

As an example, whenever a tensor diagram T has a crossing, one can apply the following *crossing removal relation* [3, Corollary 6.2.3]:

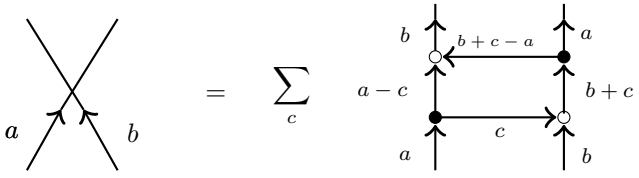

(34)

expressing a diagram invariant which has a crossing as a sum of diagram invariants in which this crossing has been removed. It follows that web invariants span the algebra $\mathrm{Sk}(\mathcal{M})$. The Ptolemy relations familiar in the theory of cluster algebras from surfaces are the $k = 2$ instance of this relation.

Although web invariants span $\mathrm{Sk}(\mathcal{M})$, they are not linearly independent. The most mysterious linear relation between them is the *square switch relation* [3, (2.10)]:

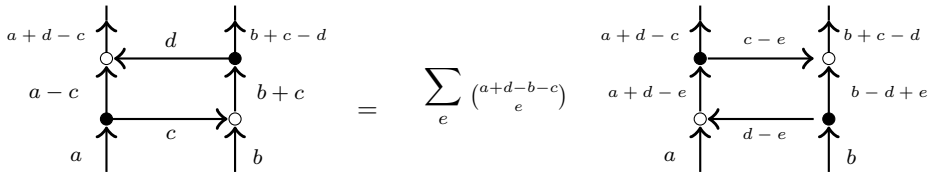

. (35)

Besides this square switch relation, there are additional skein relations which reflect "obvious" properties of the underlying exterior algebra gadgets. Among these are the *leaf migration* relations [3, (2.3), (2.7),(2.8)] and also the associativity relation [3, (2.6)]. The former relations capture signs associated with permuting vectors in an exterior product and the latter relation expresses associativity of exterior product. The composition of a black vertex (splitting up a tensor into several factors) with a white vertex (wedging these vectors back together) is a scalar multiple of the identity map and this is encoded by another skein relation [3, (2.4)]. There are also the duals of all of these relations.

The above named linear relations are the *only* relations between diagram invariants provided we make two simplifying assumptions [3]. First, we should work only with the n -gon $\mathbb{S} = D_{n,0}$ and second, we should only consider tensor diagrams with at most one edge at each marked point. In our setting, marked points can have arbitrary degree and further diagrammatic relations are needed. For example, any diagram with a 2-cycle based at a marked point will define the zero invariant. As a separate issue, we do not think that it has been proved that the relations in [3] are the only linear relations between diagrams when $\mathbb{S}$ is not an n-gon.

8.4. Tagged and pseudotagged diagrams

We now modify the calculus of tensor diagram invariants in the presence of punctures, generalizing the tagged arc calculus valid when $k = 2$.

Recall that a leg is an edge of a tensor diagram which is incident to a puncture. Denote by 2^k the power set of $[k]$.

Definition 8.11. A *pseudotagged* tensor diagram is a pair (T, φ) consisting of a tensor diagram T and a *tagging function*

$$\varphi\colon \mathrm{Legs}(T) \to 2^{[k]} \qquad \text{subject to} \qquad |\varphi(e)| = \mathrm{wt}(e) \text{ for all } e \in \mathrm{Legs}(T). \tag{36}$$

A *tagged tensor diagram* is a pseudotagged tensor diagram (T, φ) which satisfies the following additional condition: for any puncture p and any pair of legs $e, e' \in \mathrm{Legs}_p(T)$, one has a containment of subsets either $\varphi(e) \subseteq \varphi(e')$ or $\varphi(e') \subseteq \varphi(e)$.

We depict a pseudotagged tensor diagram by labeling each leg of T by the subset $\varphi(e)$ rather than by its integer label $\mathrm{wt}(e) \in [k]$. Edges which are not legs continue to carry their integer labels.

The construction of invariants for pseudotagged tensor diagrams is obtained by modifying Definition 8.4 only in the choice of vector $v_e(z)$ associated to edges e whose source is a puncture. For such a leg $e \in \mathrm{Legs}_p(T)$ choose an arbitrary permutation w_e satisfying $w_e([\mathrm{wt}(e)]) = \varphi(e)$. Acting by w_e on the affine flag F_p decoration which z assigns to p, we get a tensor $(w_e \cdot F_p)_{(\mathrm{wt}(e))}$. Then $v_{\mathrm{initial}}(\varphi, z)$ is the tensor product of these vectors over all $e \in \mathrm{Legs}(T)$, together with the vectors assigned to edges at boundary points and those assigned to coevaluation leafs. The choice of permutation w_e in this recipe is immaterial thanks to Lemma 4.3.

Definition 8.12. The invariant $[(T, \varphi)]$ encoded by a tagged tensor diagram (T, φ) is the rational function whose value on a point $z \in \mathcal{M}$ is the image of the vector $v_{\mathrm{initial}}(\varphi, z) \in \overrightarrow{W}(T)$ under the morphism $\overrightarrow{W}(T) \to \mathbb{C}$ encoded by z.

The weight of a pseudotagged diagram invariant is given by $\mathrm{wt}_p([(T, \varphi)]) = \sum_{e \in \mathrm{Legs}_p(T)} \iota_{\varphi(e)}$, a sum of indicator vectors.

We use *pseudotagged* and *tagged* as adjectives modifying the nouns diagram invariant, web invariant, etc. For example, a *tagged web invariant* is a function on $\mathcal{M}$ which can be encoded as $[(T, \varphi)]$ with T a planar tensor diagram and φ a tagging (not merely a pseudotagging) of T.

Definition 8.13. We denote by $\mathrm{Sk}^{\bowtie}(\mathrm{SL}_k, \mathbb{S})$ the algebra generated by pseudotagged invariants, thought of as a subalgebra of the field of rational functions on $\mathcal{M}$.

Since the superposition of pseudotagged diagrams is again pseudotagged, this algebra is spanned by pseudotagged diagram invariants.

Remark 8.14. Tagged tensor diagram invariants are exactly the pullbacks of diagram invariants along the Weyl group action at punctures. Indeed, if (T, φ) is a tagged tensor diagram, then the containment condition ensures that one can choose a permutation $w_p \in W$ for each puncture (as opposed to a permutation w_e for each leg e) with the property that $\varphi(e) = w_p([\mathrm{wt}(e)])$ for all legs e at p. One has then that $[(T, \varphi)] = \prod_p \psi(w_p)^*([T])$. And conversely, the pullback of any diagram invariant $[T]$ along $\prod_p \psi(w_p)$ is the invariant of a tagged diagram (T, φ) whose tagging function is $\varphi(e) := w_p([\mathrm{wt}(e)])$ whenever e is a leg at p.

By the preceding remark, we have the equality of algebras $\mathrm{Sk}^{\bowtie}(\mathrm{SL}_k, \mathbb{S}) = W^{\mathbb{M}_\circ} \cdot \mathrm{Sk}(\mathrm{SL}_k, \mathbb{S})$, i.e. the former algebra is the orbit of the skein algebra under the birational Weyl group action at punctures.

Remark 8.15. If a rational function f is the invariant of a tagged diagram (T, φ), then using the translation (29) we can deduce how many legs (T, φ) has at each puncture and how each of these is tagged. Such a deduction is not possible in general within the class

of pseudotagged diagram invariants. For example, we cannot deduce from the knowledge $\mathrm{wt}_p(f) = e_1 + e_2$ whether we have two legs tagged by the subsets $\{1\}$ and $\{2\}$ or one leg tagged by $\{1, 2\}$. In a tagged diagram, only the second of these is possible.

Remark 8.16. Let $v_1, \ldots, v_k$ be a tuple of vectors representing an affine flag F_p assigned to a puncture p by a decorated local system $z \in \mathcal{M}$. We stress that, for example, a leg $e \in \mathrm{Legs}_p(T)$ appearing in a pseudotagged diagram (T, φ) and with tagging $\varphi(e) = \{2\}$ does not represent the "second vector" v_2, but rather the vector $u_p(v_2) - v_2$. In a similar vein, one should not confuse the following two pseudotagged tensor diagram fragments:

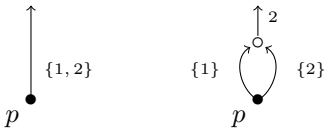

(37)

The left fragment represents $v_1 \wedge v_2$ while the right one represents $v_1 \wedge (u(v_2) - v_2) = 0$. Any (T, φ) containing the right fragment has $[(T, \varphi)] = 0$.

Example 8.17. We explain in detail that when $k = 2$, the choice of tagging a leg by the subset $\{1\}$ or $\{2\}$ is exactly the choice of plain versus notched tagging of e.

We use the notation of [10, Figure 14] for arcs in a once-punctured digon with boundary points p', q and with puncture p. We consider an arc γ connecting p to p' in this digon and denote by $\gamma^{\bowtie}$ the result of notching the end of γ at p. We denote by α and β be the arcs forming the two sides of the digon, by θ the arc connecting q to p in the digon, and by η the loop based p' and enclosing p as a once-punctured monogon. Each of the plain arcs γ, α, β, and η determines an element of the skein algebra which we denote by $[T(\gamma)]$ etc. as in Example 8.10.

We denote by $x(\gamma^{\bowtie})$ the cluster variable indexed by $\gamma^{\bowtie}$. It is defined by the exchange relation $[T(\theta)]x(\gamma^{\bowtie}) = [T(\alpha)] + [T(\beta)]$. On the other hand, we have a skein relation $[T(\theta)][T(\eta)] = [T(\gamma)]\,([T(\alpha)] + [T(\beta)]) \in \mathrm{Sk}(\mathrm{SL}_2, \mathbb{S})$. It follows that

$$x(\gamma^{\bowtie}) = \frac{[T(\eta)]}{[T(\gamma)]}. \tag{38}$$

We need to reconcile this expression for $x(\gamma^{\bowtie})$ with the result of pulling back along s_1 action at p. We have vector decorations v_p at p and $v_{p'}$ at p' respectively. In the notation of Example 8.10, $\psi^*_{s_1,p}([T(\gamma)])$ is the function $\det\begin{pmatrix} v_{p'} & u_p(w) - w \end{pmatrix}$, where u_p is the monodromy around p and w is *any* choice of vector with the property that $\det(v_p w) = 1$. Generically, we can take $w = \frac{v_{p'}}{\det(v_p v_{p'})}$. Thus

$$\psi^*_{s_1,p}([T(\gamma)]) = \det\left(v_{p'} \quad \tfrac{1}{\det(v_p v_{p'})}\,(u_p(v_{p'}) - v_{p'})\right)$$

$$= \frac{\det\begin{pmatrix} v_{p'} & u_p(v_{p'}) \end{pmatrix}}{\det\begin{pmatrix} v_p & v_{p'} \end{pmatrix}} = \frac{[T(\eta)]}{[T(\gamma)]} = x(\gamma^{\bowtie})$$

as claimed.

8.5. Cluster compatibility conjectures

Let $\mathcal{M}$ be one of the two versions of moduli space considered in this paper. We now state our two main conjectures phrasing cluster compatibility for $\mathscr{A}(\mathcal{M})$ in terms of tagged tensor diagram calculus.

Our first conjecture extends the philosophy of Fomin and the second author [8,7] in two directions: to the case of surfaces $\mathbb{S}$ with punctures, and also to the setting $k > 3$.

Conjecture 8.18. Every cluster monomial $f \in \mathscr{A}(\mathcal{M})$ is a tagged web invariant and also a tagged forest invariant. Every cluster variable f is moreover a tagged tree invariant.

Remark 8.19. It would be desirable to formulate a move on tagged diagrams which does not change the tagged invariant and with the property that a planar diagram represents a cluster variable if and only if it can be turned into a tree diagram by repeated applications of this move. This would be a higher rank analogue of the arborization move introduced when $k = 3$ in [7]. We found generalizations of this arborization move in higher rank, but we were not confident that the generalization we found has the desired properties. Certainly, this deserves further study.

Our second conjecture extends the $k = 2$ philosophy of Fomin, Shapiro, and Thurston to the $k > 2$ setting:

Conjecture 8.20. Let $T_1, \ldots, T_s$ be tensor diagrams with each $[T_i] \in \mathscr{A}(\mathcal{M})$ a cluster variable. Let φ_i be a tagging of T_i with corresponding invariant $f_i = [(T_i, \varphi_i)]$. Assume for simplicity that all the f_i's are distinct. Then $\prod_i f_i \in \mathscr{A}(\mathcal{M})$ is a cluster monomial if and only if the following three conditions hold.

First, the underlying plain monomial $\prod_i [T_i] = [\cup_i T_i] \in \mathscr{A}(\mathcal{M})$ is a cluster monomial. Second, if $[T_i] \neq [T_j]$, then $(T_i, \varphi_i) \cup (T_j, \varphi_j)$ is a tagged tensor diagram. And third, if $[T_i] = [T_j]$, then the weight vectors $\mathrm{wt}(f_i)$ and $\mathrm{wt}(f_j)$ are root-conjugate at some puncture p and coincide at all other punctures p'.

Note that Conjecture 8.18 is only a necessary condition on cluster monomials, while Conjecture 8.20 claims to be both necessary and sufficient.

Translating from the language of weight vectors to the combinatorics of legs as in (29), we can restate Conjecture 8.20 in a way that is more obviously consonant with [9]. Recall that we depict a tagged tensor diagram (T, φ) as a plain diagram each of whose legs e is decorated by a subset $\varphi(e)$. Then the second condition in Conjecture 8.20 says that if $[T_i] \neq [T_j]$, then either $\varphi_i(e) \subset \varphi_j(e')$ or $\varphi_j(e') \subset \varphi_i(e)$ for any pair of legs

$e, e' \in \mathrm{Legs}_p(T)$ and any puncture p, which generalizes the requirement that distinct arcs must be tagged the same way at punctures. The third condition says that if $[T_i] = [T_j]$, then the taggings φ_i and φ_j disagree in exactly one leg just as in the $k = 2$ case.

Lemma 8.21. *If a cluster $\{f_i\}$ satisfies the three conditions in Conjecture 8.20, then its underlying P-cluster $\{\pi_p(\mathrm{wt} f_i)\} \subset P$ at any puncture p is a basic multiset of pairwise compatible weight vectors.*

Thus, if the conditions stated in Conjecture 8.20 are indeed necessary conditions, and if every cluster variable is indeed a tagged invariant, then every cluster is in the good part of the exchange graph, and the dosp mutation results of Section 7 apply to the exchange graph itself.

Proof. The second and third conditions in Conjecture 8.20 together imply that the P-cluster at any puncture consists of pairwise compatible vectors. We must also show that the P-cluster at each puncture is basic. Suppose that f_i and f_j have root-conjugate weights by the transposition $t = (ab)$ at some puncture p. Then from the definition of evaluation of tagged tensor diagrams, the corresponding functions are related by birational Weyl group action: $f_j = \psi(t)^* f_i$ with t acting at the puncture p. Since we can compute f_i uniquely from f_j, it follows that f_j must be the unique cluster variable with this weight vector. □

Remark 8.22. Suppose in Conjecture 8.20 that the collection $\{f_i\}$ has the property that the vectors $\mathrm{wt}_p(f_i)$ and $\mathrm{wt}_p(f_j)$ are W-sortable for every puncture p and every pair i, j. (That is, the second condition in Conjecture 8.20 holds.) Then the third condition plays no role, and the first condition is clearly both necessary and sufficient for $\prod f_i$ to be a cluster monomial: the monomials $\prod_i f_i$ and $\prod_i [T_i]$ are related by pullback along the $W^{\mathbb{M}^\circ}$ action, so the two products are cluster monomials simultaneously. To emphasize: the "interesting" behavior in Conjecture 8.20 arises when at least one pair of weights at some puncture p are root-conjugate, provided we believe Conjecture 5.4 holds. We did not see how to prove either the necessity or sufficiency of the third condition in this case.

9. The flattening and spiral theorems

We state and prove the flattening theorem and the spiral theorem. Both theorems concern the relationship between pseudotagged diagram invariants and tagged diagram invariants. The flattening theorem says that every pseudotagged invariant is a linear combination of tagged invariants. The flattening relation which underlies the proof is a useful algebraic tool for computations in the cluster algebra $\mathscr{A}(\mathcal{M})$. For example, certain exchange relations in $\mathscr{A}(\mathcal{M})$ are instances of the flattening theorem. The spiral theorem establishes a setting in which certain specific pseudotagged invariants are in fact tagged invariants as predicted by our cluster compatibility conjectures. We also discuss more

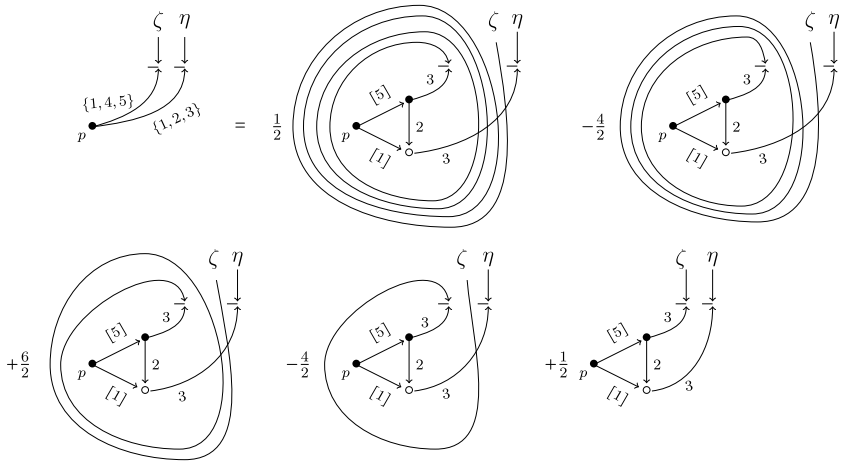

Fig. 8. Expressing a pseudotagged invariant as a $\mathbb{Q}$-linear combination of tagged invariants using Lemma 9.4 with $a = 1$ and $b = c = 2$. The two legs in the leftmost diagram violate the containment condition from Definition 8.11. The tensors $\eta, \zeta \in \bigwedge^{k-3} V$ schematically indicate "the rest" of the diagram. The numerators on the right hand side form the fourth row of Pascal's triangle. The denominator counts SYTs of $b \times c = 2 \times 2$ shape.

refined statements, e.g. whether or not the pseudotagged invariant has a tagged forest form.

9.1. The flattening theorem

Theorem 9.1. *Any pseudotagged diagram invariant is a linear combination of tagged diagram invariants.*

That is, $\mathrm{Sk}^{\bowtie}(\mathcal{M})$ is spanned by tagged diagram invariants. Fig. 8 illustrates Theorem 9.1, flattening a pair of legs tagged by the subsets $\{1, 2, 3\}$ and $\{1, 4, 5\}$ as a linear combination of five tagged diagrams, each which with one leg tagged by [5] and the other by $\{1\}$.

See [29, Section 8.4] for a discussion of similar flattening relations when $k = 2$.

Remark 9.2. We expect, but are not sure how to prove, that $\mathrm{Sk}^{\bowtie}(\mathcal{M})$ is the (upper) cluster algebra $\mathscr{A}^{\mathrm{up}}(\mathcal{M})$. For a similar statement for the "$\mathcal{X}$-space" see [32]. We have not investigated the relationship between ordinary and upper cluster algebras for $\mathscr{A}(\mathcal{M})$. This has been thoroughly explored when $k = 2$, see e.g. [27, Section 3.2] and references therein.

9.2. Proof of the flattening theorem

For the next several paragraphs we fix the following data. First, a choice of vectors $v_1, \ldots, v_k$ determining an affine flag with steps $F_{(a)} := v_1 \wedge \cdots \wedge v_a$. Second, we choose a matrix $u \in \mathrm{stab}(F)$ in the stabilizer of this flag.

Recall that a subset $S \in \binom{[k]}{a}$ determines a Grassmannian permutation $w_S \in W$ and a Young diagram $\lambda(S)$ whose number of standard Young tableaux is denoted $f^{\lambda(S)}$.

We partially order $\binom{[k]}{a}$ via the *termwise order*: an m-subset T is less than or equal to another such S provided, when we list $T = \{t_1 < t_2 < \cdots < t_a\}$ and $S = \{s_1 < s_2 < \cdots < s_a\}$, we have $t_i \leq s_i$ for all i. This partial order is graded by the length function $\ell(S) = \ell(w_S)$.

The above data provides us with a tensor

$$v_S := \bigwedge_{s \in S} v_s \in \bigwedge^a V, \tag{39}$$

with the exterior product taken in increasing order from left to right. We also get a complex number $\mathcal{W}_S$ defined by

$$(w_s \cdot F)_{(a)} = \mathcal{W}_S F_{(a)} \in \bigwedge^a V, \tag{40}$$

where $F \mapsto w_s \cdot F$ denotes Weyl group action using the permutation w_S and the chosen matrix u.

We let $\mathbb{C}\langle G \rangle$ be the group algebra of G. We view $\bigwedge^m(V)$ as a $\mathbb{C}\langle G \rangle$-module via the natural action $G \curvearrowright V$. For $M \in G$ and a complex number c, we distinguish between the group algebra element $c \cdot M \in \mathbb{C}\langle G \rangle$ and the matrix cM in which each entry of M is rescaled by c. This is important because, e.g., $(c \cdot M)(x \wedge y) = cMx \wedge My$ whereas $(cM) \cdot x \wedge y = c^2 Mx \wedge My$ inside $\bigwedge^2(V)$. Our next lemma concerns elements of the form $(u^{-1} - \mathrm{Id}_V)^m \in \mathbb{C}\langle G \rangle$ for $m \in \mathbb{N}$.

Lemma 9.3. *For any S and any $T \leq S$, we have*

$$(u-1)^{\ell(S)} v_T = \begin{cases} f^{\lambda(S)} \mathcal{W}_S F_{(a)} & \text{if } T = S \\ 0 & \text{if } T < S \end{cases} \in \bigwedge^a V. \tag{41}$$

Proof. Define numbers $z_{ji} \in \mathbb{C}$ for $1 \leq j \leq i \leq k$ by $u(v_i) = \sum_j z_{ji} v_j$.

For any $T \in \binom{[k]}{a}$ it is easy to see that $(u-1)v_T$ is a linear combination of tensors $v_{T'}$ satisfying $T' < T$ in termwise order. (The coefficients of such a linear combination are polynomial functions in the z_{ji}'s.) Thus, $(u-1)^{\ell(S)}$ annihilates the tensor v_T whenever $T < S$ as claimed.

It remains to compute $(u-1)^{\ell(S)} v_S$. We use $\lessdot$ to denote cover relations with respect to the order $<$. We expand $(u-1)v_S$ as a linear combination of tensors v_T with $T < S$. By the same reasoning as in the previous paragraph, if T does not satisfy $T \lessdot S$ then v_T is annihilated by $(u-1)^{\ell(S)-1}$, so we can ignore such terms for the rest of the computation. In a cover relation $T \lessdot S$, we have $T = S \setminus \{j+1\} \cup \{j\}$ for some $j+1 \in S$, and the coefficient of v_T in $(u-1)v_S$ is the number $z_{j,j+1}$.

Consider the set of saturated chains

$$E(S) = \{[a] =: S_0 \lessdot S_1 \lessdot \cdots \lessdot S_{\ell(S)} := S\}$$

ending at S in termwise order. The tensor v_{S_i} appears in the tensor $(u-1)v_{S_{i+1}}$ with coefficient z_{j_i,j_i+1} where $S_{i+1} = S_i \setminus (j_i+1) \cup j_i$. Note also that $v_{S_0} = F_{(a)}$. It follows that

$$(u-1)^{\ell(S)} v_S = \left(\sum_{S_0 \lessdot S_1 \lessdot \cdots \lessdot S_{\ell(S)}} \prod_{i=1}^{\ell(s)} z_{j_i,j_i+1} \right) F_{(a)} \tag{42}$$

where j_i are as defined above.

We view the parenthesized scalar on the right hand side as a weighted generating function for elements of $E(S)$, with the weight of such an element defined as the monomial in z_{j_i}'s coming from the steps of the chain. To complete the proof, we will show that this weight function is constant on $E(S)$, and moreover is equal to the number $\mathcal{W}_S$ defined in (40). The right hand side of (42) then becomes $f^{\lambda(S)}\mathcal{W}_S F_{(a)}$, completing the proof.

To argue that $\prod_{i=1}^{\ell(s)} z_{j_i,j_i+1} = \mathcal{W}_S$ is independent of the chain, begin by observing that if $S_i \lessdot S_{i+1}$ is a cover relation in such a chain, then the permutations w_{S_i} and $w_{S_{i+1}}$ are related by left action of a simple transposition swapping the values (j_i, j_i+1). Multiplying these simple transpositions $(s_{j_{\ell(S)}}, \ldots, s_{j_1})$ from right to left, we get the sequence of permutations $w_{S_1}, w_{S_2}, \ldots, w_{S_{\ell(S)}} = w_S$. In particular, $(s_{j_{\ell(S)}}, \ldots, s_{j_1})$ is a reduced word for w_S^{-1} (since reduced words are multiplied from left to right). This determines a bijection between saturated chains ending at S and reduced words for w_S^{-1}. The indices $(j_{\ell(s)}, \ldots, j_1)$ appearing in the reduced word are what is needed to compute the weight $\prod_{i=1}^{\ell(s)} z_{j_i,j_i+1}$ in (42). So we can view the parenthesized scalar factor in (42) as a weight-generating function for reduced words of w_S^{-1}, or equivalently of w_S.

The set of reduced words for w_S is connected by commutation moves. Clearly, commutation moves do not change the weight appearing in (42), so that the weight is in fact independent of the reduced word (or equivalently, of the chain $S_0 \lessdot \cdots \lessdot S_{\ell(S)}$) as claimed.

Finally we prove that the weight of a certain reduced word is $\mathcal{W}_S$, completing the proof. We write $S = \{x_1 < \cdots < x_a\}$. We have the following reduced word for w_S

$$\prod_{j=1}^{a} s_j \cdots s_{x_j - 1} \tag{43}$$

(the product is taken from left to right).

This word transforms the identity permutation to w_S by the following sequence of swaps in adjacent positions. First consider the smallest element x_j such that $x_j > j$. We can swap the number x_j leftwards (past the numbers $j, j+1, \ldots, x_j - 1$) until it is in position j. Then we swap the number x_{j+1} leftwards until it is in position $j+1$, etc.

We can compute the tensor $w_S(F)$, hence the number $\mathcal{W}_S$ using the formula (9). The effect of swapping the number x_j leftwards into position j is to take the jth vector v_j of the affine flag and replace it with the vector $(u-1)^{x_j-j}v_{x_j}$. (The vectors in positions $j+1,\ldots,x_j$ are also rescaled, but this will not affect the rest of the calculation.)

It is easy to see that

$$F_{(j-1)} \wedge (u-1)^{x_j-j}v_{x_j} = z_{j,j+1}z_{j+1,j+2}\cdots z_{x_j-1,x_j}F_{(j)}.$$

In the next step, we swap the number x_{j+1} leftwards to position $j+1$, which corresponds to replacing the $j+1$st step of the affine flag by $(u-1)^{x_{j+1}-j-1}v_{x_{j+1}}$. Again, we have

$$F_{(j)} \wedge (u-1)^{x_{j+1}-j-1}v_{x_{j+1}} = z_{j+1,j+2}\cdots z_{x_{j+1}-1,x_{j+1}}F_{(j+1)}.$$

Continuing in this way, we arrive at

$$\mathcal{W}_S = \prod_{j=1}^{a} z_{j,j+1}\cdots z_{x_j-1,x_j} \tag{44}$$

which is the weight (42) we would associate to the reduced word (43). □

Our next lemma uses the operation $\cap$ defined in (32).

Lemma 9.4. *Consider tensors $\eta \in \bigwedge^{k-a-b}(V)$ and $\zeta \in \bigwedge^{k-a-c}(V)$ where $a,b,c \in \mathbb{N}$ satisfy $a+b+c \leq k$. Let $S = [a] \cup [a+b+1, a+b+c]$. We have the following equality of complex numbers in $\bigwedge^k(V) \cong \mathbb{C}$:*

$$(F_{(a+b)} \wedge \eta)((w_S \cdot F)_{(a+c)} \wedge \zeta) = \frac{1}{f^{\lambda(S)}}\left(F_{(a+b+c)} \cap F_{(a)}\eta \cap (u^{-1}-1)^{\ell(S)}(\zeta)\right). \tag{45}$$

Fig. 8 illustrate this lemma in the language of tensor diagrams. The quantity on the left hand side is the product of two elements of $\bigwedge^k(V)$, each of which is identified with a complex number using the volume form. The parenthesized term on the right hand side can be similarly interpreted.

Proof. We compute the right hand side by shuffling b vectors from the tensor $F_{(a+b+c)} = v_1 \wedge \cdots \wedge v_{a+b+c}$ to the second factor and shuffling the complementary $a+c$ vectors to the third factor, summing over all such shuffles with an appropriate sign. If we shuffle any of the vectors $v_1,\ldots,v_a$ into the first factor then the resulting term vanishes since $v_i \wedge F_{(a)} = 0$. Thus, the terms that are shuffled into the third factor correspond to subsets $T \in \binom{[a+b+c]}{a+c}$ satisfying $[a] \subset T$. The given subset $S = [a] \cup [a+b+1, a+b+c]$ is maximal among these in the termwise partial order. For such a T, we have

$$v_T \wedge (u^{-1}-1)^{\ell(S)}(\zeta) = \sum_i (-1)^i \binom{\ell(S)}{i} v_T \wedge u^{-i}(\zeta)$$

$$
\begin{aligned}
&= \sum_i (-1)^i \binom{\ell(S)}{i} u^i v_T \wedge \zeta \\
&= (u-1)^{\ell(S)} v_T \wedge \zeta \\
&= \delta_{S,T} f^{\lambda(S)} \mathcal{W}_S F_{(a+c)} \wedge \zeta \\
&= \delta_{S,T} f^{\lambda(S)} w_S(F)_{(a+c)} \wedge \zeta.
\end{aligned}
$$

The first and third equalities are the binomial theorem. In the second, we move u from ζ to v_T using $\det u = 1$. The fourth equality is (41) where δ denotes Kronecker delta and the last equality is the definition of $\mathcal{W}_S$ (39). By this calculation, the unique T which can be shuffled to the third factor when computing the right hand side of (45) is $T = S$, leaving $[a+1, b]$ to be shuffled to the second factor (with a sign of $+1$). We get the two determinants appearing on the left hand side after canceling the $f^{\lambda(S)}$. □

Recall the indicator vector $\iota_S = \sum_{s \in S} e_s \in \mathbb{Z}^k$ of a subset $S \subset [k-1]$.

Lemma 9.5. *Let $\mathcal{S}$ be a multiset of indicator vectors ι_S with the property that $\sum_{v \in \mathcal{S}} v$ lies in the dominant Weyl chamber closure $\overline{C_{\mathrm{id}}}$. By a rewriting move, we will mean the transformation*

$$\{\iota_S, \iota_T\} \mapsto \{\iota_{S \cap T}, \iota_{S \cup T}\}$$

removing indicator vectors of sets S and T and replacing them with that of their intersection and union. Then by a finite sequence of such rewriting moves, we can transform the multiset $\mathcal{S}$ to a new multiset $\mathcal{S}'$ consisting entirely of fundamental weights.

Proof. Pick an enumeration $\iota_{S_1}, \iota_{S_2}, \ldots, \iota_{S_m}$ of the vectors in the initial multiset. We claim for $2 \leq i \leq m$ that by a sequence of rewrites one can transform the multiset $\iota_{S_1}, \ldots, \iota_{S_i}$ into the multiset of indicator vectors for the subsets

$$\cup_{T \in \binom{[i]}{j}} \cap_{t \in T} S_t \text{ for } j = 1, \ldots, i. \tag{46}$$

The base case ($i = 2$) is a single rewriting

$$\{\iota_{S_1}, \iota_{S_2}\} \mapsto \{\iota_{S_1 \cup S_2}, \iota_{S_1 \cap S_2}\}, \tag{47}$$

with the first term the $j = 1$ term and the second the $j = 2$ term.

Assuming we can rewrite the first i indicator vectors as in (46), we perform several rewrites to obtain the $i+1$ version of (46). First we rewrite the pair consisting of the $j = i$ term $\cup_{T \in \binom{[i]}{i}} \cap_{t \in T} S_t = S_1 \cap \cdots \cap S_i$ and the new term S_{i+1}. The result of this rewriting is to obtain $S_1 \cap \cdots \cap S_{i+1}$, which is a desired term in the $i+1$ version of (46), and $(S_1 \cap \cdots \cap S_i) \cup S_{i+1}$, which is not. Then we rewrite the undesired term from the previous step with the $i-1$ term in (46). One of the resulting terms is

$$(\cup_{T\in\binom{[i]}{i-1}} \cap_{t\in T} S_t) \cap ((S_1 \cap \cdots \cap S_i) \cup S_{i+1}) = \cup_{T\in\binom{[i+1]}{i}} \cap_{t\in T} S_t,$$

a desired term in the $i+1$ version of (46), and the other term is not. The rest of the proof of the inductive step proceeds similarly using the identity

$$(\cup_{T\in\binom{[i]}{j-1}} \cap_{t\in T} S_t) \cap ((\cup_{T\in\binom{[i]}{j}} \cap_{t\in T} S_t) \cup S_{i+1}) = \cup_{T\in\binom{[i+1]}{j}} \cap_{t\in T} S_t,$$

in each step. The first term of these terms appears in (46) and the second of these is the undesired term created in the previous step using the identity $(\cup_{T\in\binom{[i]}{j+1}} \cap_{t\in T} S_t) \cup (\cup_{T\in\binom{[i]}{j}} \cap_{t\in T} S_T) = \cup_{T\in\binom{[i]}{j}} \cap_{t\in T} S_T$.

Let η be the sum $\iota_{S_1} + \cdots + \iota_{S_m}$ of the indicator vectors of our initial multiset. Since the rewriting rules do not affect the sum of the indicator vectors in the collection, we can evaluate η by summing the indicator vectors in (46) in the case $i = m$. It is clear that the jth term in (46) contains the $j+1$st term, for $j = 1, \ldots, m-1$. Thus, the largest coordinates of η are supported on the $j = m$ term in (46), i.e. on the intersection $S_1 \cap S_2 \cap \cdots S_m$. Since η lies in the dominant Weyl chamber by assumption, we conclude that $S_1 \cap S_2 \cap \cdots S_m$ is an initial interval $[a]$. Continuing in this way, we conclude that each term in (46) is an initial interval, completing the proof. □

Now we can prove Theorem 9.1.

Proof. Let (T, φ) be a pseudotagged tensor diagram and p a puncture. The weight of the function $[(T, \varphi)]$ at p is the sum of indicator vectors $\iota_{\varphi([\mathrm{wt}(e)])}$. Composing with the birational Weyl group action, we may assume that $\mathrm{wt}_p([T])$ lies in the dominant Weyl chamber closure for all punctures p. We want to show that $[(T, \varphi)]$ is a linear combination of diagram invariants.

For any pair of legs e, e' incident to the same puncture, Lemma 9.4 says that we can replace the tagged tensor diagram T by a linear combination of tensor diagrams in which the legs e, e' are replaced by a pair of legs whose weights are indicator vectors $\iota_{\varphi(e)\cap\varphi(e')}$ and $\iota_{\varphi(e)\cup\varphi(e')}$. Moreover, each of the tensor diagrams in this linear combination agrees with T outside of a small neighborhood of the puncture p.

This transformation on legs is exactly the rewriting rule from Lemma 9.5. Thus, after sufficiently many applications of Lemma 9.4 we will be able to express $[(T, \varphi)]$ as a linear combination of tagged tensor diagrams T' whose legs are tagged by initial subsets $\varphi(e) = [\mathrm{wt}(e)]$. That is, each of the tensor diagrams T' is plainly tagged at p. As this recipe only changes the tensor diagrams nearby p, we can continue this reduction process until we are plainly tagged at every puncture. □

9.3. The spiral theorem

Consider an r-tuple of tagged tensor diagrams (T, φ_i) for $i = 1, \ldots, r$, different taggings of the same underlying tensor diagram T. Suppose that the weights of the invariants

$f_i := [(T, \varphi_i)]$ are pairwise root-conjugate at some puncture p and coincide at all other punctures $p' \neq p$. Thus, there is a unique leg $e \in \mathrm{Legs}_p(T)$ such that $\varphi_i(e') = \varphi_j(e')$ for all $i, j \in [r]$ and all $e' \in \mathrm{Legs}(T) \setminus e$. The taggings $\varphi_i(e)$ and $\varphi_j(e)$ do not coincide.

Next, suppose that the underlying plain diagram invariant $[T]$ is a cluster variable. Then Conjecture 8.20 predicts that the product $\prod_i f_i$ is a cluster monomial. We have $\prod_i f_i = [\cup_i (T, \varphi_i)]$ but the diagram on the right hand side is pseudotagged rather than tagged. Conjecture 8.18 predicts that this monomial should be a tagged invariant, moreover a tagged forest and tagged planar invariant. The spiral theorem partially confirms this.

Theorem 9.6. *Let $f_i = [(T, \varphi_i)]$, for $i = 1, \ldots, r$, be taggings which are root-conjugate at exactly one puncture as above and let $e \in T$ be the unique leg at which these tagged diagrams disagree. Suppose that T is a tree diagram. Then the following statements hold:*

- *The monomial $\prod_{i=1}^r f_i$ is a tagged tree invariant.*
- *The monomial $\prod_{i=1}^r f_i$ can be computed by a tagged diagram which is the result of gluing r many copies of the tensor diagram $T \setminus e$ to a tagged and planar diagram fragment.*

We elaborate on the second part of this theorem, which is slightly technical. We view the superposition diagram $\cup_{i=1}^r (T, \varphi_i)$ as r many root-conjugate tagged legs plugged into several copies of the fragment $\eta := T \setminus e$. The second part of this theorem asserts that we can replace these root-conjugate legs by a more complicated tagged tensor diagram fragment which is planar. We think of this result as saying that the corresponding monomial is given by a diagram which is tagged and "planar near p," although this is of course silly since *any* diagram is planar in a very small neighborhood of a puncture.

Remark 9.7. We would be happiest if we could strengthen the second part of this theorem to conclude that $\prod_i f_i$ is a tagged planar invariant by adding an appropriate planarization hypothesis on the underlying plain diagram T. There are two issues here: first, the diagram $T \setminus e$ might have legs at p, and these legs will cross the planar fragment near p appearing in the second part of the theorem statement. Second, the union of several copies of $T \setminus e$ will typically have many self-crossings, and we need to argue that these self-crossing can be planarized. The first of these issues is not hard to handle: a straightforward generalization of the arborization move from [8] is sufficient to planarize all crossings between the copies of $T \setminus e$ and the planar fragment introduced below. The second of these issues is a bit more subtle; we remove it in the case $k = 3$ in Proposition 9.11.

We illustrate Theorem 9.6 in Fig. 9, which schematically depicts three diagrams D_{super}, $D_{\mathrm{maelstorm}}$, and D_{snail} appearing in the upper left, upper right, and bottom

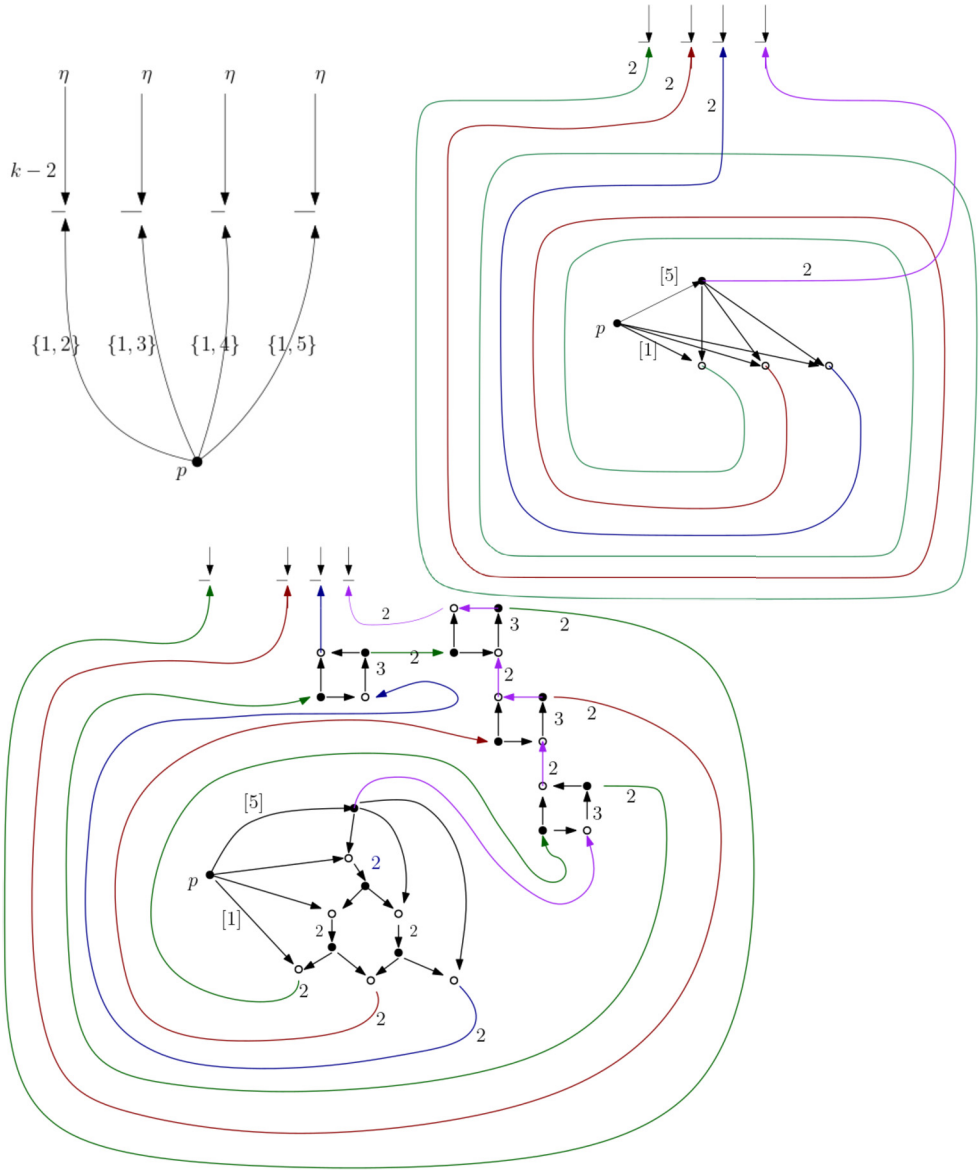

Fig. 9. The superposition, maelstrom, and snail shell fragmentsdiagrams from the proof of the spiral theorem. All three diagrams determine the same invariant when plugged into several copies of the tensor diagram fragment $\eta := T \setminus e$.

of the figure. Several copies of the tensor diagram fragment $T \setminus e$ enter from the top of each of these pictures. The first diagram D_{super} is pseudotagged and depicts the union $\cup_i(T, \varphi_i)$. The latter two diagrams are tagged. We will prove below that $[D_{\text{pseudo}}] = [D_{\text{maelstrom}}] = [D_{\text{snail}}]$.

9.4. Proof of Theorem 9.6

Fix an affine flag F and a matrix $u \in \mathrm{stab}(F)$. The following algebraic identity underlies the passage from D_{super} to $D_{\mathrm{maelstrom}}$ in Fig. 9.

Lemma 9.8. *Let $a, r \in \mathbb{N}$ with $a + r \leq k$ and let $\eta \in \bigwedge^{k-a-1} V$. Then*

$$\prod_{j=a}^{a+r-1} ((s_{a+1}s_{a+2}\cdots s_j \cdot F)_{(a+1)} \wedge \eta) = F_{(a+r)} \cap (F_{(a)} \wedge \eta) \cap (F_{(a)} \wedge u^{-1}\eta) \cap \cdots \cap (F_{(a)} \wedge u^{2-r}\eta) \cap u^{1-r}\eta. \tag{48}$$

And dually, when $\eta \in \bigwedge^{k-a-r+1}(V)$, one has

$$\prod_{j=a+1}^{a+r} ((s_{r-1}\cdots s_{j+1}s_j \cdot F)_{(a+r-1)} \wedge \eta) = F_{(a)} \wedge (F_{(a+r)} \cap \eta) \wedge (F_{(a+r)} \cap u^{-1}\eta) \wedge \cdots \wedge (F_{(a+r)} \cap u^{2-r}\eta) \wedge u^{1-r}\eta. \tag{49}$$

The left hand side of these equalities represents an r-tuple of root-conjugate legs in a pseudotagged superposition of tagged diagrams $\cup_i (T_i, \varphi)$. The ground set of the root-conjugacy class is $[a+1, \ldots, a+r]$. The two versions of the equality correspond to the two choices of sign $\varepsilon = \pm 1$ of the root-conjugacy class.

The right hand side of either equality, when interpreted in terms of tensor diagrams, is computed by a tensor diagram fragment whose legs satisfy the containment condition for tagged diagrams. Thus, these equalities allow us to replace an r-tuple of root-conjugate legs at a puncture with an r-tuple of legs satisfying the containment condition, provided each of these legs is connected to the same tensor diagram fragment $T \setminus e$.

The application of the first equality in this lemma is illustrated in the passage from the upper left to the upper right diagram in Fig. 9, in the case $a = 1$ and $r = 4$.

Proof. We prove (48) by induction on $r \geq 1$; the proof of (49) is very similar. For the base case, the right hand side should be interpreted as $F_{(a+1)} \cap u^{r-1}\eta = F_{(a+1)} \wedge u^0\eta$, so the identity is trivial. We abbreviate

$$\zeta_r = (F_{(a)} \wedge \eta) \cap (F_{(a)} \wedge u^{-1}\eta) \cap \cdots \cap (F_{(a)} \wedge u^{2-r}\eta) \in \bigwedge^{k-r+1} (V)$$

so that the right hand side of (48) takes the form $F_{(a+r)} \wedge (\zeta_r \cap u^{1-r}\eta)$. It is easy to argue that $F_{(a)} \wedge (\zeta_r \cap u^{1-r}\eta) = \zeta_{r+1}$.

Assuming the formula for a given $r \geq 1$, we have

$$\prod_{j=a}^{a+r} (s_{a+1}\cdots s_j \cdot F)_{(a+1)} \wedge \eta = \left(F_{(a+r)} \wedge (\zeta_r \cap u^{1-r}\eta)\right) \left((s_{a+1}\cdots s_{a+r} \cdot F)_{(a+1)} \wedge \eta\right)$$

$$
\begin{aligned}
&= F_{(a+r+1)} \cap (F_{(a)} \wedge (\zeta_r \cap u^{1-r}\eta)) \cap (u^{-1}-1)^r \eta \\
&= F_{(a+r+1)} \cap \zeta_{r+1} \cap (u^{-1}-1)^r \eta \\
&= F_{(a+r+1)} \cap \zeta_{r+1} \cap u^{-r}\eta
\end{aligned}
$$

establishing the inductive claim. The first of these equalities is the inductive hypothesis, the second is Lemma 9.4 when $S = [a] \cup \{r+1\}$, the third is the relation between ζ_r and ζ_{r+1} we noted above, and the fourth can be argued as follows. We claim that if $i < r$, then $\zeta_{r+1} \cap u^{-i}\eta = 0$. Indeed, by reordering terms in the definition of ζ_{r+1} we can compute this quantity by first performing $F_{(a)}u^{-i}\eta \cap u^{-i}\eta$ (since the first of these factors appears in ζ_{r+1}). To compute $F_{(a)}u^{-i}\eta \cap u^{-i}\eta$ we need to shuffle $a+1$ vectors from the first factor over to the second factor, hence we must shuffle at least one vector from $u^{-i}\eta$ over, hence $F_{(a)}u^{-i}\eta \cap u^{-i}\eta = 0$. This completes the proof of the fourth equality. □

Next we state a lemma which underlies the passage from the $D_{\text{maelstrom}}$ to D_{snail} in Fig. 9.

A type $+1$ *maelstrom fragment* is a fragment of a plain tensor diagram which generalizes the upper right diagram in the figure and computes the right hand side of (48). It has a leg e_0 of weight $a+r$ and legs $e_1, \ldots, e_r$ of weight a at the puncture p. It has edges $e_0 \to t(e_i)$ for $i \in [r]$, and dangling edges e_i' of weight $a+1$, with e_i' emanating from $t(e_i)$ and swirling around the puncture $i-1$ many times. Each of these dangling edges is then paired with a copy of the fragment $T \setminus e$ to create a tensor diagram whose invariant we call the *maelstrom invariant*. There is a similar type -1 maelstrom fragment computing the right hand side of (49).

The maelstrom fragment has two types of crossings: crossing between legs e_i and e_j and crossings between the dangling edges e_i' and e_j' swirling around the puncture. The former type of crossings can be planarized using the crossing removal relation; all but one of the terms on the right hand side of such a crossing relation vanish. A *snail shell fragment* is the result of removing these crossings and also replacing each crossing between edges e_i' and e_j' by a square fragment as indicated by example in the passage from the right diagram to the bottom diagram in the figure. As above, we pair the dangling edges of the snail shell fragment with several copies of the fragment $T \setminus e$ to get a tensor diagram and corresponding *snail shell invariant*.

Lemma 9.9. *The snail shell invariant coincides with its corresponding maelstrom invariant.*

That is, the snail shell invariant also computes the right hand side of (48) or (49).

Proof. We first argue that one can replace each of the crossings between dangling edges e_i' and e_j' by a square. One can subsequently planarize the crossings between legs using the inverse arborization moves, as we have claimed.

Let $z \in \mathcal{M}$ determine an affine flag $F = F_p$ and monodromy $u = u_p$. Since u fixes $F_{(a)}$, when we evaluate the maelstrom diagram on z, the tensors flowing in to any of the spiraling crossings are of the form $v \wedge F_{(a)}$ and $w \wedge F_{(a)}$ for appropriate vectors $v, w \in V$. In such an instance, the right hand side of the crossing removal relation (34) only has two terms because the other terms define the zero invariant:

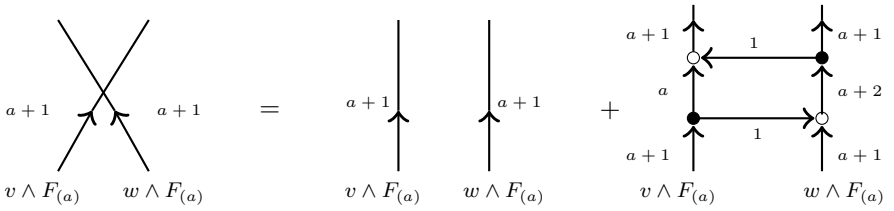

. (50)

In the rightmost diagram in this equation, the only nonzero contributions occur when the tensor $F_{(a)}$ flows to the edge of weight a and the vector v flows rightward along the edge of weight 1. That is, there are no choices on where the tensor $F_{(a)}$ must flow, and we can reduce to the case that $a = 0$. Once we set $a = 0$, the crossing removal relation becomes

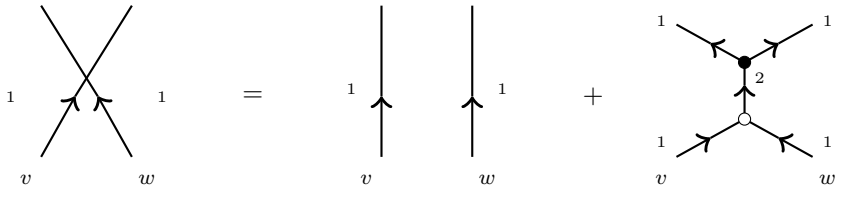

. (51)

We refer to the three terms appearing in this relation as the crossing term, the resolve term, and the bridge term. The maelstrom versus snail shell diagrams amount to choosing either the crossing term or the bridge term at every crossing. We can express the snail shell invariant as a linear combination of tensor diagrams in which we either take the crossing term or the resolve term. Regardless of how we make these binary choices, the resulting tensor diagram will have b edges which emanate from the leg of weight $[r]$ and spiral the puncture some number of times before meeting up with a copy of $T \setminus e$. The total number of times these edges spiral, namely $0 + 1 + \cdots + r - 1$, is conserved.

In any of these terms, the resulting invariant is antisymmetric in the r spiraling edges since they all emanate from the same black vertex. On the other hand, if any two of these edges spiral the same number of times around the puncture, then the corresponding invariant is also symmetric in these edges, which are all plugged into the same fragment $T \setminus e$. Thus, such an invariant is the zero invariant. Since the total amount of spiral around the puncture is conserved, one can see that in order for the r legs to spiral a distinct number of times, we must take the crossing term (not the resolve term) at *every* crossing, i.e. the maelstrom and snail shell diagrams determine the same invariant. □

With these lemmas in hand, we can prove Theorem 9.6.

Proof. Translating as needed by the birational Weyl group action, we may assume that $\mathrm{wt}_p \prod_i f_i$ lies in the closed dominant Weyl chamber at every puncture p. Moreover, if p is the puncture at which the $\mathrm{wt}_p f_i$'s are root-conjugate we may assume that the ground set (16) of the pairwise root-conjugate vectors $\mathrm{wt}_p f_i$ takes the form $[a+1, a+r]$, and that $[a+1, c]$ is a block of the ordered set partition $\Pi_{\mathrm{wt}_p \prod_i f_i}$ for some $c \geq a+r$. Let ε be the sign of the root-conjugacy class as in Lemma 5.3. Using the translation (29), when $\varepsilon = +1$, one can see that T has a unique leg e of weight $a+1$ and all its other legs at p either have weight at most a or at least $a+r$. A similar statement holds when $\varepsilon = -1$, but in this case the leg e has weight $a+r-1$. Because we have translated to the dominant Weyl chamber, our goal is to show that $\prod_i f_i$ is a planar invariant and a forest invariant. We will argue this when $\varepsilon = +1$; the other case is similar.

The r-fold superposition $\cup_{i=1}^r T$ has r many copies of the leg e. We can tag these r legs by the subsets $[a+1], [a] \cup \{2\}, \ldots, [a] \cup \{a+r\}$ (tagging all other legs by a fundamental weight) to obtain a pseudotagged tensor diagram whose invariant is $\prod_{i=1}^r f_i$. We denote this pseudotagged diagram by $(\cup_{i=1}^r T, \varphi)$. It is the upper left diagram in Fig. 9.

We denote by $T_{\mathrm{maelstrom}}$ (resp. T_{snail}) the plainly tagged diagram which results by replacing the r many root-conjugate legs in $(\cup_i T, \varphi)$ by the maelstrom (resp. snail shell) fragment and plugging this fragment into r many copies of the fragment $\eta := T \setminus e$. By Lemmas 9.8 and 9.9 we have $[(\cup_i T, \varphi)] = [T_{\mathrm{maelstrom}}] = [T_{\mathrm{snail}}]$.

The diagram $T_{\mathrm{maelstrom}}$ is not a superposition of smaller diagrams and is without interior cycles. Thus, it is a tree diagram which computes the monomial $\prod_i f_i$. This proves the first assertion in the theorem. The snail shell fragment is planar and the diagram $T_{\mathrm{maelstrom}}$ is obtained by gluing this planar fragment to r many copies of $T \setminus e$. This proves the second assertion in the theorem. □

9.5. Undoing the spiral when $k = 3$

We conclude this section by strengthening the above results to the case $k = 3$, thereby extending the conjectures from [8,7] to surfaces with punctures. When $k = 3$, the weight of a leg at a puncture is either 1 or 2, and any such leg can be tagged in three possible ways.

We refer to [8,7] for a discussion of the arborization algorithm converting a planar SL_3 web diagram into its *arborized form*. Conjecturally, this algorithm has the following properties when $\mathbb{S}$ has no punctures: a non-elliptic web diagram is a cluster variable if and only if its arborized form is a tree diagram; it is a cluster monomial if and only if its arborized form is a union of tree diagrams corresponding to the irreducible factors of the cluster monomial. We now extend these conjectures in the presence of punctures:

Conjecture 9.10. Let $[(T, \varphi)] \in \mathscr{A}(\mathcal{M})$ be a tagged non-elliptic SL_3 web invariant. Then $[(T, \varphi)]$ is a cluster variable if and only if the arborized form of T is a tree diagram. Moreover, $[(T, \varphi)]$ is a cluster monomial if and only (T, φ) can be transformed to a superposition of tagged tree diagrams (T_i, φ_i) by a sequence of arborization moves and

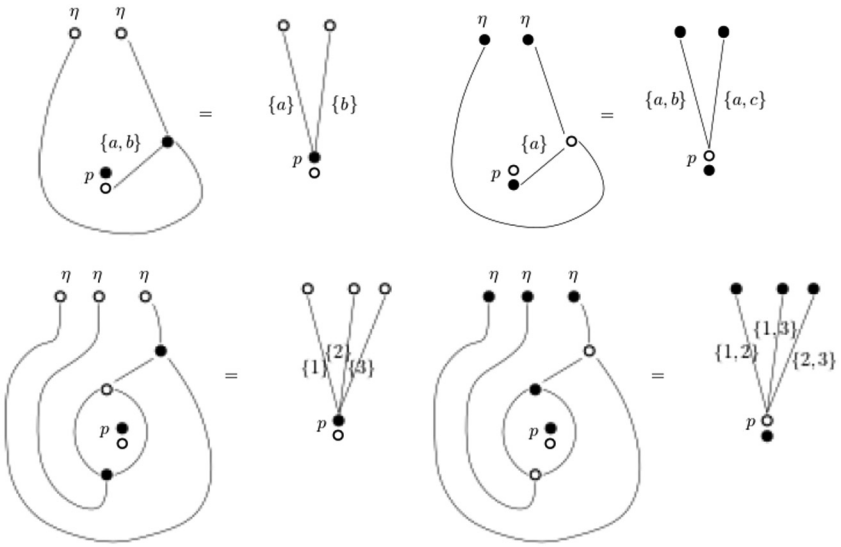

Fig. 10. Additional diagrammatic moves for tagged tensor diagrams when $G = \mathrm{SL}_3$. The black and white dots at the puncture p represent the first and second steps of the affine flag at p; the symbol η is a tensor diagram fragment.

the four moves in Fig. 10. In this case, the various $[(T_i, \varphi_i)]$ are the irreducible factors appearing in $[(T, \varphi)]$.

The four extra moves in the figure correspond to the passage from the snail shell fragment to the superposition of root-conjugate legs (i.e., from the bottom to the upper left diagram in Fig. 9). There are four moves corresponding to the four possible dosps with a non-singleton block, namely $ab|c$, $a|bc$, $|123^+|$, and $|123^-|$ from left to right and top to bottom.

Proposition 9.11. *Continuing the setup of Theorem 9.6 in the special case that $k = 3$, suppose that T is the arborized form of a non-elliptic web invariant. Then $\prod_i f_i$ is a tagged web invariant.*

Proof. Suppose that T is the arborized form a non-elliptic web diagram W. Denote by W_r the r-thickening of W as defined in [7, Definition 10.8]. Then W_r is again a non-elliptic web invariant. Recall that T_{snail} is the result of gluing the r dangling edges of the snail shell fragment to $\cup_i T$. We denote by W_{snail} the result of instead gluing the r dangling edges of the snail shell fragment to W_r.

By the "thickening theorem" [7, Theorem 10.9], the r-fold superposition $\cup_i T$ planarizes into W_r by a sequence of inverse arborization moves. We recall the proof of this theorem now and refer heavily to Figure 61 from [7] which underlies the proof. By assumption, we can planarize the tree diagram T into the planar diagram W by a sequence of inverse arborization moves ("planarization moves"). The thickening theorem says that we can planarize $\cup_i T$ into W_r by "thickening" each of these planarization moves (each

step in the sequence converting T into W corresponds to several moves turning $\cup_i T$ into W_r), working from the "inside out."

The aforementioned Figure 61 corresponds to the case $r = 3$. The left part of that figure has six identical tensor fragments entering from the bottom, three identical fragments entering from the upper left, three more entering from middle left, three more from middle right, and three more from upper right. The six fragments entering the bottom are not connected to any of the other fragments. We take the opportunity here to point out that, contrary to what is written in the proof of the thickening theorem, one can planarize the left figure into the right figure without any square moves. (One should planarize the triple of crossings at the top of the figure, then planarize the nine crossings in the middle, and finally planarize the right triple of crossings.)

Note that $\cup_i T$ coincides with T_{snail} outside of a neighborhood of the puncture p and likewise W_r coincides with W_{snail} away from p. We claim that we can planarize the crossings in T_{snail} to obtain W_{snail} via the same sequence of steps which planarizes the crossings in $\cup_i T$ to obtain W_r. Indeed, note that the distinguished edge $e \in \mathrm{Legs}_p(T)$ is the unique leg of its weight at p. Thus, when we carry out the steps outlined in the previous two paragraphs, the fragment containing e appears in one of the triples entering from the left or right (not from the bottom). Therefore, the planarization moves which turn the left diagram in Figure 61 into the right diagram are unaffected if we replace the three copies of the edge e by the three dangling edges in the seashell fragment.

Let $T' \subset T$ be a maximal planar subtree containing the distinguished edge e (start with the edge e and grow the tree one edge at a time provided the added edge does not cross any previously added edges). When we superimpose r copies of the tree T, we create self-crossings between the copies of T'. One may check that it is possible to planarize these self-crossings via arborization moves in which e never appears in the bottom of the arborization fragment.

If we first perform the planarization moves from the previous paragraph and subsequently perform those from two paragraphs previous, then we planarize $\cup_i T$ into W_r without the distinguished edge e ever appearing in the bottom of the arborization fragment. Thus, replacing the r distinguished edges e with the snail shell fragment and performing the same sequence of planarization moves, we transform T_{snail} into W_{snail}. □

10. Examples of finite mutation type

We study the $\mathscr{A}(\mathcal{M})$'s of finite mutation type. Our first proposition classifies these.

Proposition 10.1. *Suppose that $\mathbb{S}$ has at least one puncture and $k \geq 3$. Then the cluster algebra $\mathscr{A}_{\mathrm{SL}_k,\mathbb{S}}$ has finite mutation type if and only $\mathbb{S}$ is a once-punctured bigon and $k = 3$. Likewise, $\mathscr{A}'_{\mathrm{SL}_k,\mathbb{S}}$ has finite mutation type if and only if either $\mathbb{S}$ is a once-punctured bigon and $k = 3, 4$ or $\mathbb{S}$ is a once-punctured triangle and $k = 3$.*

Proof. We call a triangle *fully interior* if it has three distinct sides none of which is a boundary arch. When a fully interior triangle appears in a triangulation, it contributes a copy of the quiver Q_k with all of the vertices considered mutable vertices and with none of its vertices identified.

The quiver Q_4 is mutation infinite, and every Q_k when $k \geq 4$ contains this quiver as a full subquiver. The quiver Q_3 is mutation finite, but if we glue this quiver to that for an adjacent triangle (regardless of whether zero, one, or two sides of this adjacent triangle are boundary intervals carrying frozen variables), the resulting quiver is mutation infinite. Any fully interior triangle appearing in a triangulation will have such an adjacent triangle, so we can rule out the $k = 3$ cases with a fully interior triangle as well.

Recall that $\mathbb{S}$ has at least one puncture. One can check that any surface besides a once-punctured n-gon has a fully interior triangle, and from the preceding paragraph therefore admits a full subquiver of infinite mutation type. So it remains to consider the cases $\mathbb{S} = D_{n,1}$ for some $n \geq 2$.

One may check that the result of gluing three copies of Q_3, with the right edge of the first copy glued to the left edge of the second copy and the right edge of the second copy glued to the left edge of the third copy, has infinite mutation type. The same is true for gluing three copies of Q_3^1, and for gluing two copies of Q_4 along a shared edge. One can also check that $\mathscr{A}(\mathrm{SL}_k, D_{2,1})$ has infinite type when $k = 4$ hence when $k \geq 4$. In a similar way one checks that $\mathscr{A}'(\mathrm{SL}_k, D_{3,1})$ has infinite mutation type when $k \geq 4$ and $\mathscr{A}'(\mathrm{SL}_k, D_{2,1})$ has infinite mutation type when $k \geq 5$. □

Remark 10.2. The cluster algebra $\mathscr{A}(\mathrm{SL}_2, \mathbb{S})$ always has finite mutation type. We assumed that $\mathbb{S}$ has at least one puncture because that is our focus in the present paper. If we relax this assumption, there are additional finite mutation type examples arising when k is small and $\mathbb{S}$ is an n-gon. For example, $\mathscr{A}'(\mathrm{SL}_3, D_{n,0})$ has finite mutation type when $n \leq 9$. We are not aware of finite mutation type examples outside of those listed in the above proposition and cases when $\mathbb{S}$ is an n-gon, but we did not thoroughly rule these out.

One may verify that the cluster algebras $\mathscr{A}'(\mathrm{SL}_k, D_{3,1})$ and $\mathscr{A}(\mathrm{SL}_k, D_{2,1})$ are related by a quasi cluster isomorphism as predicted in Remark 3.5. This quasi cluster isomorphism respects the notions of being an arborizable web invariant in either side, and preserves the weight of cluster variables at punctures. Thus, we only need to understand the three finite mutation $\mathscr{A}'$-version examples listed in the above proposition. We study these three examples in the remainder of this section.

The cluster algebra $\mathscr{A}'(\mathrm{SL}_3, D_{2,1})$ has finite cluster type D_4. The other two examples both have the same cluster type as $\mathscr{A}(\mathrm{SL}_2, S_{0,4})$, also known as elliptic D_4 type.

Our next proposition summarizes our results in this section:

Proposition 10.3. *Conjectures 5.4, 8.18, and 8.20 hold for the cluster algebras* $\mathscr{A}'(\mathrm{SL}_3, D_{n,1})$ *when* $n = 2, 3$. *Conjecture 5.4 holds for* $\mathscr{A}'(\mathrm{SL}_4, D_{2,1})$, *and Conjectures 8.18 and 8.20 can be verified through an explicit finite computation.*

Before outlining the proof of this proposition, we need to introduce the following ingredient. Both of the elliptic D_4 examples have infinitely many clusters. To understand them, we introduce an appropriate group of quasi cluster automorphisms of the cluster algebra, modulo which there are only finitely many clusters. Provided we also check that this group preserves the assertions in our conjectures, this reduces our conjectures to a finite verification.

We number the boundary points of the once-punctured n-gon in cyclic order. A choice of point in the moduli space $\mathcal{M}$ affords us a vector v_i at the ith boundary point, as well as an affine flag F at the puncture p and a monodromy u around p.

We denote by ρ denote the *cyclic shift* cluster automorphism induced by rotating boundary points $i \mapsto i + 1 \mod n$. We let τ denote a fixed reflection in the n-gon, so that $\langle \rho, \tau \rangle$ is a dihedral group.

The cluster algebra has a canonical initial seed since there is a unique taut triangulation of a once-punctured n-gon. We use the notation A, B, C for the initial cluster variables associated to the arc connecting the first boundary point to the puncture, so that the initial cluster variable A has weight a at the puncture in multiplicative notation, the initial cluster variable B has weight ab, and the initial cluster variable C has weight abc. Note that C only is present when $k = 4$. The cluster variables associated to the arc connecting the second boundary point to p are cyclic shifts of these, namely $\rho^* A$, $\rho^* B$, and $\rho^* C$. And so on.

Definition 10.4. We define a birational automorphism $\sigma \in \mathrm{Bir}(\mathcal{A}_{\mathrm{SL}_3, D_{3,1}})$ as follows. The automorphism changes the decorations at boundary points according to

$$v_1 \mapsto v_2 \qquad v_2 \mapsto (v_1 \wedge v_2) \cap (v_3 \wedge u^{-1} v_1) \qquad v_3 \mapsto v_3 \tag{52}$$

while preserving the local system and the affine flag F.

We have similarly a birational automorphism $\sigma \in \mathrm{Bir}(\mathcal{A}_{\mathrm{SL}_4, D_{2,1}})$ defined via

$$v_1 \mapsto v_2 \qquad v_2 \mapsto (v_1 \wedge v_2) \cap (u^{-1} v_1 \wedge u^{-1} v_2 \wedge u^{-2} v_1).$$

Lemma 10.5. *The transformation* σ *is a quasi cluster automorphism of the corresponding cluster algebra. It acts on the cluster algebra while preserving the sets of diagram invariants, web invariants, tree invariants, and forest invariants.*

Proof. One checks that σ is a quasi cluster automorphism by providing an explicit sequence permutation-mutation sequence which matches Definition 10.4. The required sequences are

$$(B\,\rho^2(A)\,\rho(A)\,\rho(B)) \circ \mu_B\mu_{\rho A} \text{ for the once-punctured triangle, } k=3 \tag{53}$$

$$(A\,\rho(A)\,\rho(B)) \circ (C\,\rho(C)) \circ \mu_C\mu_{\rho A}\mu_{\rho B}\mu_C \text{ for the once-punctured bigon, } k=4. \tag{54}$$

In both cases, we wrote the required permutation of vertices in cycle notation.

If T is a tensor diagram, then one can compute $\sigma_*([T])$ by "plugging in" the boundary edges of T to a tensor-diagrammatical gadget which computes the formulas (52). This plugging in process does not create any interior cycles on boundary vertices, so σ preserves the sets of forest and tree invariants. The crossings which are created by plugging in can be planarized using the crossing removal relation (only one of the terms in the crossing removal relation survives), so σ also preserves the set of web invariants. □

Now we sketch the proof of Proposition 10.3 with further details filled in on a case by case basis below.

Proof. In all three cases, there are only finitely many P-clusters and the Conjecture 5.4 can be verified by hand and by inspection. We have listed these P-clusters up to W-action in (17), (56), and (58).

In the finite type example, we compute below all cluster variables and clusters explicitly as tagged tensor diagram invariants. One can see the validity of our two cluster combinatorics conjectures by inspection.

In the two elliptic D_4 examples, consider the group generated by the quasi cluster automorphism σ identified above, together with the cyclic shift map ρ, the duality map $*$, and the Weyl group action at the puncture. This group acts on the cluster algebra by (quasi) cluster automorphisms and preserves the sets of tagged web invariants and tagged forest invariants.

We argue in both cases elliptic D_4 examples that there are only finitely many clusters modulo the action of this group. When $k=3$ example, we identify explicitly these finitely many clusters and argue the following extra steps. First, every cluster monomial in this finite list of clusters is a tagged web invariant and a tagged forest invariant. Second, if x, y are cluster variables in this finite list, and if the weights of x and y at the puncture are not W-sortable, then x and y are in fact root-conjugate taggings of the same underlying tensor diagram. Third, if x appears in our finite list, then x is cluster compatible with any of its root-conjugate taggings. (It is not important that a cluster witnessing this fifth assertion appears in our finite list.) The first of these statements proves Conjecture 8.18 while the second and third statements prove Conjecture 8.20. □

Remark 10.6. It would be possible to carry out the same three steps alluded to in the last paragraph of the above proof in the case of $\mathscr{A}'(\mathrm{SL}_4, D_{2,1})$, thereby verifying our cluster combinatorics conjectures for this cluster algebra. The number of clusters which needed to be explicitly computed was fairly large so we did not carry this out.

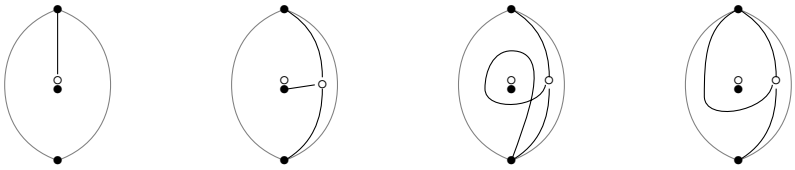

Fig. 11. The tensor diagrams underlying mutable and frozen variables in $\mathscr{A}'(\mathrm{SL}_3, D_{2,1})$, which has finite cluster type D_4.

10.1. The case of $\mathscr{A}'(\mathrm{SL}_3, D_{2,1})$

This is a type D_4 cluster algebra with two frozen variables, 16 cluster variables, and 50 clusters. We depict the underlying tensor diagrams (up to dihedral and W- action) in Fig. 11. The diagrams B and A both can be tagged in three ways and have two dihedral images. The diagram x has weight zero at the puncture and has four dihedral images. The two dihedral images of the last diagram f are the frozen variables. We have drawn x in its tree form but we can obtain its planar form by applying the crossing removal relation and noting that one of the two terms vanishes because it has a boundary 2-cycle.

Each of the four variables $\{A, B, \rho(A), \rho(B)\}$ is compatible with exactly two out of the four cluster variables $\{x, \rho(x), \tau(x), \tau\rho(x)\}$ in both senses of compatibility (namely, cluster compatibility and also the compatibility notion that two web invariants are compatible when their product is again a web invariant). Specifically B is compatible with $\rho(x)$ and $\rho(\tau(x))$, $\rho(B)$ is compatible with x and $\tau(x)$, A is compatible with x and $\rho(\tau(x))$, and $\rho(A)$ is compatible with $\rho(x)$ and $\tau(x)$.

Fig. 3 shows all 50 clusters in this cluster algebra, grouped in concentric circle "levels" with 6, 12, 24, 8 clusters respectively. Equation (17) lists the P-clusters. Clusters in the outermost level have P-cluster in the first row of (17) up to W-action, those in the next level have P-cluster in rows two or three of (17) up to W-action, those in the next level have P-cluster in rows four or five, and those in the innermost have P-cluster in the last two rows.

Mutation from the first to second level is an instance of the flattening relation (45) with $\ell(S) = 1$, e.g. exchanging $\rho(B)$ for $s_2(B)$. Mutating from the second to the third level is an instance of the skein relations, e.g. exchanging $\rho(A)$ for $\rho\tau(x)$. There are also mutations within the third level which implement the dosp mutation $12|3 \leftrightarrow 12|3$. Mutation from the third to fourth level is an instance of $\ell(S) = 2$ flattening relation followed by additional skein relation, e.g. exchanging A for $s_2 s_1(B)$. Specifically, the flattening relation yields a linear combination $[T]-2[T']+[T'']$ after applying the binomial theorem to $(u^{-1}-1)^2$. One of the invariants $[T'']$ is equal to 0. The tensor diagram $[T]$ has a self-crossing and can be expressed as $[T] = [T'''] + 3[T']$ using the crossing removal relation and the fact that $\mathrm{tr}(u) = 3$. The right hand side of the flattening relation then becomes a sum of two terms $[T'''] + [T']$ and these two terms match the two terms

appearing the exchange relation. Finally, there are mutations within the fourth level exchanging two dihedral images of x (these are consequences of the skein relations).

10.2. The case of $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$

This cluster algebra has the same cluster type as the cluster algebra associated to a four-punctured sphere $S_{0,4}$. The underlying quiver mutation class consists of four quivers. There are six combinatorial types of tagged triangulations cf. [1]. We use the nomenclature for these six combinatorial types given in [1] and note that types I and VI have the same underlying quiver, as do types II and V. We distinguish between *unlabeled* and *labeled* triangulations – in the former we consider the arcs in a triangulation as a set while in the latter we think of these arcs as an ordered tuple. Any two tagged unlabeled triangulations of the same combinatorial type are related by an element of the tagged mapping class group $\mathrm{MCG}^{\bowtie}(S_{0,4})$.

We denote by T_0 the following labeled triangulation of $S_{0,4}$, where we use $p_1, \ldots, p_4$ to denote the four punctures and think of the sphere as $\mathbb{R}^2 \cup \{\infty\}$.

$$T_0 = \quad \begin{array}{l} 1 : A \qquad 2 : B \\ 3 : \rho A \qquad 4 : \rho B \\ 5 : \rho^2 A \qquad 6 : \rho^2 B \end{array} \tag{55}$$

To the right, we have identified the 6 initial cluster arcs for $S_{0,4}$ with the 6 initial cluster variables for $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$ in such a way that the mutable parts of the quivers are identified. We denote by T_1 the result of notching all arcs in T_0 at the puncture p_1.

Fixing the above matching of initial seeds, we obtain a bijection between cluster variables (resp. clusters) for $\mathscr{A}'(\mathrm{SL}_3, \mathbb{D}_{3,1})$ and tagged arcs (resp. tagged triangulations) in $S_{0,4}$. If Σ is a seed for $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$ we denote by $T(\Sigma)$ its corresponding tagged triangulation. In particular, every cluster for $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$ has one of the six aforementioned combinatorial types.

Figs. 12 and 13 together illustrate this bijection between cluster variables in specific instances. The bijection is equivariant with respect to the group of dihedral symmetries of $D_{3,1}$, where in the case of $S_{0,4}$ we view the punctures p_1, p_2, p_3 as the vertices of a triangle whose puncture is p_4. We have broken the two figures up so that the SL_3 tensor diagrams appearing in Fig. 12 have nonzero weight at the puncture (hence, W acts on them), while those in Fig. 13 have zero weight. We computed the tensor diagrams in this table by performing mutations using skein relations and the flattening relation. We omit the details.

Altogether, there are six plain SL_3 tensor diagrams appearing in these figures versus 4 plain SL_2 arcs. Each plain arc has three dihedral images and can be tagged in four

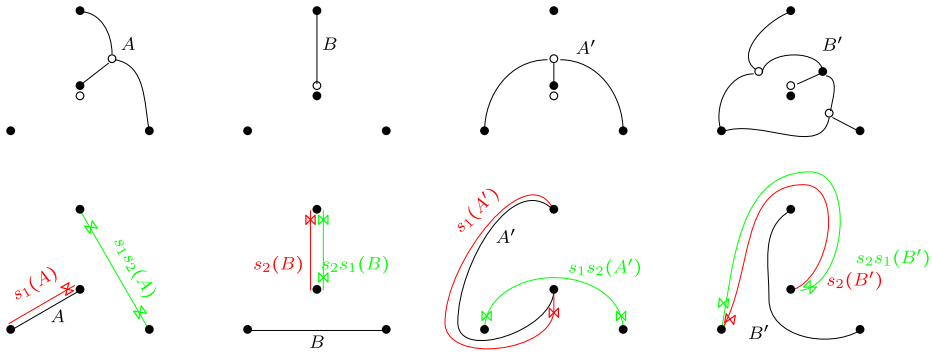

Fig. 12. The top row shows plain SL_3 tensor diagrams A, B, A', B' in the once-punctured triangle, each of which has three images under SL_3-Weyl group action at the puncture. In the bottom row, we depict the three tagged arcs in $S_{0,4}$ which correspond to each of these W-orbits.

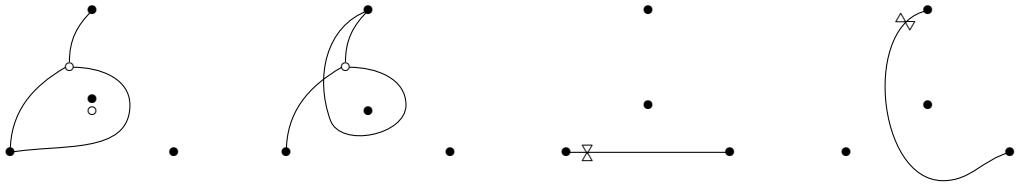

Fig. 13. The first and second SL_3 tensor diagrams in the once-punctured triangle correspond respectively to the third and fourth tagged arcs in $S_{0,4}$. Each of these has six dihedral images.

ways, for a total of 48 tagged arcs appearing in the two figures. The SL_3 diagrams in the first figure have 3 dihedral images and can be tagged in 3 ways, while those in the second figure have six dihedral images and cannot be tagged, again for a total of 48 diagrams.

Each quasi cluster automorphism of $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$ induces an element of the cluster modular group for $S_{0,4}$. We list some instances of this correspondence in the table below:

$\mathrm{Aut}(\mathcal{A}')$	$\mathrm{MCG}^{\bowtie}(S_{0,4})$	Sequence
ρ	120° rotation about p_4	$(135)(246)$
σ	half-twist $\overset{p_2}{\underset{p_1}{\leftrightarrows}}$	$(2534)\mu_{23}$
s_1	tag-change at p_4	$(35)\mu_{1351}$
s_2	N/A	$(46)\mu_{2462}$
$w_0 * \rho$	tag-change at all punctures	

The third column shows the permutation-mutation sequence out of the initial seed which induces the corresponding automorphism. The notation μ_{23} means to mutate at vertex 2 and then at vertex 3. The quasi automorphism in the last row is the Donaldson-Thomas transformation, whose mutation sequence we have omitted due to its length. The quasi automorphism in the fourth row does not correspond to a tagged mapping class, but rather an "exotic" element of the cluster modular group for $S_{0,4}$ matching up a type I triangulation with a type VI triangulation.

We now carry out the three additional steps from Proposition 10.3, thereby verifying all of our conjectures for $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$.

Proof. We begin by explaining that there are only finitely many clusters modulo H' action, where H' is the group defined in the proof of the above proposition. Then we carry out the three extra steps from the proof of the proposition. There is a group homomorphism $B_3 \to \mathrm{MCG}(S_{0,4})$ from the braid group on three strands to the mapping class group of a four-punctured sphere. The image of this homomorphism is the subgroup of mapping classes which fix the fourth puncture, see [5, Section 9.2]. This subgroup has index four inside $\mathrm{MCG}(S_{0,4})$ with cosets distinguished by which puncture is brought to p_4. Looking at the above table, the elements σ and $\rho\sigma\rho^{-1}$ are the images of the Artin generators under this homomorphism, thus the group $H := \langle \sigma, \rho \rangle$ has index four in $\mathrm{MCG}(S_{0,4})$. There are clearly only finitely many clusters modulo the action of this group. There will be even fewer clusters modulo the action of the larger group H'.

To carry out the three extra steps from the proof of the proposition, we explicitly identify a finite list of clusters which exhaust all clusters modulo H'-action.

One can check that the $\mathrm{MCG}(S_{0,4})$-orbit of the unlabeled triangulation T_0 coincides with its H-orbit. (The mutations needed to bring a given p_i, $i = 1, 2, 3$ to the North pole are the same mutations needed to perform a half-twist. The two mapping classes differ only by a permutation of variables.)

One can also check that there are six ways to mutate from a type I triangulation to a type II triangulation, four ways to mutate from a type II to a type IV, one way to mutate from a type IV to a type V, two ways to mutate from a type V to a type VI, and one way to mutate from a type V to a type III. All other mutations move weakly "backwards" to an earlier combinatorial type. Altogether, given a particular type I triangulation T, we get $6 \cdot 4 \cdot 1 \cdot 2 + 6 \cdot 4 \cdot 1 = 72$ mutation sequences which efficiently pass from a given type I triangulation to a triangulation of any other type. Note that some steps in these mutation sequences commute so these 72 mutations sequences do not yield 72 distinct clusters. We refer to the clusters reachable from T by one of these 72 mutations sequences as the *club* of T.

We claim that every tagged triangulation of $S_{0,4}$ is in the club of either T_0 or T_1 modulo H'-action. Indeed, suppose that T has type I and let T' be it underlying plain triangulation. By two paragraphs previous, there exists an element $h_1 \in H$ such that $h_1(T')$ is a relabeling of T_0. We can choose $h_2 \in \langle \mathrm{DT}, s_1 \rangle$ (i.e., in the group generated by tag-changing at p_4 or at all the p_i's) such that $h_2 h_1(T) \in \{T_0, T_1\}$ as unlabeled tagged triangulations. Thus, if T has type I, then we can move it using H' to either T_0 or T_1. If we consider a tagged triangulation T' of some other type, then T' is in the club of some type I triangulation T. Choosing h_1 and h_2 as in the previous sentences so that $h_2 h_1(T) \in \{T_0, T_1\}$, then $h_2 h_1(T')$ is in the club of either T_0 or T_1 as claimed.

The clubs of T_0 and T_1 serves as our list of finitely many clusters modulo H'-action.

Now we carry out the first extra step, i.e. we check that every cluster monomial drawn from the clubs of T_0 and T_1 is a tagged web invariant and a tagged forest invariant. This is a straightforward but lengthy check using Figs. 12 and 13. Every tagged triangulation in these two clubs is built out of the 48 tagged arcs dihedrally related to these two figures. One must verify that when a collection of such tagged arcs are pairwise compatible then so are their corresponding tensor diagrams, where the notion of compatibility for SL_3 tensor diagrams is that their product is a tagged web invariant. The proof of the second extra step is illustrated in Fig. 12: in the clubs of T_0 and T_1, the cluster variables which have nonzero weight are listed in the top of the figure, and each of these is compatible with its three taggings (because the three tagged arcs below each diagram are compatible tagged arcs).

The proof of the third extra step is another straightforward check using Fig. 12. For example, the tensor diagram A has weight ω_1 at the puncture, so it is not W-sortable with tensor diagrams of weight e_2, e_3, or e_2+e_3. Looking right in the figure, the diagrams that have this weight and which are not root-conjugate to A are $s_1s_2(A')$, $s_1(A')$, $s_2s_1(B')$, and $s_2s_1(B)$. We can see that the tagged arc corresponding to A is not compatible with any of the four tagged arcs which correspond to these cluster variables. □

The P-clusters in $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$ are listed below up to W-action:

P-cluster	dosp
$e_1 \times 3, e_1 + e_2 \times 3$	$1\vert 2\vert 3$
$e_1 \times 2, e_1 + e_2 \times 4$	
$e_1 \times 4, e_1 + e_2 \times 2$	
$e_1 \times 2, e_1 + e_2 \times 3, 0$	
$e_1 \times 3, e_1 + e_2 \times 2, 0$	
$e_1 \times 2, e_1 + e_2 \times 2, 0 \times 2$	
$e_1, e_2, e_1 + e_2 \times 4$	$12\vert 3$
$e_1, e_2, e_1 + e_2 \times 3, 0$	
$e_1, e_2, e_1 + e_2 \times 2, 0 \times 2$	
$e_1, e_2, e_1 + e_2, 1 \times 3$	

P-cluster	dosp
$e_1 \times 4, e_1 + e_2, e_1 + e_3$	$1\vert 23$
$e_1 \times 3, e_1 + e_2, e_1 + e_3, 1$	
$e_1 \times 2, e_1 + e_2, e_1 + e_3, 0 \times 2$	
$e_1 \times 1, e_1 + e_2, e_1 + e_3, 0 \times 3$	
$e_1, e_2, e_3, 0 \times 3$	123^+
$e_1 + e_2, e_1 + e_3, e_2 + e_3, 0 \times 3$	123^-

(56)

Remark 10.7. Using the above argument, it is not hard to see that every cluster variable in $\mathscr{A}'(\mathrm{SL}_3, D_{3,1})$ is a tagged tree invariant, and that every cluster monomial factors as a product of these using arborization moves and the four extra moves from Fig. 10, as predicted in Conjecture 9.10. We do not know how to prove the converse directions of these statements, e.g. that a web diagram which arborizes to a tree diagram indeed determines a cluster variable.

10.3. The case of $\mathscr{A}'(\mathrm{SL}_4, D_{2,1})$

This cluster algebra also has the same cluster type as a four-punctured sphere. We follow the style of argument from the previous section making the needed modifications. We take our initial triangulation T_0 to be the following labeled tagged triangulation of $S_{0,4}$, which is a type V triangulation:

$T_0 =$ 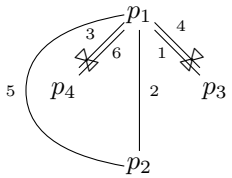

$1 : A \qquad 2 : B$
$3 : C \qquad 4 : \rho C$
$5 : \rho B \qquad 6 : \rho A$

(57)

We identify the initial tagged arcs with initial cluster variables for $\mathcal{A}'_{\mathrm{SL}_4, D_{2,1}}$ as indicated to the right. This induces the following correspondence between quasi cluster automorphisms of $\mathcal{A}'_{\mathrm{SL}_4, D_{2,1}}$ and elements of the cluster modular group of $\mathscr{A}(S_{0,4})$:

$\mathrm{Aut}(\mathcal{A}')$	$\mathrm{MCG}^{\bowtie}(S_{0,4})$	Sequence
ρ	lift of (34)	(16)(34)(25)
$\sigma \circ \rho$	Dehn twist about p_1, p_4	$(152)\mu_{4214}$
$\rho \circ \sigma$	Dehn twist about p_1, p_3	$(625)\mu_{3563}$
$s_1 \circ s_3$	lift of (12)(34)	
s_2	tag change at p_2	$(25)\mu_{25}$
$*$	tag change at both p_3, p_4	(14)(36)
	tag-change at all punctures	

There is a group homomorphism from $\mathrm{MCG}(S_{0,4})$ to the symmetric group on four symbols keeping track of how the punctures are permuted. In the middle column above, "lift of (34)" indicates a mapping class whose image under this homomorphism is the transposition (34), etc. The two Dehn twists in this table are pure mapping classes, i.e. their image is the identity permutation.

We now explain that there are only finitely many clusters modulo the action of the group defined in the proof of Proposition 10.3.

Proof. The two compositions $\rho \circ \sigma$ and $\sigma \circ \rho$ correspond to Dehn twists about simple closed curves with geometric intersection number two. Any two such mapping classes generate the pure mapping class group of $S_{0,4}$, see [5, Section 4.2.4]. There are clearly finitely many tagged triangulations of $S_{0,4}$ modulo the pure mapping class group. □

We note that we can reduce the size of the finite check needed to verify our conjectures by precomposing with ρ, $s_1 \circ s_3$, s_2, and $*$, as these quasi cluster transformations correspond to certain non-pure mapping classes and tag-changing transformations.

The P-clusters for $\mathscr{A}'(\mathrm{SL}_4, D_{2,1})$ are listed below up to W action, this time in multiplicative notation for space-saving purposes:

P-cluster	dosp
$a \times 2, ab \times 2, abc \times 2$	$1\|2\|3\|4$
$a, b, ab \times 2, abc \times 2$	$12\|3\|4$
$a, b, ab, abc \times 3$	
$a, b, ab, abc \times 2, 1$	
$a \times 2, ab \times 2, abc, abd$	$1\|2\|34$
$a \times 3, ab, abc, abd$	
$a \times 2, ab, abc, abd, 1$	
$a \times 2, ab, ac, abc \times 2$	$1\|23\|4$
$a, ab, ac, abc \times 3$	
$a \times 3, ab, ac, abc$	
$a, ab, ac, abc \times 2, 1$	
$a \times 2, ab, ac, abc, 1$	
a, b, ab, ab, abc, abd	$12\|34$
$a, b, ab, abc, abd, 1$	
$a, b, abc, abd, 1 \times 2$	

P-cluster	dosp
$a, b, c, abc, abc, 1$	$123^+\|4$
$a, b, c, abc, 1 \times 2$	
$a, b, c, abc \times 3$	
$ab, ac, bc, abc \times 3$	$123^-\|4$
$ab, ac, bc, abc \times 2, 1$	
$a \times 3, ab, ac, ad$	$1\|234^+$
$a \times 2, ab, ac, ad, 1$	
$a \times 2, abc, abd, acd, 1$	$1\|234^-$
$a, abc, abd, acd, 1 \times 2$	
$a \times 3, abc, abd, acd$	
$abc, abd, acd, bcd, 1 \times 2$	1234^-
$a, b, c, d, 1 \times 2$	1234^+

(58)

References

[1] E. Barnard, E. Meehan, N. Reading, S. Viel, Universal geometric coefficients for the four-punctured sphere, Ann. Comb. 22 (1) (2018) 1–44.
[2] S.C. Billey, W. Jockusch, R.P. Stanley, Some combinatorial properties of Schubert polynomials, J. Algebraic Comb. 2 (4) (1993) 345–374.
[3] S. Cautis, J. Kamnitzer, S. Morrison, Webs and quantum skew Howe duality, Math. Ann. 360 (1–2) (2014) 351–390.
[4] D. Douglas, Z. Sun, Tropical Fock-Goncharov coordinates for SL_3-webs on surfaces I: construction, arXiv:2011.01768, 2021.
[5] B. Farb, D. Margalit, A Primer on Mapping Class Groups, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012.
[6] V.V. Fock, A. Goncharov, Moduli spaces of local systems and higher Teichmüller theory, Publ. Math. Inst. Hautes Études Sci. 103 (2006) 1–211.
[7] S. Fomin, P. Pylyavskyy, Tensor diagrams and cluster algebras, Adv. Math. (2016) 717–787.
[8] S. Fomin, P. Pylyavskyy, Webs on surfaces, rings of invariants, and cluster algebras, Proc. Natl. Acad. Sci. USA 111 (27) (2014) 9680–9687.
[9] S. Fomin, M. Shapiro, D. Thurston, Cluster algebras and triangulated surfaces. Part I: cluster complexes, Acta Math. 201 (1) (2008) 83–146.
[10] S. Fomin, D. Thurston, Cluster algebras and triangulated surfaces. Part II: lambda lengths, Mem. Am. Math. Soc. 255 (1223) (2018), v+97 pp.
[11] C. Fraser, Braid group symmetries of Grassmannian cluster algebras, Sel. Math. New Ser. 26 (2) (2020) 17.
[12] C. Fraser, Quasi-homomorphisms of cluster algebras, Adv. Appl. Math. 81 (2016) 40–77.
[13] C. Frohman, A. Sikora, SU(3)-skein algebras and webs on surfaces, Math. Z. 300 (2022) 33–56, arXiv:2002.08151.

[14] A. Goncharov, L. Shen, Donaldson-Thomas transformations for moduli spaces of G-local systems, Adv. Math. 327 (2018) 225–348.
[15] A. Goncharov, L. Shen, Geometry of canonical bases and mirror symmetry, Invent. Math. 202 (2) (2015) 487–633.
[16] A. Goncharov, L. Shen, Quantum geometry of moduli spaces of local systems and representation theory, arXiv:1904.10491, 2020.
[17] J.E. Grabowski, Graded cluster algebras, J. Algebraic Comb. 42 (4) (2015) 1111–1134.
[18] H.K. Kim, SL_3-laminations as bases for PGL_3 cluster varieties for surfaces, Memoirs of the AMS, in press, arXiv:2011.14765, 2021.
[19] R. Inoue, T. Ishibashi, H. Oya, Cluster realizations of Weyl groups and higher Teichmüller theory, Sel. Math. 27 (37) (2021), arXiv:1902.02716.
[20] R. Inoue, T. Lam, P. Pylyavskyy, On the cluster nature and quantization of geometric R-matrices, Publ. Res. Inst. Math. Sci. 55 (1) (2019) 25–78.
[21] T. Ishibashi, W. Yuasa, Skein and cluster algebras of marked surfaces without punctures for $\mathfrak{sl}_3$, arXiv:2101.00643, 2021.
[22] R. Kedem, Q-systems as cluster algebras, J. Phys. A 41 (19) (2008).
[23] G. Kuperberg, Spiders for rank 2 Lie algebras, Commun. Math. Phys. 180 (1) (1996) 109–151.
[24] I. Le, Cluster structures on higher Teichmüller spaces for classical groups, Forum Math. Sigma 7 (2019) e13.
[25] I. Le, C. Fraser, Tropicalization of positive Grassmannians, Sel. Math. New Ser. 25 (5) (2019) 75.
[26] T. Magee, Littlewood-Richardson coefficients via mirror symmetry for cluster varieties, Proc. Lond. Math. Soc. (3) 121 (3) (2020) 463–512.
[27] M.R. Mills, On the relationship between green-to-red sequences, local-acyclicity, and upper cluster algebras, arXiv:1804.00479, 2018.
[28] G. Muller, Skein and cluster algebras of marked surfaces, Quantum Topol. 7 (3) (2016) 435–503.
[29] G. Musiker, R. Schiffler, L. Williams, Bases for cluster algebras from surfaces, Compos. Math. 149 (2) (2013) 217–263.
[30] P. Pylyavskyy, Zamolodchikov integrability via rings of invariants, J. Integrable Syst. 1 (1) (January 2016), xyw010.
[31] G. Schrader, A. Shapiro, A cluster realization of Uq(sln) from quantum character varieties, Invent. Math. 216 (3) (2019) 799–846.
[32] L. Shen, Duals of semisimple Poisson-Lie groups and cluster theory of moduli spaces of G-local systems, Int. Math. Res. Not. (2021), in press, arXiv:2003.07901.
[33] B. Sturmfels, Algorithms in Invariant Theory, Springer-Verlag, 1993.