

POSETS FOR F -POLYNOMIALS IN CLUSTER ALGEBRAS FROM SURFACES

VINCENT PILAUD, NATHAN READING, AND SIBYLLE SCHROLL

ABSTRACT. We prove a simple formula for arbitrary cluster variables in the marked surfaces model. As part of the formula, we associate a labeled poset to each tagged arc, such that the associated F -polynomial is a weighted sum of order ideals. Each element of the poset has a weight, and the weight of an ideal is the product of the weights of the elements of the ideal. In the unpunctured case, the weight on each element is a single \hat{y}_i , in the usual sense of principal coefficients. In the presence of punctures, some elements may have weights of the form \hat{y}_i/\hat{y}_j . Our search for such a formula was inspired by the Fundamental Theorem of Finite Distributive Lattices combined with work of Gregg Musiker, Ralf Schiffler, and Lauren Williams that, in some cases, organized the terms of the F -polynomial into a distributive lattice. The proof consists of a simple and poset-theoretically natural argument in a special case, followed by a geometric argument using a cover of the surface to prove the general case.

Cluster variable formulas were given by James Propp [Pro02], Gabriel Carroll and Gregory Price [CP03], and Ralf Schiffler [Sch08] for the cluster algebra of type A associated to the triangulations of a polygon. These formulas were later extended for cluster algebras arising from (tagged) triangulations of a surface [FST08, FT18a], first by Gregg Musiker and Ralf Schiffler in terms of perfect matchings on tile graphs [MS10] for unpunctured surfaces, and later by Gregg Musiker, Ralf Schiffler and Lauren Williams in terms of perfect matchings on snake graphs [MSW11] for any bordered marked surface with or without punctures. Further reformulations also appeared in [CP15, Yur19a, Yur19b]. Note that, while these formulas are often referred to as “cluster expansion formulas” in the literature, we prefer the term “cluster variable formulas” which avoids confusion with the geometric notion of cluster expansion in the cluster fan [FZ03].

The objective of this paper is to prove a more condensed cluster variable formula for cluster algebras on marked surfaces. We present the formula independently of the formula for snake graphs, but lying under the surface is the observation by James Propp [Pro02, Theorem 2] that the set of perfect matchings on a graph forms a distributive lattice. (This observation was already mentioned in the context of snake graphs by Gregg Musiker, Ralf Schiffler, and Lauren Williams [MSW13, Theorem 5.2]. It was also a crucial ingredient in [CS21], where it was used to characterize the canonical submodule lattice of a string module. The present paper got its start from conversations among the authors about the results of [CS21].)

The fact that the monomials in the cluster variable form a distributive lattice calls to mind the Fundamental Theorem of Finite Distributive lattices (FTFDL), due to Birkhoff [Bir37] and found, for example, in [Sta12, Theorem 3.4.1]. (The relevant direction here is that any finite distributive lattice is isomorphic to the inclusion order on order ideals in its poset of join-irreducible elements.) One is led to ask what the poset P of join-irreducible elements of this lattice is, and whether P is relevant to the cluster variable formula. Optimistically, one would hope to write down a poset P_α directly from a tagged arc α , and would hope that the F -polynomial associated to α is a sum of monomials indexed by order ideals in P_α . More optimistically, one would hope that the monomials could be recovered from order ideals in P_α simply by giving a weight to each element of P_α such that the monomial corresponding to an order ideal of P_α is just the product of the weights of its elements. That is precisely what happens. Most optimistically, each element e of P_α is

VP is partially supported by the French project CHARMS (ANR 19CE40 0017), and by the French–Austrian project PAGCAP (ANR 21 CE48 0020 & FWF I 5788). NR is partially supported by the National Science Foundation under award numbers DMS-1500949 and DMS-2054489 and by the Simons Foundation under award number 581608. SS is partially supported by the DFG through the project SFB/TRR 191 Symplectic Structures in Geometry, Algebra and Dynamics (Projektnummer 281071066- TRR 191).




FIGURE 1. A tagged arc and the corresponding labeled poset.

associated to a tagged arc in the initial tagged triangulation, and the weight of e is precisely the monomial \hat{y}_i associated to that arc (in the usual principal-coefficients notation of [FT18b]). This most optimistic possibility occurs whenever the surface has no punctures, and more generally when the base triangulation has no self-folded triangles. In the general case, the situation is nearly as good: Each element e of P_α is again associated to an arc, but in some cases when i and j index the edges of a self-folded triangle, the weight on an arc is of the form \hat{y}_i/\hat{y}_j .

As usual, we may as well take the initial tagged triangulation T to have all arcs tagged plain, except possibly at some punctures p incident to exactly two tagged arcs of T , identical except for opposite taggings at p . (If T does not have this property, we can modify it by reversing all taggings at certain punctures, the only cost being that we must make the same modification of taggings in all arcs α for which we compute the cluster variable x_α .) Also as usual, we will pass between T and the ordinary triangulation T° with self-folded triangles replacing pairs of coinciding arcs with opposite tagging at one puncture, considering T and T° to be different ways of viewing the same combinatorial data.

Details on the construction of the poset P_α appear in the next section. For now, we state the general idea and give an example. The poset P_α has a central piece that we will call P_α^{cr} . The undirected graph underlying the Hasse diagram of P_α^{cr} is a path, recording essentially the sequence of arcs of T° that α crosses. In many cases when α is tagged notched at one or more endpoints, we adjoin to P_α^{cr} one or more chains, with each adjoined chain recording the sequence of arcs that are incompatible with α because of its tagging at that endpoint. When α agrees, up to tagging, with an arc in T , then the construction “degenerates” in various ways that can be precisely described.

Example 1. The left picture in Figure 1 shows an example of a marked surface (a 4-punctured disk) with a triangulation T° (with arcs numbered 1 through 11) and a tagged arc α (shown thicker, in purple). The corresponding cluster variable is $\frac{x_5 x_6 x_8}{x_1 x_4 x_7 x_9}$ times the weighted sum of all order ideals in the poset P_α shown in the right picture. The weight of an element is \hat{y}_i if the element is labeled i . The weight of the element labeled $\frac{2}{3}$ is $\frac{\hat{y}_2}{\hat{y}_3}$. The poset P_α^{cr} is like the poset shown, but deleting the chain labeled $4 < 1 < 8 < 11$ on the right. We explain this example in full detail in Section 1.

The unpunctured case of the cluster variable formula can also be obtained from [MSW13, Theorem 5.4]. Jon Wilson [Wil20] constructed similar weighted posets and proved similar cluster algebra variable formulas in the “nondegenerate” cases (with some restrictions on the surface), proving his results based on the cluster variable formulas in terms of the snake graphs of [MSW11]. Some of these posets appeared more recently in [OY22], and similar posets also appear in [Wen23], where they are used to compute Donaldson-Thomas F -polynomials.

We conclude this introduction with an outline of the paper and a discussion of background references.

Section 1. We state the formula (Theorem 1.1) for the cluster variable associated to a tagged arc and illustrate it with numerous examples. We also make a standard, slight reduction of the theorem, showing that it needs only be proved for a subset of the tagged arcs.

Section 2. We prove in a special case that the cluster variable formula satisfies the correct exchange relations. The point is that certain exchange relations correspond to simple, natural decompositions of the set of order ideals of the corresponding posets (Section 2.3), and to natural relations between the \mathbf{g} -vectors (Section 2.4).

Section 3. We prove the analog of Theorem 1.1 in the coefficient-free case. This coefficient-free analog is a consequence of Theorem 1.1 and is not needed for the proof of Theorem 1.1. However, it is included here as a simplified version of the proof of Theorem 1.1. More specifically, the coefficient-free case uses only the fairly standard hyperbolic geometry of marked surfaces (avoiding the complications of opened surfaces and the corresponding lambda lengths and tropical lambda lengths). It is hoped that the separate treatment of this simpler coefficient-free case will serve to illustrate a key idea of the proof before introducing the complications of the full proof. The basic structure of the proof is an induction on the number of elements of P_α . The key idea is to pass to a cover of part of the surface and lift the hyperbolic metric to the cover, so that the exchange relations established in Section 2 serve as the inductive step. (In the simplest case of plain-tagged arcs, the combinatorics of passing to a cover is similar to the construction in [MSW11, Section 7]. The construction here for arbitrary arcs and the lifting of the hyperbolic metric appear to be new.)

Section 4. We prove the full version of Theorem 1.1. The proof structure is the same as the simplified proof in Section 3 and reuses many of the tools, but with the full machinery of opened surfaces, lambda lengths, and tropical lambda lengths.

Section 5. The construction that takes an arc and defines the associated cluster variables as a laminated lambda length [FT18a, Definition 15.3] is valid more generally for tagged geodesics connecting marked points (relaxing conditions such as the requirement that the geodesic not cross itself). We extend Theorem 1.1 to these more general tagged geodesics as Theorem 5.1. We also give a combinatorial characterization of tagged geodesics as Proposition 5.3.

Throughout the paper, we assume the most basic background on marked surfaces in the sense of Sergey Fomin, Michael Shapiro, and Dylan Thurston [FST08, FT18a]. Specifically, we use without explanation the material on (tagged) arcs, (tagged) triangulations signed adjacency matrices from [FST08, Sections 2–4 & 7] (summarized in [FT18a, Section 5]) and the material on laminations and shear coordinates from [FT18a, Sections 12–13]. Later, in support of the proof, we will review results of [FT18a] that use hyperbolic geometry and tropical hyperbolic geometry to realize cluster variables ([FT18a, Sections 7–8] for the simpler coefficient-free case in Section 3 and [FT18a, Sections 9–10 & 14–15] for the full proof in Section 4). We also assume the most basic background on cluster algebras of geometric type, F -polynomials, and \mathbf{g} -vectors from [FT18b, Sections 2–3 & 5–6]. Aside from this background, the paper is completely self-contained.

1. THE THEOREM

Let (\mathbf{S}, \mathbf{M}) be a marked surface and let T be a tagged triangulation of (\mathbf{S}, \mathbf{M}) with all arcs tagged plain, except possibly at some punctures p incident to exactly two tagged arcs of T , identical except for opposite taggings at p . The corresponding ordinary triangulation is denoted T° . Let $B(T)$ be the signed adjacency matrix of T , with entries $b_{\alpha\gamma}$ indexed by arcs α and γ in T . We associate a principal-coefficients cluster algebra $\mathcal{A}_\bullet(T)$ to (\mathbf{S}, \mathbf{M}) and T as usual. Specifically, extend $B(T)$ by appending the identity matrix below it. Let $\{x_\gamma \mid \gamma \in T\}$ be the initial cluster variables, and let $\{y_\gamma \mid \gamma \in T\}$ be the tropical variables (which are also the coefficients at the initial seed). For each $\gamma \in T$, let $\hat{y}_\gamma = y_\gamma \prod_{\beta \in T} x_\beta^{b_{\beta\gamma}}$.

Given a tagged arc α in (\mathbf{S}, \mathbf{M}) , let x_α be the associated cluster variable. For each tagged arc α , we will define the Laurent monomial \mathbf{g}_α in the initial cluster variables, the poset P_α , and the map $w : P_\alpha \rightarrow \{\hat{y}_\beta \mid \beta \in T\} \cup \{\hat{y}_\beta/\hat{y}_\gamma \mid \beta \neq \gamma \in T\}$ that appear in the following theorem.

Theorem 1.1. *The principal-coefficients cluster variable associated to a tagged arc α is*

$$(1) \quad x_\alpha = \mathbf{g}_\alpha \cdot \sum_I \prod_{e \in I} w(e),$$

where the sum is over all order ideals I in P_α .

The situation is simpler in the unpunctured case, or more generally when the triangulation T° has no self-folded triangles, because in that case the weights $w(e)$ are all of the form \hat{y}_α for $\alpha \in T$. We write $\text{arc}(e)$ for the tagged arc $\alpha \in T$ such that $w(e) = \hat{y}_{\text{arc}(e)}$. The following corollary is a special case of Theorem 1.1.

Corollary 1.2. *If the triangulation T° has no self-folded triangles, then the principal-coefficients cluster variable associated to an arc α is*

$$x_\alpha = \mathbf{g}_\alpha \cdot \sum_I \prod_{e \in I} \hat{y}_{\text{arc}(e)},$$

where the sum is over all order ideals I in P_α .

In the remainder of this section, we provide the precise definitions of the Laurent monomial \mathbf{g}_α , the poset P_α , and the weight map w that appear in Theorem 1.1 and Corollary 1.2.

1.1. Definition of the Laurent monomial \mathbf{g}_α . To define \mathbf{g}_α , we define a curve $\kappa(\alpha)$ that coincides with α except near the endpoints of α . Specifically, if α has an endpoint p at a marked point on the boundary of \mathbf{S} , then $\kappa(\alpha)$ ends on a boundary segment, at a point q very near p , chosen so that the path along the boundary from p to q keeps \mathbf{S} on the left. If α has an endpoint at a puncture p , where α is tagged plain, then $\kappa(\alpha)$ spirals clockwise into p . If α is tagged notched at p , then $\kappa(\alpha)$ spirals counterclockwise into p . (This definition follows [Rea14, Section 5] and is a modification of [FT18a, Definition 17.2].)

We define

$$(2) \quad \mathbf{g}_\alpha = \prod_{\gamma \in T} x_\gamma^{-b_\gamma(T, \kappa(\alpha))},$$

where $b_\gamma(T, \kappa(\alpha))$ stands for the γ -entry of the shear coordinates of $\kappa(\alpha)$ with respect to T .


Example 2. Figure 2 shows $\kappa(\alpha)$ for α as in Example 1. Also shown are the contributions $(-1, 0$ or 1 , abbreviated as $-$, \circ , or $+$) to $b_\gamma(T, \kappa(\alpha))$ for each γ . As part of the definition of shear coordinates, to find the shear coordinate $b_9(T, \alpha)$, we must reverse the spiral on α inside the self-folded triangle containing arc 9 and reverse the roles of arcs 9 and 10 (as shown in the right picture in Figure 2). Similarly, to find the shear coordinate $b_2(T, \alpha)$, we must reverse the roles of arcs 2 and 3, but since α has no spiral inside the self-folded triangle containing arc 2, this amounts to setting $b_2(T, \alpha)$ equal to $b_3(T, \alpha)$. Keeping in mind the negative sign in (2), we have justified the assertion in Example 1 that $\mathbf{g}_\alpha = \frac{x_5 x_6 x_8}{x_1 x_4 x_7 x_9}$.

The following proposition is [Rea14, Proposition 5.2]. The proposition also appears as [FT17, Lemma 8.6], where it is generalized to orbifolds.

Proposition 1.3. *Given a tagged arc α , the \mathbf{g} -vector of x_α has γ -entry $-b_\gamma(T, \kappa(\alpha))$ for all $\gamma \in T$.*

In light of Proposition 1.3, we call \mathbf{g}_α the **\mathbf{g} -vector** of α . The proofs of Proposition 1.3 in [Rea14, Proposition 5.2] and [FT17, Lemma 8.6] both rely on general structural results on cluster algebras. Proposition 1.3 is an immediate corollary of Theorem 1.1, by the definition of the \mathbf{g} -vector. Since we do not use Proposition 1.3 in any arguments in this paper, the proof of Theorem 1.1 given here in particular constitutes a new, direct proof of Proposition 1.3.

By [FT18b, Corollary 6.3], the content of Theorem 1.1 is Proposition 1.3 plus the assertion that $\sum_I \prod_{e \in I} w(e)$ is the F -polynomial of x_α in the variables \hat{y}_i . (When we apply [FT18b, Corollary 6.3], the denominator is 1 because we are in the principal-coefficients case.)


 FIGURE 2. Finding the monomial \mathbf{g}_α .

1.2. **Definition of the poset P_α and the weight map w .** We now proceed to define P_α and w . The key point is that we can describe α completely by giving the sequence of arcs of T° that α crosses and noting whether α is tagged notched or plain at its endpoints.

There are degenerate cases in the construction of P_α , occurring when α coincides, up to tagging, with an arc in T . We assume for the moment that we are not in a degenerate case.

To define P_α , we first define a smaller poset P_α^{cr} . The tagged arc α is defined up to isotopy, and (as is standard) we assume up to isotopy that as we move along α in a monotone fashion, when α crosses an arc γ of T° , it crosses a different tagged arc before ever crossing γ again. In particular, there is a finite set of points where α crosses arcs of T° . This set of points is the ground set of P_α^{cr} . (The endpoints of α are explicitly excluded from the ground set of P_α^{cr} .) Two points in the ground set are called *adjacent* if one can move from one point to the other along α without crossing any other arc of T° . The cover relations in P_α^{cr} come from pairs of adjacent points: Given adjacent points e and f with e in an arc β of T and f in an arc γ of T° , as one moves along α from e to f , one cuts off a corner of a triangle of T° . We set $e > f$ if the cut corner is to the right of α (as we move along α from e to f) or $e < f$ if the cut corner is on the left. We define P_α^{cr} to be the transitive closure of these cover relations.

In most cases, the weight function on P_α^{cr} has $w(e) = \hat{y}_\beta$, where β is the arc of T° containing e . There are two kinds of exceptions, both having to do with self-folded triangles. First, if β is the interior edge of a self-folded triangle of T° and γ is the exterior edge of the self-folded triangle and if α passes through γ before and after passing through β at e , then $w(e) = \frac{\hat{y}_\beta}{\hat{y}_\gamma}$. (In keeping with the usual convention for passing between the tagged triangulation T and the ordinary triangulation T° , the ordinary arc γ corresponds to the tagged arc that agrees with β but is tagged notched at the interior point of the self-folded triangle.) Second, suppose α is tagged notched at a puncture p inside a self-folded triangle with interior edge β and exterior edge γ . If $e \in \gamma$ is the first point where α crosses an arc of T° after leaving p , then $w(e) = \hat{y}_\beta$.

In many cases, we define $P_\alpha = P_\alpha^{\text{cr}}$. Specifically, in the non-degenerate cases, this happens if and only if both endpoints of α have the following property: Either α is tagged plain at the endpoint or the endpoint is a puncture inside a self-folded triangle.

Before defining P_α in the cases where it is larger than P_α^{cr} , we give an example of the simpler case (Example 3). In this example (as in Example 1 and all further examples), we suppress the actual elements of P_α (which are certain points in \mathbf{S}) and instead show labels that indicate weights: The arcs of T are numbered, and an element of P_α is labeled either with an integer i to indicate a weight \hat{y}_i or with a formal quotient $\frac{i}{j}$ to indicate a weight $\frac{\hat{y}_i}{\hat{y}_j}$.




FIGURE 3. An example where $P_\alpha = P_\alpha^{\text{cr}}$.

Example 3. This example is just like Example 1, except that α is tagged plain at the top-middle puncture in the picture, as shown in left picture of Figure 3. In this case, $\mathbf{g}_\alpha = \frac{x_5 x_6 x_8}{x_1 x_7 x_9}$ and $P_\alpha = P_\alpha^{\text{cr}}$ is shown on the right of Figure 3.

The Hasse diagram of P_α^{cr} is, as a graph, isomorphic to a path. Thus it has two ends, given by the last points where α crosses an arc of T° before its endpoints. At each endpoint p of α that is tagged notched and is not inside a self-folded triangle, we augment P_α^{cr} by adding another chain C . This chain C is constructed in a similar manner to P_α^{cr} , as follows. From a point on α near the endpoint p , we circle around p , crossing all arcs of T° that are incompatible with α because of the disagreement in tagging at the endpoint. The elements of C are the points of crossing, they are ordered according to whether they cut corners left or right just as in P_α^{cr} (necessarily with either all left or all right), and they are labeled just as the points in P_α^{cr} . Let $e \in P_\alpha^{\text{cr}}$ be the last point where α crosses an arc of T° before reaching p . We attach C to P_α^{cr} with the top element of C covering e and the bottom element of C covered by e . (Note that because p is not inside a self-folded triangle, C has at least two elements, so its top and bottom elements are different.) In the simplest case, where α actually is one of the arcs in T , then $\mathbf{g}_\alpha = x_\alpha$ and P_α is the empty poset.

Example 1 revisited. We continue and explain Example 1, which contrasts with Example 3. We take T° and α as shown in the left picture of Figure 1. As explained in Example 2, $\mathbf{g}_\alpha = \frac{x_5 x_6 x_8}{x_1 x_4 x_7 x_9}$. The poset P_α^{cr} is the poset shown in Figure 3. Since α is tagged notched at the top-middle puncture, we adjoin a 4-element chain C , labeled 4, 1, 8, 11 from bottom to top, as shown in the right picture of Figure 1. There is no chain attached to the other end of P_α^{cr} , because that endpoint is the interior vertex of a self-folded triangle.

Remark 1.4. We offer a heuristic explanation for the addition of a chain at an endpoint where α is tagged notched (at a puncture not inside a self-folded triangle). The heuristic is equally good in all such cases, but we explain it using Example 1 and Figure 1. In the surfaces model, when an arc is tagged notched, one often thinks of it as “going around” the puncture rather than ending there. As α approaches the puncture (at the top of Figure 1), we have recorded $7 < 5$ in P_α^{cr} , but instead of stopping there and later adding a chain to make P_α , we make P_α directly by continuing around the puncture (in either direction) and adding elements for the arcs we cross. Thus we record $5 < 11 > 8 > 1 > 4$. Then continuing along α in the opposite direction, we see that the element labeled 4 is covered by the same element as before, labeled 5.




FIGURE 4. An example where P_α is P_α^{cr} plus two chains.

Example 4. We also give an example where neither endpoint of α is inside a self-folded triangle and both endpoints are tagged notched, so that two chains are adjoined to P_α^{cr} to make P_α . For T° and α as shown in Figure 4, we have $\mathbf{g}_\alpha = \frac{x_6}{x_4 x_5}$. The poset P_α^{cr} is a chain labeled $3 < 2 < 1 < 3 < 6$. At the top element of P_α^{cr} , we adjoin a chain labeled $4 < 7$, and the bottom, we adjoin a chain labeled $5 < 4 < 6$, to obtain P_α as shown in Figure 4. The cover relation $4 < 7$ in one adjoined chain “disappears” in the Hasse diagram (hence we show it as a dotted line), because $4 < 6 < 7$ in P_α .

We now consider the degenerate cases, where α coincides, up to tagging, with an arc in T . If α is actually an arc in T , then the \mathbf{g}_α is the initial cluster variable x_α and P_α is the empty poset, whose only order ideal is the empty set, with weight 1. The other degenerate cases are where α coincides with an arc of T but has different taggings.

First, consider the case where α is tagged notched at only one endpoint and otherwise agrees with an arc β of T that is not the interior edge of a self-folded triangle in T° . Then α and β are compatible, so β is not in the chain C coming from arcs of T that are incompatible with α because of tagging at that endpoint. In this case, P_α consists only of the chain C , as illustrated in Example 5.

Example 5. The arc α shown in Figure 5 agrees with arc 1 except for its tagging at one point. Thus P_α is the chain C illustrated in the figure. In this case, $\mathbf{g}_\alpha = \frac{x_2}{x_3}$.

We next consider the case where α coincides with the interior edge of a self-folded triangle in T° . In other words, α coincides with two arcs β, β' of T that are the same up to tagging, but α is tagged differently than both. In this case, since α does not cross any arcs in T° , the poset P_α^{cr} is empty. Exactly one of β and β' is compatible with α , without loss of generality β' . As before, P_α is only the chain C consisting of the arcs that are incompatible with α at its notched endpoint, but in this case, two elements labeled \hat{y}_β appear in the chain, one at the top and one at the bottom, as illustrated in Example 6.

Example 6. In Figure 6, the arc α shown coincides with arc 7 except for tagging. It also coincides with arc 6 except for tagging, when the latter is represented as a tagged arc rather than an ordinary arc. The arc α is compatible with arc 6 and incompatible with arc 7. The poset P_α^{cr} is empty, and P_α is a chain as shown in the figure. In this case, $\mathbf{g}_\alpha = \frac{1}{x_7}$.

Remark 1.5. Note that the heuristic of Remark 1.4 applies equally well to the cases treated in Examples 5 and 6, the only difference being that the poset P_α^{cr} is empty.

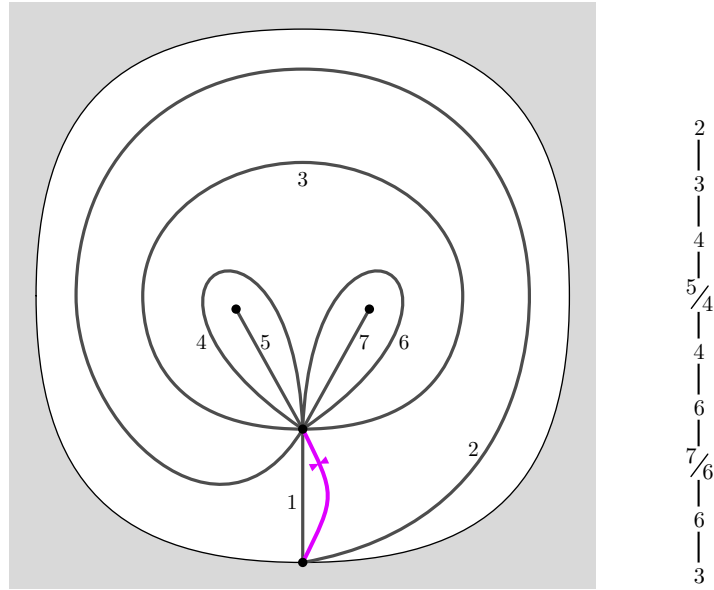


FIGURE 5. An example where P_α^{cr} is empty but a chain is “attached” to it.

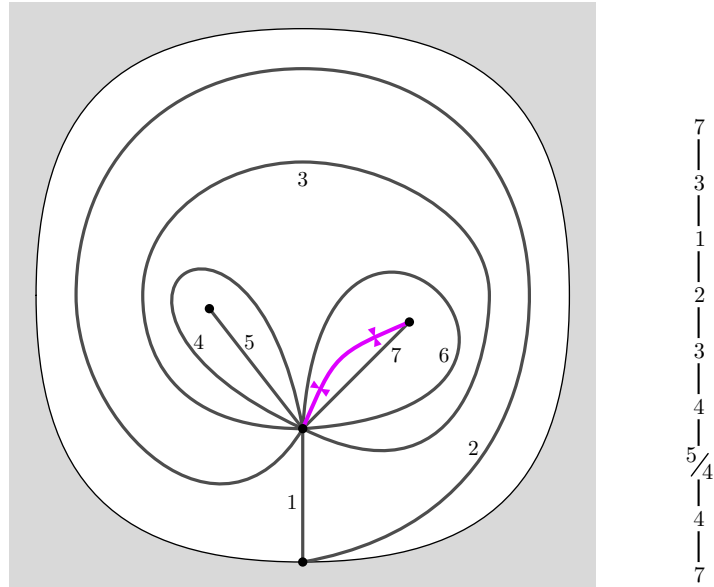


FIGURE 6. An example where α coincides with the interior edge of a self-folded triangle.

Finally, if α agrees, except for tagging, with an arc β of T that is not the interior of a self-folded triangle in T° but is tagged notched at both endpoints, then two chains are “attached” to the empty poset P_α^{cr} . Two elements weighted \hat{y}_β are in each chain, and the chains intersect and interact. Each chain is obtained as before by following a path around the puncture, except that the path starts and ends by crossing β , so that an element weighted \hat{y}_β is at both the top and bottom of the chain. The poset P_α is a union of the two chains, identified at their top and bottom elements, with two additional cover relations, each having the atom in one chain below the coatom in the other chain, as illustrated in Example 7.




FIGURE 7. An example where P_α^{cr} is empty but two chains are “attached” to it.

Example 7. In Figure 7, the arc α coincides with arc 6 except that α is tagged notched at both endpoints (neither endpoint being inside a self-folded triangle). The poset P_α is the union of two chains labeled $6 < 1 < 2 < 6$ and $6 < 3 < 4 < 5 < 6$. The top elements of the chains are identified, as are the bottom elements, and there are additional cover relations as shown. In this case, $\mathbf{g}_\alpha = \frac{1}{x_5}$.

Remark 1.6. We offer a heuristic explanation of the form of P_α in the cases where α agrees with an arc in T (not the interior edge of a self-folded triangle) but is notched at both endpoints. We will give the explanation specifically in the context of Example 7, but the explanation works equally well in general. The idea is to think of P_α as the union of four “polygons”, similar to the polygons that have appeared in earlier examples.

Since the arc α ends at the left puncture with the wrong tagging, it is incompatible with arc 1 of T . Imagine that the arc α , instead of ending at the left puncture, was incompatible with arc 1 by actually crossing arc 1. In this case, P_α^{cr} would have an element labeled 1, to which we would adjoin a chain to make the polygon $6 < 3 < 4 < 5 > 1 > 6$. But also, the notch in α at the left puncture makes it incompatible with arc 2, so, drawing α below arc 6 and letting it cross arc 2 instead of ending at the puncture, we obtain instead the polygon $3 < 4 < 5 < 6 > 2 > 3$. By the same process at the right puncture, we obtain the polygons $1 < 2 < 6 > 5 > 1$ and $6 < 1 < 2 > 3 > 6$. The poset P_α is the union of these four polygons.

We now offer a more complicated example of the case where α agrees with an arc in T (not the interior edge of a self-folded triangle) but is notched at both endpoints.

Example 8. In Figure 8, the arc α coincides with arc 3 except that α is tagged notched at both endpoints (neither endpoint being inside a self-folded triangle). The poset P_α is the union of two chains labeled

$$3 < 2 < 1 < 3 < 6 < 7/6 < 6 < 4 < 5/4 < 4 < 3$$

and

$$3 < 6 < 7/6 < 6 < 4 < 5/4 < 4 < 3 < 2 < 1 < 3,$$

with top elements identified and bottom elements identified, and with additional cover relations as shown. In this case, $\mathbf{g}_\alpha = \frac{1}{x_3}$.

To conclude the discussion of the case where α agrees with an arc in T (not the interior edge of a self-folded triangle) but is notched at both endpoints, we give an example where one of the two chains contains only three elements, so that cover relations “disappear”. This happens when α agrees with an arc in T that is in two triangles of T that combine to form a once-punctured digon. It is impossible for both chains to contain only three elements, because if so, then \mathbf{S} is a sphere with three punctures, which is ruled out in the definition of a marked surface.




FIGURE 8. Another example where P_α^{cr} is empty but two chains are “attached” to it.




FIGURE 9. P_α^{cr} is empty, two chains are “attached” to it, and cover relations “disappear”.

Example 9. In Figure 9, P_α is the union of two chains labeled $1 < 2 < 1$ and $1 < 3 < 4 < 1$ with top elements identified and bottom elements identified. Because of the additional cover relations that connect the atom in one chain to the coatom in the other chain, two cover relations between elements labeled 1 and 2 disappear and the cover relation between elements labeled 3 and 4 also disappears (hence we show them as dotted lines). In this example, $\mathbf{g}_\alpha = \frac{1}{x_1}$.

The cover relation involving 3 and 4 disappears because the chain labeled $1 < 3 < 4 < 1$ has only four elements. In cases where one chain has three elements and the other has more than four, the only cover relations that disappear are those in the chain with three elements.

1.3. A slight reduction of the theorem. We now reduce Theorem 1.1 to the case where α has the following property: When α has an endpoint that is the interior vertex of a self-folded triangle in T° , the tagging of α is plain at that endpoint. This kind of argument is standard in the marked surfaces model.

First, we mention a general combinatorial symmetry that is built into the marked surfaces model [FST08, FT18a]. (Indeed, we have already used this symmetry to assume that the tagged

triangulation T has all arcs tagged plain, except possibly at punctures p incident to exactly two tagged arcs of T with opposite taggings at p .)

Suppose T is a tagged triangulation and suppose p is a puncture. For any tagged arc γ , let γ' be the arc obtained by switching the tagging of γ at p (or leaving γ unchanged if it does not have an endpoint at p). Let T' be the tagged triangulation $\{\gamma' : \gamma \in T\}$. Suppose we are given a tagged arc α and an expression for the cluster variable x_α with principal coefficients at T , in terms of the initial cluster variables x_γ for $\gamma \in T$ and the tropical variables y_γ for $\gamma \in T$. If we take that expression and replace every x_γ with $x_{\gamma'}$ and every y_γ with $y_{\gamma'}$, we obtain an expression for the cluster variable $x_{\alpha'}$ with principal coefficients at T' , in terms of the initial cluster variables $x_{\gamma'}$ for $\gamma' \in T'$ and the tropical variables $y_{\gamma'}$ for $\gamma' \in T'$.

This symmetry exists because the signed adjacency matrix of T and T' are the same by definition (except for being indexed by arcs γ or arcs γ') and because flips of tagged arcs commutes with tag-switching at p . Thus, if we obtain α by a sequence of flips starting with T and write an expression for x_α using exchange relations, the corresponding sequence of flips starting with T' writes the same expression for $x_{\alpha'}$.

Now, suppose T is a tagged triangulation with all arcs tagged plain, except possibly at some punctures p incident to exactly two tagged arcs of T , identical except for opposite taggings at p . Suppose p is one such puncture, so that p is the interior vertex of a self-folded triangle in the corresponding ordinary triangulation T° . If β and γ are the two tagged arcs associated to the self-folded triangle, then the operation of switching tags at such a p simply switches β and γ . (That is, $\beta' = \gamma$ and $\gamma' = \beta$.)

We see that if α is tagged notched at a puncture p that is the interior vertex of a self-folded triangle, then an expression for x_α in terms of the initial cluster variables and principal coefficients at T can be obtained from any such expression for $x_{\alpha'}$ by replacing all instances of x_β by x_γ and vice versa. We check that the right side of Equation (1) has the same symmetry. First, $\kappa(\alpha)$ and $\kappa(\alpha')$ are related by reversing their spirals at p . By the definition of shear coordinates with respect to tagged triangulations [FT18a, Definition 12.1], we see that \mathbf{g}_α and $\mathbf{g}_{\alpha'}$ are related by replacing all instances of x_β by x_γ and vice versa. By inspection of the definition in Section 1.2, we also see that the weighted posets P_α and $P_{\alpha'}$ are related in the same way. (In the non-degenerate case, see Example 3 and in the degenerate case, see Example 6.)

We have established the following reduction of Theorem 1.1.

Proposition 1.7. *Assume that Equation (1) holds for every tagged arc α with the following property: When α has an endpoint that is the interior vertex of a self-folded triangle in T° , the tagging of α is plain at that endpoint. Then Equation (1) holds for every tagged arc.*

2. TIDY EXCHANGE RELATIONS

Given a finite poset P with weights $w : P \rightarrow \{\hat{y}_\beta \mid \beta \in T\} \cup \{\hat{y}_\beta/\hat{y}_\gamma \mid \beta \neq \gamma \in T\}$, write

$$F(P) = \sum_I \prod_{e \in I} w(e),$$

where the sum is over all ideals I in P . (This notation suppresses the dependence of $F(P)$ on the map w .) The content of Theorem 1.1 is that for any tagged arc α , the cluster variable indexed by α has \mathbf{g} -vector monomial \mathbf{g}_α and F -polynomial $F(P_\alpha)$, where the weight function on P_α is w as defined in Section 1.2. By [FT18b, Proposition 5.1], the F -polynomials of cluster variables are uniquely determined by the condition that initial cluster variables have F -vector 1 and the exchange relations [FT18b, (5.3)].

Write F_α for the F -polynomial of a tagged arc α . In this section, we prove that for certain special choices of α and T and an initial arc $\gamma \in T$, the exchange relation that writes $F_\alpha \cdot F_\gamma$ in terms of other F -polynomials F_β also holds with F_α , F_γ , and each F_β replaced by $F(P_\alpha)$, $F(P_\gamma)$ and the corresponding polynomials $F(P_\beta)$. We also show that this exchange relation records a natural decomposition of the set of order ideals of α . Finally, we show that the relationship between $\mathbf{g}_\alpha \cdot \mathbf{g}_\gamma$ and the \mathbf{g}_β is correct.



FIGURE 10. Exchange relations for crossing arcs.

2.1. Exchange relation on F -polynomials. Exchange relations on F -polynomials in general are given in [FT18b, Proposition 5.1]. In this section, we interpret them in the surface model using the principal-coefficients version of the material in [FT18a, Sections 12–13] that realizes geometric coefficients in terms of laminations. (See also [FT18a, Theorem 15.6].)

Throughout the section, we assume that the tagged triangulation T has all taggings plain. Thus also the ordinary triangulation T° has no self-folded triangles. Since in this case T and T° coincide, we will only use the notation T in this section. In this section, we have no need of hyperbolic geometry, but rather we use results of [FT18a] that were proved using hyperbolic geometry.

Given an arc $\gamma \in T$, the *elementary lamination* associated to γ is a curve L_γ that is defined precisely as $\kappa(\gamma)$ was, except with the right and left directions swapped. Thus L_γ agrees with γ except very close to the endpoints, where L_γ turns slightly to the right to end on a boundary segment or spiral counterclockwise into a puncture. (In this section, we only allow plain taggings. In general, if γ is tagged notched at one of its endpoints, then L_γ spirals clockwise into that puncture.) Starting with principal coefficients at T , the extended exchange matrix at another tagged triangulation T' has coefficient rows given by shear coordinates of elementary laminations with respect to T' . Specifically, in the coefficient row indexed by $\gamma \in T$, the entry in the column indexed by $\alpha \in T'$ is $b_\alpha(T', L_\gamma)$, the α -entry of the shear coordinates of L_γ with respect to T' .

Since the coefficients of principal-coefficients extended exchange matrices are determined by shear coordinates of elementary laminations, we can describe exchange relations between F -polynomials geometrically in terms of arcs and elementary laminations. Two distinct tagged arcs admit an exchange relation if and only if one of the following descriptions hold:

- The two arcs intersect exactly once in their interiors and, if they share endpoints, have the same tagging at shared endpoints; or
- The two arcs share an endpoint where their taggings disagree, do not intersect in their interiors, and, if they also share the other endpoint, they have the same tagging there.

Figure 10 illustrates the exchange relation for arcs that intersect in their interior. The arcs being exchanged are shown in purple and are horizontal and vertical. Each picture of Figure 10 shows four more arcs, colored red and blue and forming a quadrilateral whose diagonals are the purple arcs, and some orange curves crossing the quadrilateral that represent elementary laminations associated to arcs of T . There are some complications in general that we are able to avoid here by assuming that the purple arcs have, between them, four distinct endpoints. (Otherwise, some of the blue or red arcs may coincide or some red or blue arc may need to be replaced, in exchange relations, by the product of two tagged arcs that coincide except for the tagging at one endpoint.) There is one complication that we will not be able to avoid: If any endpoints of purple arcs are tagged notched, then the orange curves agree with the elementary laminations for the arcs of T , except that all spirals at those endpoints are reversed.

Elementary laminations (possibly with some spirals reversed) play a role in the exchange relation if they cross a blue arc, then the purple arcs, then the other blue arc, or if they similarly cross a red arc, then the purple arcs, then the other red arc. Since we have assumed that T has no notched taggings, the elementary laminations associated to arcs in T can be taken to be disjoint from each other. (This is *almost* true in any case; in general, two of these elementary laminations only intersect when they are associated to a pair of arcs that coincide except for tagging at one endpoint.) Because the elementary laminations are disjoint from each other, it is impossible to have some laminations crossing blue-purple-blue and at the same time other laminations crossing red-purple-red. (In the general case, this is also true, for only slightly more complicated reasons.) Since we have the choice of which pair of arcs to color red and which to color blue, *we take*




FIGURE 11. Exchange relations for arcs with disagreeing tags.

the convention that the orange curves only have blue-purple-blue crossings, and the pictures in Figure 10 are colored accordingly.

We adopt the notation F_{purple} for the product of the F -polynomials of two purple arcs shown in Figure 10, and similarly F_{blue} and F_{red} . Each orange curve shown is associated to an arc $\gamma \in T$, and \hat{y}_{orange} refers to the product of the variables \hat{y}_γ for all $\gamma \in T$ whose associated orange curve has a blue-purple-blue intersection with the quadrilateral. In general, a factor \hat{y}_γ can appear in \hat{y}_{orange} more than once, but we will consider those exchange relations where each \hat{y}_γ appears at most once. In light of our coloring convention, the *F -polynomial exchange relation* is

$$(3) \quad F_{\text{purple}} = F_{\text{blue}} + \hat{y}_{\text{orange}} \cdot F_{\text{red}}.$$

Figure 11 illustrates the exchange relation for arcs that share a vertex where their taggings disagree. Each picture of the figure shows the two arcs (in purple and horizontal) being exchanged and a red and blue arc. We avoid complications by assuming that the two purple arcs have three distinct endpoints. The same rule on reversing spirals of elementary laminations applies when non-shared endpoints of the purple arcs are tagged notched. The notation F_{purple} again refers to the product of the F -polynomials of two purple arcs shown. The figure also shows dotted red and blue arcs (coinciding with one of the purple arcs). The notation F_{blue} refers to the F -polynomial of the one *non-dotted* blue arc and F_{red} refers to the F -polynomial of the one *non-dotted* red arc. The dotted red and blue arcs are an aid to understanding the orange curves, specifically because we care about blue-purple-blue or red-purple-red crossings, with one of these blues or reds being dotted. It is again impossible to have both laminations crossing blue-purple-blue and laminations crossing red-purple-red, so we again take the convention that orange curves have only blue-purple-blue crossings. The notation \hat{y}_{orange} again refers to the product of the variables \hat{y}_γ for all $\gamma \in T$ whose associated orange curve crosses blue-purple-blue with one of these blues being dotted. (Again, in general, a factor \hat{y}_γ can appear in \hat{y}_{orange} more than once, but we will avoid that case.) Once again, the exchange relation is given by (3).

Remark 2.1. Equation (3) gives the exchange relation on F -polynomials, interpreting [FT18b, Proposition 5.1] in terms of elementary laminations. Later, in the proof of Theorem 1.1, we will also need the principal-coefficients case of the exchange relation on cluster variables, which can be obtained from Equation (3) by replacing each F -polynomial by the corresponding cluster variable and replacing \hat{y}_{orange} by y_{orange} , the product of the y_γ for all $\gamma \in T$ whose associated orange curve has a blue-purple-blue intersection.

2.2. Tidyness. We now describe conditions on T and α that turn certain exchange relations into decompositions of the set of order ideals of P_α . We have not tried to minimize these conditions, because we can easily satisfy them using the covering arguments in later sections. We assume as usual that an isotopy representative of α has been chosen to minimize intersections with arcs of T .

We say that (T, α) is *tidy* if:

- T consists of plain-tagged arcs;
- each triangle of T has three distinct vertices and three distinct edges (so that in particular each arc of T has two distinct endpoints);
- α has two distinct endpoints;
- no arc in T intersects the interior of α more than once;
- no arc in T both intersects the interior of α and shares an endpoint with α ;
- if an arc of T has the same two endpoints as α , then α coincides with that arc except possibly for tagging.




FIGURE 12. Exchange relations for tidy arcs and initial arcs.

This tidiness condition implies the following simplifications:

- T has no self-folded triangles, so the labels of P_α are all of the form \hat{y}_γ for $\gamma \in T$.
- There are no repeated labels in P_α , except in the case where α coincides with an arc β (not the interior edge of a self-folded triangle) and α has both ends notched, so that the label \hat{y}_β appears twice (as the minimal and maximal element of P_α , see *e.g.* Figure 7).
- An arc $\gamma \in T$ is exchangeable with α if and only if it crosses α or it shares with α an endpoint notched in α (but doesn't coincide with α up to tagging). In other words, there is an exchange relation exchanging x_α and x_γ if and only if \hat{y}_γ appears exactly once as a label in P_α .
- For any arc $\gamma \in T$, the corresponding elementary lamination contributes at most once to \hat{y}_{orange} .

We now consider the exchange relations for exchanging arcs α and γ in the special case when (T, α) is tidy and $\gamma \in T$ is exchangeable with α . The top picture of Figure 12 shows the case where α and γ cross in their interiors (α is the horizontal purple arc, while γ is the vertical purple arc). Also pictured (in gray) are the arcs of T that α crosses or shares an endpoint with, and (in black) additional arcs that complete some triangles in T , and (in orange) the elementary lamination L_γ (which fixes our color choices according to the convention that it should cross the quadrilateral as blue-purple-blue). In this case, it is impossible for α to coincide with an arc of T except for tagging, because if so, α could not cross another arc γ of T . It is possible that α is tagged notched at either or both endpoints, in which case, the red and blue arcs at those endpoints are also notched. However, since (T, α) is tidy, the endpoints of γ are not tagged notched.

The bottom picture of Figure 12 shows the case where α and γ share an endpoint where they disagree on tagging, with additional grey and black arcs and orange curve as in the top picture. In this case, it is possible that α is also tagged notched at the endpoint it does not share with γ , in which case the red and blue arcs at that endpoint are also notched. Again in this case, the endpoints of γ are not tagged notched.

It is important to consider the question of how general the pictures in Figure 12 are. Here are the ways that (T, α) can differ from the pictures in the tidy case.

- The sequence of triangles between endpoints of α can vary. (This variation corresponds to the ways that P_α^{cr} can be any poset whose Hasse diagram, as a graph, is a path.)
- One or both of the endpoints of α may be on the boundary. When an endpoint p of α is on the boundary, one or more of the triangles incident to p are deleted and correspondingly two of the arcs of T incident to p become boundary segments. In this case, α is tagged plain at p , so the arcs incident to p are irrelevant to the definition of P_α .
- The number of triangles incident to each endpoint p of α may vary (at least 2 when p is a puncture and at least 1 when p is on the boundary).
- Some black edges may be identified with each other and/or some marked points (excluding the endpoints of α) may be identified with each other.




FIGURE 13. Exchange relations for tidy arcs and initial arcs, with digons.

- One or both endpoints of α may be the puncture in a once-punctured digon formed as the union of two triangles of T . Figure 13 is a version of Figure 12 with these digons at both endpoints of α in each picture. The case with digons is not qualitatively different, but instead simply achieves the lower-bound on the number of arcs in T incident to endpoints of α . (For that reason, we will continue to draw the cases without digons, as in Figure 12.)
- α may coincide with an arc of T .

2.3. Tidy exchange relation on posets. Assume now that (T, α) is tidy and consider $\gamma \in T$ exchangeable with α . We define P_{red} to be the weighted poset P_β when there is a single red arc β , or the disjoint union of the two weighted posets P_β and $P_{\beta'}$ when there are two red arcs β and β' in the exchange relation. Note that $F(P_{\text{red}})$ is then the product of the $F(P_\beta)$ for the (one or two) red arcs in the exchange relation. We define P_{blue} similarly. (We could also define P_{purple} , but since P_γ is the empty poset, P_{purple} would always equal P_α .)

Proposition 2.2. *Suppose (T, α) is tidy and $\gamma \in T$ is exchangeable with α , and denote by e_γ the element of P_α labeled \hat{y}_γ . Then*

- P_{blue} is obtained from P_α by deleting all elements weakly above e_γ ,
- P_{red} is obtained from P_α by deleting all elements weakly below e_γ ,
- \hat{y}_{orange} is the product of the weights of the elements weakly below e_γ in P_α .

Hence, the weighted sum $F(P_\alpha)$ of order ideals in P_α decomposes into

- the weighted sum $F(P_{\text{blue}})$ of order ideals in P_α that do not contain e_γ and
- the weighted sum $\hat{y}_{\text{orange}} \cdot F(P_{\text{red}})$ of order ideals in P_α that contain e_γ .

This yields the **poset exchange relation**

$$(4) \quad F(P_\alpha) = F(P_{\text{blue}}) + \hat{y}_{\text{orange}} \cdot F(P_{\text{red}}).$$

Example 10. To illustrate Proposition 2.2 and its proof, we give a series of examples in Figures 14–21. Each figure has purple, red, and blue arcs as in Figure 12 or Figure 13. The top-left picture shows the arcs of T that intersect the purple arc α , colored orange if the corresponding laminations (with spirals reversed as appropriate) have blue-purple-blue intersections and colored gray otherwise. The top-right picture shows the corresponding laminations, with the same color coding. The purple arc γ that is being exchanged with α will always appear orange in the pictures, so we identify γ explicitly in the figures. The arc γ can also be located using the positions of the red and blue arcs. In other cases where an arc should be two colors (either red and orange or purple and orange), both colors are shown side by side. Below the pictures, we identify γ , give the monomial \hat{y}_{orange} , and picture the posets P_α , P_{blue} , and P_{red} .




FIGURE 14. First example of Proposition 2.2.




FIGURE 15. Second example of Proposition 2.2.

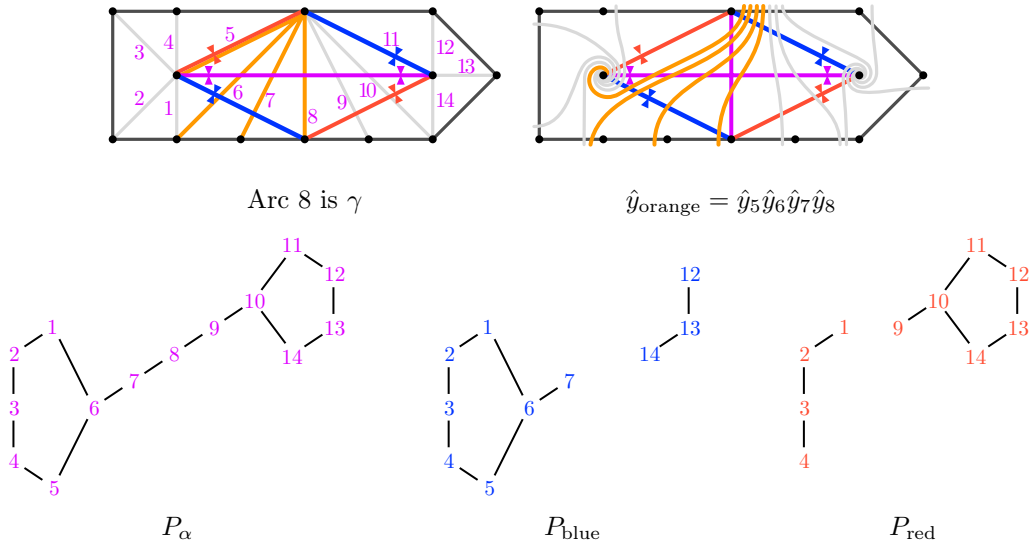


FIGURE 16. Third example of Proposition 2.2.




FIGURE 17. Fourth example of Proposition 2.2.




FIGURE 18. Fifth example of Proposition 2.2.




FIGURE 19. Sixth example of Proposition 2.2.




FIGURE 20. Seventh example of Proposition 2.2.




FIGURE 21. Eighth example of Proposition 2.2.

Proof of Proposition 2.2. We continue the notation e_δ for the element of P_α labeled \hat{y}_δ , for $\delta \in T$. In the one case where there are two elements of P_α labeled \hat{y}_δ , namely the case where α coincides with δ except that α has both endpoints tagged notched, we take e_δ to mean the bottom element of P_α , not the top element.

For the statement about \hat{y}_{orange} , we need to show that, for every arc $\delta \in T$, we have $e_\delta \leq e_\gamma$ if and only if the lamination associated to δ (with spirals reversed where it shares an endpoint with α that is notched in α) has a blue-purple-blue intersection, or equivalently, \hat{y}_δ is a factor of \hat{y}_{orange} . It will also be clear that these factors appear with multiplicity 1. (In the case where α coincides with δ except that α has both endpoints notched, the top element of P_α labeled \hat{y}_δ is of course never below e_γ .) One of the purple arcs is α , so if the lamination associated to $\delta \in T$ has a blue-purple-blue intersection, then in particular δ has some intersection with α , so we will

consider only arcs of T that intersect α . It is impossible for α to coincide with γ except for tagging, because in that case \hat{y}_γ occurs as a label in P_α either 0 or 2 times. This contradicts the hypothesis that (T, α) is tidy and $\gamma \in T$ is exchangeable with α .

For the statement about P_{red} , suppose ρ is one of the red arcs. We can choose an isotopy representative of ρ such that every element e_δ of P_ρ is also an element of P_α . Furthermore, two elements of P_ρ form a cover relation in P_ρ if and only if they form a cover relation in P_α . If ρ and σ are the two red arcs, then P_ρ and P_σ are disjoint subsets of P_α , and P_{red} is their disjoint union. We will show that each e_δ appears in P_{red} if and only if $e_\delta \not\leq e_\gamma$ in P_α . When we have showed that, since we obtain P_{red} from P_α by deleting a lower order ideal, and since the cover relations in P_{red} are the same as in P_α , we can conclude that the partial order on P_{red} is the restriction of the partial order on P_α , as desired.

There are two main cases to consider: The arc γ may cross α or it may share an endpoint with α where α is tagged notched. In either case, if $\delta \in T$ does not intersect α or if δ shares an endpoint of α where α is tagged plain, then neither P_α nor P_{red} contain an element e_δ , and the associated lamination can only make red-purple-blue intersections. See Figures 14, 15, 17, and 18. Thus we will not consider such arcs δ below.

In both cases, and in every case for δ , we give a simple condition on δ that is easily equivalent to $e_\delta \leq e_\gamma$, to the lamination associated to δ having a blue-purple-blue intersection, and to e_δ not being an element of P_{red} . To avoid unnecessary repetition of words, in every case, we simply state below what that condition is.

Case 1. γ crosses α . In this case, α does not coincide with an arc of T .

Subcase 1.1. δ crosses α . See Figures 14, 15, and 16. As we move from e_γ along α to e_δ , each arc that we cross has a left and right endpoint. The condition is that γ and δ share their right endpoints and each arc of T crossed in between γ and δ also shares that right endpoint.

Subcase 1.2. δ shares an endpoint p with α such that α is tagged notched at p . See Figure 16. Let t be the triangle in T that has p as a vertex and contains part of α in its interior. (This t is bounded by the last arc of T crossed by α and by two arcs of T that share the endpoint p with α .) The condition is that all three of the following statements hold: every arc crossed by α between e_γ and p shares its right endpoint with γ ; the arc δ is an edge of t ; and δ shares its right endpoint with γ . (The notion of “right endpoint” for δ is from the point of view of someone moving along α from e_γ to p . Thus p is the *left* endpoint of δ .) Equivalently, there is a much simpler condition: δ coincides with the red arc whose endpoint is p , except that that red arc is notched at p where δ is plain. (To consider blue-purple-blue intersections in this case, the reversal of spirals at the notched endpoint is crucial.)

Case 2. γ shares an endpoint p with α where α is tagged notched. This includes the cases where α coincides with δ except for tagging, and in those cases, no arcs in T cross α .

Subcase 2.1. δ shares the endpoint p . See Figures 17, 18, 19, 20, and 21. The condition is that one can move from γ to δ clockwise around p without crossing α .

We clarify the condition in cases where δ coincides with α except for tagging: We have already supposed that α is tagged notched at p . If α coincides with δ and is tagged plain at its other endpoint, then there is no element e_δ in P_α or P_{red} . If α is tagged notched at both ends and otherwise coincides with δ , recall that the notation e_δ means the bottom element of P_α , not the top element. (In any case, the top element is not below e_γ .) In this case, we consider the condition to be fulfilled (because, moving from γ to δ clockwise around p , we arrive at α but do not cross it).

Subcase 2.2. δ crosses α . See Figures 17, 18, and 19. Again, arcs crossed by α have a left endpoint and a right endpoint from the point of view of someone moving along α from e_δ to p . Let t be the triangle in T that has p as a vertex and contains part of α in its interior. The condition is that the following three statements hold: the arc γ is an edge of t ; every arc crossed by α between e_δ and p shares its left endpoint with γ ; and δ shares its left endpoint with γ . (The “left endpoint” of γ is the other endpoint, not p .)




FIGURE 22. The γ -entry of \mathbf{g}_α and \mathbf{g}_{blue} , when γ crosses α , Case 1 in the proof of Proposition 2.3.1.

Subcase 2.3. α is also tagged notched at its other endpoint q and δ shares the endpoint q with α . See Figures 18, 19, 20, and 21. Let u be the triangle of T having q as a vertex and containing part of α in its interior. We use the terms left and right from the point of view of someone moving along α from q to p . The condition is that the following four statements hold: the arc γ is an edge of t ; the arc δ is an edge of u ; every arc of T crossed by α (if there are any) shares its left endpoint with γ ; and δ shares its left endpoint with γ . (The left endpoint of δ is the other endpoint, not q .)

This concludes our proof of the statement about P_{red} and the statement about \hat{y}_{orange} . Reversing the orientation of \mathbf{S} exchanges the red arcs with the blue arcs and replaces the partial order P_α with the dual partial order. The statement about P_{blue} thus follows from the statement about P_{red} and this orientation-reversal symmetry. Finally, decomposing the order ideals of P_α depending on whether or not they contain e_γ yields the poset exchange relation of Equation (4). \square

2.4. \mathbf{g} -vectors in tidy exchange relations. We next describe the relation among the \mathbf{g} -vectors appearing in the exchange. We again assume that (T, α) is tidy and consider $\gamma \in T$ exchangeable with α . Write \mathbf{g}_{blue} for the product of monomials \mathbf{g}_β associated to the blue arcs, and similarly \mathbf{g}_{red} . The notation of specialization to “ $y = 1$ ” in the following proposition is shorthand for setting every tropical variable to 1.

Proposition 2.3. *If (T, α) is tidy and $\gamma \in T$ is exchangeable with α , then*

1. $\mathbf{g}_{\text{blue}} = x_\gamma \cdot \mathbf{g}_\alpha$ and
2. $\mathbf{g}_{\text{red}} = x_\gamma \cdot \mathbf{g}_\alpha \cdot \hat{y}_{\text{orange}}|_{y=1}$.

The proof of Proposition 2.3 has a similar feel as that of Proposition 2.2. We separate the proofs of the two items.

Proof of Proposition 2.3.1. We show that $\mathbf{g}_{\text{blue}} = x_\gamma \cdot \mathbf{g}_\alpha$, in two cases, depending on whether γ crosses α or shares an endpoint of α where α is tagged notched. In either case, let Q be the quadrilateral consisting of γ and the two triangles of T having γ as an edge.

Case 1. γ crosses α . Write β and β' for the two blue arcs. The curve $\kappa(\beta)$ follows $\kappa(\alpha)$ from near one endpoint of α through the same intersections with triangles of T until they both enter Q through the same side of Q and then both cross γ before $\kappa(\beta)$ ends near an endpoint of γ and $\kappa(\alpha)$ continues. The curve $\kappa(\beta')$ has the same relationship with $\kappa(\alpha)$, but starting from the other endpoints. We see that the shear coordinates of $\kappa(\beta)$ and $\kappa(\beta')$ add up to a vector that agrees with the shear coordinates of $\kappa(\alpha)$ except at the γ -entry. It remains only to check, in all four subcases of how $\kappa(\alpha)$ can cross Q , that the γ -entry of the sum of the shear coordinates of $\kappa(\beta)$ and $\kappa(\beta')$ is one less than the γ -entry of the shear coordinates of $\kappa(\alpha)$. This simple check is shown in Figure 22. In the figure, $\kappa(\alpha)$ and the contribution to its shear coordinates at γ are shown in purple, while $\kappa(\beta)$ and $\kappa(\beta')$ and the contributions to their shear coordinates at γ are shown in blue.

Case 2. γ shares an endpoint of α where α is tagged notched. Let β be the one blue arc. Then $\kappa(\beta)$ follows $\kappa(\alpha)$ from near the other endpoint (not the shared endpoint) of α through the same intersections with triangles of T until they both enter Q through the same side of Q . At that point, $\kappa(\beta)$ turns right and leaves Q just to the left of the other endpoint (not the shared endpoint) of γ and either hits a boundary segment or spirals clockwise into that endpoint, while $\kappa(\alpha)$ turns left and spirals counterclockwise into the shared endpoint. Thus the shear coordinates of $\kappa(\beta)$




FIGURE 23. The γ -entry of \mathbf{g}_α and \mathbf{g}_{blue} , when γ shares an endpoint with α , notched in α , Case 2 in the proof of Proposition 2.3.1.

agree with the shear coordinates of $\kappa(\alpha)$ except at the γ -entry. We check that the γ -entry of the shear coordinates of $\kappa(\beta)$ is one *less* than the γ -entry of the shear coordinates of $\kappa(\alpha)$. The check is again simple and is shown in Figure 23, where there are two subcases corresponding to the two possible sides through which $\kappa(\beta)$ and $\kappa(\alpha)$ might enter Q . \square

Proof of Proposition 2.3.2. To show that $\mathbf{g}_{\text{red}} = x_\gamma \cdot \mathbf{g}_\alpha \cdot \hat{y}_{\text{orange}}|_{y=1}$, it is convenient to rewrite the identity as $\mathbf{g}_{\text{red}}/\mathbf{g}_\alpha = x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$.

Given an orange arc (an arc δ such that $e_\delta \leq e_\gamma$ in P_α), there are two triangles of T containing δ as an edge. Since $\hat{y}_\delta = y_\delta \prod_{\beta \in T} x_\beta^{b_{\beta\delta}}$, the Laurent monomial $\hat{y}_\delta|_{y=1}$ is the product $x_\beta x_{\beta'}$ for the two edges β and β' (one in each triangle) that are counterclockwise of δ , divided by the similar product of the two edges that are clockwise, except that if any of these edges are boundary segments, then we leave them out. Thus $\hat{y}_{\text{orange}}|_{y=1}$ is the product of these contributions, over all orange edges.





Because the exponent vector of \mathbf{g}_α is the negative of the shear coordinates of $\kappa(\alpha)$ and \mathbf{g}_{red} is the negative of the (sum of the) shear coordinates of κ applied to the red arc(s), we can rephrase the assertion that $\mathbf{g}_{\text{red}}/\mathbf{g}_\alpha = x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$ in terms of shear coordinates. Specifically, taking the shear coordinates of $\kappa(\alpha)$ minus the (sum of the) shear coordinate(s) of κ applied to the red arc(s), the entry at each arc must equal the corresponding exponent in $\hat{y}_{\text{orange}}|_{y=1}$, plus 1 in the entry for γ .

We will verify that this equality between the difference of shear coordinates and the exponents on $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$. We again argue in two cases, depending on whether γ crosses α or shares an endpoint of α where α is tagged notched.

Case 1. γ crosses α . From the proof of Proposition 2.2 (Case 1), we see that every orange arc δ shares an endpoint with γ . Indeed, for each endpoint p of γ , there is a sequence (possibly a singleton) containing γ and all arcs of T that are clockwise of γ at p up to a particular arc (not skipping any arcs). The sequence does not extend beyond the red arc, but its last entry might be the red arc (in which case, the red arc and α are tagged notched at the other end of the red arc, opposite from p). If δ is this last orange arc at p , consider the triangle t that has δ as an edge and is clockwise of δ at p . Let η be the edge of t that shares the endpoint p with δ and let ζ be the other edge. Similarly, let η' and ζ' be the analogous arcs associated with the other endpoint of γ .

There are three subcases, depending on whether there is a non-singleton sequence of arcs at neither end of γ , at one end, or at both. Figure 24 shows, in these three cases, all of the contributions, positive and negative, to the exponents of cluster variables in $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$. These contributions are on the orange arcs and the arcs η , ζ , η' , and ζ' .

The exponent of x_η in $\hat{y}_{\text{orange}}|_{y=1}$ is 1 and the exponent of x_ζ is -1 . Since δ is the last orange arc, $\kappa(\alpha)$ does not cross η , but rather enters the triangle t through ζ . (If δ coincides with the red arc, we are using the fact that the other end of the red arc is notched.) There are two possibilities for where $\kappa(\alpha)$ comes from before crossing ζ . Figure 25 shows that the shear coordinates of $\kappa(\alpha)$ minus the sum of the shear coordinates of κ applied to the red arcs has the correct entry at ζ and η . (It is possible that η is a boundary segment, in which case the contribution to η disappears both in the shear coordinates and in $\hat{y}_{\text{orange}}|_{y=1}$.) The same is true replacing ζ and η with ζ' and η' . Since $\kappa(\alpha)$ coincides with κ of one of the red arcs until they both cross ζ , they have the same shear coordinates for every arc of T crossed by $\kappa(\alpha)$ before it crosses ζ .

FIGURE 24. Exponents in $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$, Case 1 in the proof of Proposition 2.3.2.FIGURE 25. Shear coordinates at ζ and η , Case 1 in the proof of Proposition 2.3.2.FIGURE 26. Shear coordinates of α on the orange arcs, Case 1 in the proof of Proposition 2.3.2.FIGURE 27. Exponents in $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$ and shear coordinates of α and the red arc, Subcase 2.1 in the proof of Proposition 2.3.2.

It remains only to show that the shear coordinates of $\kappa(\alpha)$ minus the sum of the shear coordinates of κ applied to the red arcs has the correct entries at the orange arcs. Indeed, κ applied to the red arcs has is disjoint from the orange arcs, so the point is the shear coordinates of $\kappa(\alpha)$. This shear coordinate computation is shown in Figure 26, in the same three cases as Figure 24, and with the same result.

Case 2. γ shares an endpoint of α where α is tagged notched. Again, the proof of Proposition 2.2 (this time Case 2), indicates that every orange arc δ shares an endpoint with γ . Let p be the endpoint common endpoint of γ and α (where α is notched) and let q be the other endpoint of γ (not of α). There are five subcases.

Subcase 2.1. γ is the only orange arc. This case is illustrated in Figure 27. The left picture of the figure shows all contributions to the exponent vector of $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$. The curve $\kappa(\alpha)$ crosses the curve labeled ζ after leaving p and there are two possibilities for where $\kappa(\alpha)$ goes after crossing ζ . The middle and right picture verify, for these two possibilities, that the shear coordinates of $\kappa(\alpha)$ minus the shear coordinates of κ of the red arc have entries that agree with the exponent vector




FIGURE 28. Exponents in $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$ and shear coordinates of α and the red arc, Subcase 2.2 in the proof of Proposition 2.3.2.




FIGURE 29. Exponents in $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$ and shear coordinates of α and the red arc, Subcase 2.3 in the proof of Proposition 2.3.2.




FIGURE 30. Exponents in $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$ and shear coordinates of α and the red arc, Subcase 2.4 in the proof of Proposition 2.3.2.

of $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$. The configuration shown in the right picture includes the possibility that α coincides with an arc of T except for being notched at one side.

Subcase 2.2. All orange arcs share the endpoint p and there are at least 2 non-orange arcs at p . This case is illustrated and verified in Figure 28, where the pictures are analogous to those in Figure 27. In this case, the configuration shown in the middle picture includes the possibility that α coincides with an arc of T except for being notched at both sides and the right picture again includes the possibility that α coincides with an arc of T except for being notched at one side.

Subcase 2.3. All orange arcs share the endpoint p and there is exactly 1 non-orange arc at p . This case is illustrated and verified in Figure 29, with pictures analogous to previous cases. Again, the middle picture includes the possibility that α coincides with an arc of T except for being notched at both sides and the right picture includes the possibility that α coincides with an arc of T except for being notched at one side.

Subcase 2.4. All arcs at p are orange and there is exactly one orange arc at q aside from γ . This case is Figure 30. The arc γ and the other orange arc with an endpoint q bound a triangle whose third side is also orange. The middle picture includes the possibility that α coincides with an arc of T except for being notched at both sides.

Subcase 2.5. All arcs at p are orange and there is more than one orange arc at q aside from γ . This case is Figure 31, which is like the previous figures, except that the top picture shows the exponent vector of $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$ and the bottom two pictures show the shear coordinates of $\kappa(\alpha)$ and of κ applied to the red arc. \square




FIGURE 31. Exponents in $x_\gamma \cdot \hat{y}_{\text{orange}}|_{y=1}$ and shear coordinates of α and the red arc, Subcase 2.5 in the proof of Proposition 2.3.2.

3. HYPERBOLIC GEOMETRY AND THE COEFFICIENT-FREE CASE

In this section, we prove the analog of Theorem 1.1 in the coefficient-free case. This is a specialization of Theorem 1.1 by setting all of the tropical variables y_γ to 1. The coefficient-free case is not used in the proof of Theorems 1.1 in Section 4, but rather serves to illustrate the remaining key idea of the proof using the simpler of the two hyperbolic-geometric settings of [FT18a]. We continue the notation (\mathbf{S}, \mathbf{M}) for a marked surface and T for a fixed tagged triangulation, again assumed to have all arcs tagged plain, except possibly at some punctures p incident to exactly two tagged arcs of T , identical except for opposite taggings at p . We also continue the notation T° for the ordinary triangulation corresponding to T .

The coefficient-free case of the theorem associates the same poset P_α to a tagged arc α , but with a different weighting. For $e \in P_\alpha$, we define the weight $\tilde{w}(e)$ to be the monomial obtained from $w(e)$ by setting all tropical variables y_γ to 1. Thus when $w(e)$ is \hat{y}_γ for some $\gamma \in T$, then $\tilde{w}(e) = \prod_{\delta \in T} x_\delta^{b_\delta \gamma}$. When $w(e) = \frac{\hat{y}_\beta}{\hat{y}_\gamma}$, then $\tilde{w}(e) = 1$ because in this case $b_\delta \beta = b_\delta \gamma$ for all $\delta \in T$.

Theorem 3.1. *The coefficient-free cluster variable associated to a tagged arc α is*

$$(5) \quad x_\alpha = \mathbf{g}_\alpha \cdot \sum_I \prod_{e \in I} \tilde{w}(e),$$

where the sum is over all order ideals I in P_α .

The entire discussion in Section 1.3 holds in the coefficient-free case verbatim (but omitting the tropical variables y_γ or setting them all to 1), so in particular we have the following version of Proposition 1.7.

Proposition 3.2. *Assume that Equation (5) holds for every tagged arc α with the following property: When α has an endpoint that is the interior vertex of a self-folded triangle in T° , the tagging of α is plain at that endpoint. Then Equation (5) holds for every tagged arc.*

3.1. Lambda lengths. In preparation for the proof of Theorem 3.1, we explain how cluster variables can be realized in terms of certain hyperbolic lengths. (Proofs of statements left unjustified here are found in [FT18a, Chapters 7–8].) More precisely, the initial cluster variables can be taken as indeterminates (constrained to be positive) whose values determine certain hyperbolic lengths in the surface and thus determine a hyperbolic metric and a choice of certain curves called horocycles. Although it is the initial cluster variables that determine the metric, it is more pleasant to explain the construction in the reverse order, starting from a metric, choosing horocycles, and measuring certain lengths that determine the initial cluster variables.

Given (\mathbf{S}, \mathbf{M}) and T , choose a finite-area hyperbolic metric with constant curvature -1 on $\mathbf{S} \setminus \mathbf{M}$ with a cusp at each point in \mathbf{M} , with the property that each boundary segment is a geodesic. Since

the arcs in T° are defined up to isotopy, we can also assume that they are geodesics. At each marked point p , choose a **horocycle** h_p , meaning a curve that is orthogonal to every geodesic that limits to p . This hyperbolic metric together with the choice of horocycles is a **decorated hyperbolic structure** on (\mathbf{S}, \mathbf{M}) . Given a geodesic γ that limits to marked points p and q , the distance along the geodesic between h_p and h_q is finite (in contrast to the distance along the geodesic between p and q). The **signed length** $l(\gamma)$ is this finite distance, counted as positive if the segment of γ between h_p and h_q is outside the disks they define, zero if h_p , h_q , and γ all intersect at the same point, and otherwise negative if the segment of γ between h_p and h_q is inside the disks they define. The **lambda-length** of γ is defined to be $\lambda(\gamma) = \exp(l(\gamma)/2)$.

The first key result about lambda lengths [FT18a, Theorem 7.4] is that an arbitrary assignment of a positive number to each arc in T° and to each boundary segment of (\mathbf{S}, \mathbf{M}) uniquely determines a decorated hyperbolic structure such that the arcs and boundary segments are geodesics and the assigned numbers are the lambda lengths. (In the language of [FT18a, Theorem 7.4], there is a bijection from the set of such assignments to the **decorated Teichmüller space**, the space of decorated hyperbolic structures. Indeed, [FT18a, Theorem 7.4] says that the map is a homeomorphism.)

When p is a puncture, the horocycle h_p has a finite positive length $L(h_p)$ in the hyperbolic metric, and $L(h_p)$ specifies h_p uniquely among horocycles at p . The **conjugate horocycle** \bar{h}_p is the unique horocycle at p whose length satisfies $L(\bar{h}_p)L(h_p) = 1$. The **lambda length** of a **tagged arc** γ is defined just as in the untagged case, except that at each endpoint p where γ is tagged notched, we use the conjugate horocycle \bar{h}_p in place of h_p .

The lambda lengths of the tagged arcs in T determine the lambda lengths of the ordinary arcs in T° and vice versa. Specifically, suppose β and γ are tagged arcs in T that coincide except that their taggings are both plain at one endpoint q but disagree at the other endpoint p , and suppose that δ is the ordinary arc in T° with both endpoints at q and follows β and γ closely to form a loop around p . Then

$$(6) \quad \lambda(\delta) = \lambda(\beta)\lambda(\gamma).$$

(This is [FT18a, Lemma 8.2].) Thus in particular, we can arbitrarily choose positive lambda lengths x_γ for the tagged arcs in $\gamma \in T$ and for the boundary segments to uniquely determines a hyperbolic metric with constant curvature -1 and a choice of horocycles. The chosen lambda lengths can be thought of as indeterminates which then determine the lambda lengths of all other tagged arcs. Taking these chosen lambda lengths to be the initial cluster variables for an exchange pattern with initial exchange matrix $B(T)$, the cluster variables in the exchange pattern are indexed by tagged arcs in (\mathbf{S}, \mathbf{M}) . Each tagged arc has a geodesic representative whose lambda length is the associated cluster variable. This is part of [FT18a, Theorem 8.6]. That theorem does not take coefficient-free cluster algebras but rather boundary coefficients, but we can recover the coefficient-free case by setting the lambda length of each boundary component to be 1.

We can now interpret Theorem 3.1 as the assertion that a given tagged arc α in (\mathbf{S}, \mathbf{M}) has lambda length $\mathbf{g}_\alpha \cdot \sum_I \prod_{e \in I} \tilde{w}(e)$, where as always the sum is over all order ideals I in P_α .

3.2. The tile cover of a tagged arc. The key to Theorem 3.1 is to construct a tagged arc α' in a new surface $(\mathbf{S}', \mathbf{M}')$, equipped with a triangulation T' and a decorated hyperbolic structure such that the lambda length $\lambda(\alpha')$ is the same as $\lambda(\alpha)$, but with simpler combinatorics (that is, so that (T', α') is tidy in the sense of Section 2.2). This will allow us to use Proposition 2.2 in an inductive proof of the theorem.

A given triangle t of T° may not have three distinct edges and three distinct vertices. Define a **tile** of t to be an ideal triangle with three distinct edges and three distinct vertices and with a hyperbolic metric such that the interior of the tile is isometric to the interior of t . The metric on the interior of the tile uniquely determines a metric on the edges, and since we have assumed that the arcs in T° are geodesic, the edges of the tile are also geodesic. The tile again has cusps at its vertices. The tile is naturally equipped with horocycles at each of its vertices, namely the images of the horocycles on vertices of t . If two triangles t and u of T° share an edge, then we can

identify a tile of t and a tile of u along that edge in such a way that their decorated hyperbolic structures agree on the edge and define a decorated hyperbolic structure on their union.

Given a tagged arc α , the *tile cover* of α is a tagged arc α' in a surface $(\mathbf{S}', \mathbf{M}')$ with a decorated hyperbolic structure and a triangulation T' , constructed as follows and illustrated in Example 11 and Figures 32–38. (Strictly speaking, we should say “a” tile cover, since there are choices to be made at the end of the construction. We say “the” tile cover because the relevant parts of the tile cover are completely determined by α and T .) The crucial property of the tile cover is that (T', α') is tidy in the sense of Section 2. We only construct the tile cover when α has the following property: If α has an endpoint p that is the interior vertex of a self-folded triangle in T° , then α is tagged plain there. (We will only need this case because of Proposition 3.2 in the coefficient-free case and Proposition 1.7 in the principal-coefficients case.)

First, assume α does not coincide, except for tagging, with an arc of T° . Then α can be decomposed into a finite collection of segments, each of which has its endpoints on arcs of T° (possibly at marked points) and otherwise not intersecting the arcs of T° . Thus each segment is contained in some triangle of T° . For each of these segments construct a tile from the triangle containing it. Since α can intersect the same triangle multiple times, we might construct multiple tiles for a given triangle. Furthermore, at each endpoint of α that is not the interior point of a self-folded triangle in T° , construct a tile for each triangle-corner of a triangle at that endpoint. A triangle of T° can have 0, 1, 2, or 3 corners at that endpoint, so again some triangles may become multiple tiles. At each endpoint of α that is the interior point of a self-folded triangle in T° , construct two copies of the tile for the self-folded triangle. Then identify all of these triangles in the natural way: For each endpoint of α , identify the tiles on their edges around the endpoint in the order that the triangle corners appear around that endpoint. (When an endpoint p of α is the interior point of a self-folded triangle in T° , arrange the two tiles for that triangle about p to form a digon whose two edges correspond to the exterior edge of the self-folded triangle. If an endpoint of α is the puncture in a once-punctured digon formed by two triangles in T° , identify the tiles for the two triangles so that the same is true of the corresponding endpoint of α in T' .) Then continue identifying edges to join the segments of α in order and lift α to a tagged arc α' (the union of the lifts of the segments) in the union of the tiles. The union of the tiles is topologically a disk.

If α coincides, except for tagging, with an arc of T° , then instead of making tiles for segments of α , make two tiles, for the triangles on both sides of α . If α coincides with the interior edge of a self-folded triangle, then the two tiles are identical. Again at each endpoint where α , construct a tile for each triangle corner (without making additional copies of the tiles for the triangles on both sides of α). The edges of these tiles are then identified in the natural way around each endpoint of α . If α coincides with the interior edge of a self-folded triangle in T° , then the tiles from the two sides of α are glued together to form a digon. If α coincides with an interior edge of a punctured digon in T° , then α' also coincides with an interior edge of a punctured digon in T' . The union of the tiles is again a disk.

In either case, embed the disk into some larger marked surface $(\mathbf{S}', \mathbf{M}')$ with a tagged triangulation T' containing the tiles as triangles, with edges of tiles constituting boundary segments of \mathbf{S}' if and only if they come from boundary segments of \mathbf{S} . This can be done in such a way that T consists of plain-tagged arcs and each triangle of T has three distinct vertices and three distinct edges (as is already true of the tiles). Crucially, (T', α') is tidy.

Example 11. We illustrate the combinatorics of the tile cover in Figures 32–38. In each figure, the left picture shows α as a tagged arc in (\mathbf{S}, \mathbf{M}) . The marked surface (\mathbf{S}, \mathbf{M}) and the tagged triangulation T are the same in Figures 32–33 and in Figures 34–37. The ordinary triangulation T° is pictured, and the labeling of arcs of T° implies a labeling of the tagged arcs of T as usual. (In Figure 32 for example, 2 labels the tagged arc that agrees with 3 but with a notch at highest puncture shown, and similarly 10 labels the tagged arc that agrees with 11 but with a notch.) In each figure, the right picture shows the tile cover α' in $(\mathbf{S}', \mathbf{M}')$. The picture shows only the part of $(\mathbf{S}', \mathbf{M}')$ that is constructed explicitly. Gray shading indicates certain triangle edges that are boundary segments in $(\mathbf{S}', \mathbf{M}')$. All other edges shown are arcs, not boundary segments. Figure 34




FIGURE 32. The tile cover of a tagged arc.




FIGURE 33. The tile cover of a tagged arc.

continues the example of Figure 5. Figure 35 continues the example of Figure 6 and illustrates the reduction in Section 1.3: In Figure 35, the arc α is no longer tagged notched at the right puncture, and correspondingly, the roles of arcs 6 and 7 have been exchanged. Figure 36 continues the example of Figure 8.

Remark 3.3. The definition of the tile cover uses more tiles than are required for the results of this section: We could have left out the triangles surrounding endpoints of α that are tagged plain. The extra triangles around punctures will be needed in Section 4 when we prove the full version of Theorem 1.1.

The tiles that form the tile cover of α come equipped with a decorated hyperbolic structure: By construction, each tile gets this structure from a triangle of T° , and the structure agrees on shared edges of tiles. The chosen horocycles at marked points in \mathbf{M} lift to tiles along with the hyperbolic metric, and horocycles in adjacent tiles join to make horocycles in $(\mathbf{S}', \mathbf{M}')$. In order to apply the results of Section 2 to the tile cover of α , we need to correctly determine the lambda lengths of the arcs of T' in this decorated hyperbolic structure to $(\mathbf{S}', \mathbf{M}')$. (In fact, it is only necessary to determine the lambda lengths of the edges of tiles, because these determine the hyperbolic metric on the union of the tiles.) This is essentially straightforward, but with one small wrinkle coming from self-folded triangles.




FIGURE 34. The tile cover of a tagged arc.




FIGURE 35. The tile cover of a tagged arc.

As before, let x_β stand for the lambda length of a tagged arc $\beta \in T$, so that the initial cluster is $\{x_\beta \mid \beta \in T\}$. We emphasize that β is a *tagged* arc in T , not an ordinary arc in T° . For most arcs β , the distinction is meaningless, but if β is tagged notched at one endpoint (and thus by our assumptions on T coincides with another arc γ of T that is tagged plain), then β is *not* the ordinary arc (loop) that forms the outer edge of a self-folded triangle of T° . However, it is that loop that lifts to become an arc in the triangulation T' of $(\mathbf{S}', \mathbf{M}')$. In this situation, Equation (6) says that the loop has lambda length $x_\beta x_\gamma$. Thus the lambda lengths of tile edges are as follows: If the edge comes from a boundary segment of (\mathbf{S}, \mathbf{M}) , then its lambda length is 1. If the edge comes from an arc β that is *not* the outer edge of a self-folded triangle of T° , then its lambda length is x_β . If the edge comes from an arc δ that *is* the outer edge of a self-folded triangle of T° ,




FIGURE 36. The tile cover of a tagged arc.




FIGURE 37. The tile cover of a tagged arc.

then its lambda length is $x_\beta x_\gamma$, where β and γ are the two tagged arcs of T corresponding to that self-folded triangle.

Example 12. Continuing Example 11, we describe the lambda lengths of the arcs and boundary segments shown in tile cover (the right picture) in Figure 38. We use the numbering of arcs of T° shown in the left picture. The boundary segments (indicated by gray shading) each have lambda length 1. Arcs labeled 2 have lambda length $x_1 x_2$. Any other arc labeled i has lambda length x_i .




FIGURE 38. The tile cover of a tagged arc.

3.3. Proof of the coefficient-free version of the main theorem. The last ingredients in the proof of Theorem 3.1 are three lemmas that relate a tagged arc α to its tile cover α' . All three lemmas have the hypothesis that wherever α has an endpoint that is the interior vertex of a self-folded triangle in T° , the tagging of α is plain at that endpoint. In each lemma, α' is the tile cover of α in $(\mathbf{S}', \mathbf{M}')$.

Lemma 3.4. *The lambda length $\lambda(\alpha)$ in (\mathbf{S}, \mathbf{M}) equals the lambda length $\lambda(\alpha')$ in $(\mathbf{S}', \mathbf{M}')$.*


Proof. We have decomposed α into a finite collection of segments according to the triangulation T° , constructed a tile in T' for each segment with the same local hyperbolic metric, and transported the horocycles in the corresponding corners of the tiles. Crucially, whenever α is tagged notched at one of its endpoints p , the length of the horocycle h_p is the same as the length of the horocycle at the corresponding point in the tile cover. (The neighborhood of every puncture p maps isometrically to the neighborhood of the corresponding puncture in the tile cover, *except* when p is the interior vertex of a self-folded triangle of T° . However, by hypothesis, α must be tagged plain at any such puncture.) Therefore, in any case where the lambda length of α involves a conjugate horocycle, the lift of that conjugate horocycle to the tile cover is precisely the conjugate horocycle in $(\mathbf{S}', \mathbf{M}')$. We see that the length $l(\alpha)$ of the segment of α in (\mathbf{S}, \mathbf{M}) between the horocycles corresponding to its endpoints equals the length $l(\alpha')$ of the segment of α' in $(\mathbf{S}', \mathbf{M}')$ between the horocycles corresponding to its endpoints. Thus $\lambda(\alpha') = \exp(l(\alpha')/2) = \exp(l(\alpha)/2) = \lambda(\alpha)$. \square

Lemma 3.5. *The g-vector \mathbf{g}_α in (\mathbf{S}, \mathbf{M}) equals the g-vector $\mathbf{g}_{\alpha'}$ in $(\mathbf{S}', \mathbf{M}')$.*

Proof. The monomial $\mathbf{g}_{\alpha'}$ is the product of the lambda lengths $\lambda(\gamma')$ raised to the exponent $-b_{\gamma'}(T', \kappa(\alpha'))$ over all arcs $\gamma \in T'$. (Recall that $b_{\gamma'}(T', \kappa(\alpha'))$ is the entry, in position γ , of the shear coordinate of $\kappa(\alpha')$ with respect to T' .) Since T' has all arcs tagged plain, for any $\gamma \in T'$, there is a quadrilateral in T' with γ' as a diagonal. There is a contribution $\lambda(\gamma')^{\pm 1}$ to $\mathbf{g}_{\alpha'}$ every time $\kappa(\alpha')$ crosses that quadrilateral through opposite sides, but by construction, $\kappa(\alpha')$ makes at most one crossing of every quadrilateral in T' . For every quadrilateral crossed, there is a corresponding pair of triangles in T (in (\mathbf{S}, \mathbf{M})), crossed by α , and sharing an edge that is an arc $\gamma \in T^\circ$. Possibly α crosses γ many times, but this crossing makes the same contribution ± 1 as the contribution from α' crossing γ' .

If γ is not the exterior edge of a self-folded triangle in T° , then $\lambda(\gamma)$ is x_γ and, by construction $\lambda(\gamma')$ is also x_γ . If γ is the exterior edge of a self-folded triangle in T° , then let δ be the interior edge. There are two possibilities: Either this crossing of γ by $\kappa(\alpha)$ happens as $\kappa(\alpha)$ passes through the self-folded triangle, or this crossing happens as $\kappa(\alpha)$ spirals into the interior vertex of the triangle.

The case where $\kappa(\alpha)$ passes through the triangle is illustrated in Figure 39 (Example 13). In this case, by the definition of the shear coordinate in the presence of self-folded triangles, this crossing of γ contributes the same ± 1 to the shear coordinates at the entries γ and δ . Thus the contribution of this crossing of γ is $x_\gamma x_\delta$. By construction $\lambda(\gamma')$ is also $x_\gamma x_\delta$. (In this case


 FIGURE 39. Computing \mathbf{g}_α and $\mathbf{g}_{\alpha'}$.

where $\kappa(\alpha)$ passes through the self-folded triangle, it is impossible for γ to be the interior edge of a self-folded triangle in T° , because if so, any crossing of γ by α lifts to a crossing of γ' by α that contributes nothing to the shear coordinate.)


The case where $\kappa(\alpha)$ spirals into the interior edge appears in Figure 39 but also more specifically in Figure 40. This crossing either contributes $+1$ to the γ -entry of the shear coordinates of $\kappa(\alpha)$ and contributes nothing to the δ -entry or contributes nothing to the γ -entry and contributes -1 to the δ -entry, thus contributing either x_γ or x_δ^{-1} to \mathbf{g}_α . In either case, since α' is tagged plain at this endpoint, there is a contribution -1 to the shear coordinate of one of the arcs in T' that is a lift of the interior edge. (This occurs because $\kappa(\alpha')$ crosses that arc of T' and then hits an arc of T' that does not come from a tile.) When the crossing contributes $+1$ to the γ -entry of the shear coordinates of $\kappa(\alpha)$, the crossing by $\kappa(\alpha')$ contributes $+1$ to the shear coordinate given by the lift of γ , and the net contribution to $\mathbf{g}_{\alpha'}$ is $(x_\gamma x_\delta) \cdot x_\delta^{-1} = x_\gamma$. When the crossing contributes -1 to the δ -entry of the shear coordinates of $\kappa(\alpha)$, the crossing by $\kappa(\alpha')$ contributes 0 to the shear coordinate given by the lift of γ , and the net contribution to $\mathbf{g}_{\alpha'}$ is x_δ^{-1} . \square

Example 13. Figure 39 shows the computations of \mathbf{g}_α and $\mathbf{g}_{\alpha'}$, continuing the example of Figure 33. Both endpoints of α are at the same puncture, the interior vertex of a self-folded triangle. Following $\kappa(\alpha)$ to the right from this point and keeping in mind the minus sign in the exponent in Equation (2), the contributions to \mathbf{g}_α are $x_{11} x_{14}^{-1} x_7 x_2^{-1} x_3^{-1} x_1 x_9^{-1} x_{11}$. If we follow α' in the same direction (from the top of the picture on the right), the contributions to $\mathbf{g}_{\alpha'}$ are $x_{11} x_{14}^{-1} x_7 (x_2 x_3)^{-1} x_1 x_9^{-1} x_{11}$. Figure 40 shows two much simpler examples that illustrate the two possibilities for shear coordinates of $\kappa(\alpha)$ and $\kappa(\alpha')$ when α has an endpoint at the interior vertex of a self-folded triangle.

Lemma 3.6. $F(P_\alpha)|_{y=1} = F(P_{\alpha'})|_{y=1}$, where the notation of specializing to “ $y = 1$ ” is shorthand for setting every tropical variable to 1.

Proof. The posets are isomorphic by construction, and we check that corresponding ideals have the same weight in the sense of \tilde{w} (setting the tropical variables to 1 in w). Indeed, it is simple to check that corresponding elements have the same weight.

If γ is an arc in T° , then $\hat{y}_\gamma|_{y=1}$ is the product $\prod_{\beta \in T} x_\beta^{b_{\beta\gamma}}$. Here $b_{\beta\gamma}$ is the $\beta\gamma$ -entry of the exchange matrix (*i.e.* signed adjacency matrix) associated to T . If γ is not the interior edge of a self-folded triangle, then there is a positive contribution to the exponent of x_β for every triangle of T° in which β precedes γ in clockwise order, and, if β is the interior edge of a self-folded

FIGURE 40. Computing \mathbf{g}_α and $\mathbf{g}_{\alpha'}$ (self-folded triangle possibilities).

triangle with exterior edge δ , a positive contribution to the exponent of x_β for every triangle of T° in which β precedes δ in clockwise order. Negative contributions to the exponent of x_β are the same, but with “clockwise” replaced with “counterclockwise” throughout. (If γ is the exterior edge of a self-folded triangle and β is the interior edge, then the positive and negative contributions to the exponents of x_β cancel.) If γ is the interior edge of a self-folded triangle, then by the definition of the signed adjacency matrix, the product $\prod_{\beta \in T} x_\beta^{b_{\beta\gamma}}$ is the same as for the exterior edge.

For a lift of γ to the edge γ' of a tile in the tile cover $(\mathbf{S}', \mathbf{M}')$, the corresponding product is $\prod_{\beta' \in T'} \lambda(\beta')^{b_{\beta'\gamma'}}$. The contributions to this product have the same description, except that there are no self-folded triangles in T' . Instead, when β' is the lift of the exterior edge β of a folded triangle whose interior edge is δ , the lambda length $\lambda(\beta')$ is $x_\beta x_\delta$. When γ the exterior edge of a self-folded triangle with interior edge δ , the tile for that triangle has edges with lambda lengths $x_\gamma x_\delta$, x_δ , and x_δ . Thus the contributions of x_δ to $\prod_{\beta' \in T'} \lambda(\beta')^{b_{\beta'\gamma'}}$ cancel, just as they cancel in $\prod_{\beta \in T} x_\beta^{b_{\beta\gamma}}$. When γ is the interior edge of a self-folded triangle with exterior edge δ , the lift γ' is the edge of two tiles each with edges of lambda lengths x_γ , x_γ and $x_\gamma x_\delta$, and the edges with lambda lengths $x_\gamma x_\delta$ are adjacent in the quadrilateral formed by the two tiles. Thus all of the contributions to $\prod_{\beta' \in T'} \lambda(\beta')^{b_{\beta'\gamma'}}$ cancel, so that the product is 1.

Suppose e is an element of P_α and e' is the corresponding element of $P_{\alpha'}$. The above considerations show that if $w(e) = \hat{y}_\gamma$ (so that γ is not the interior edge of a self-folded triangle), then $\tilde{w}(e)$ and $\tilde{w}(e')$ both equal $\prod_{\beta \in T} x_\beta^{b_{\beta\gamma}}$. They also show that if $w(e) = \hat{y}_\gamma / \hat{y}_\delta$ (so that γ is the interior edge of a self-folded triangle with exterior edge δ), then $\tilde{w}(e)$ and $\tilde{w}(e')$ both equal 1. \square

Finally, we prove the coefficient-free case of the main result.

Proof of Theorem 3.1. By Proposition 3.2, we need only prove the theorem for tagged arcs α that are plain at any endpoint that is the interior vertex of a self-folded triangle. We write (5) as $\lambda(\alpha) = \mathbf{g}_\alpha \cdot F(P_\alpha)|_{y=1}$, with notation as in Lemma 3.6.

Suppose α is a tagged arc in (\mathbf{S}, \mathbf{M}) . We argue by induction on the number of elements of P_α that $\lambda(\alpha) = \mathbf{g}_\alpha \cdot F(P_\alpha)|_{y=1}$. Specifically, the inductive hypothesis is that the formula holds for any tagged arc α' in *any* triangulated marked surface such that $P_{\alpha'}$ has fewer elements than P_α . The base of the induction is when P_α is empty. In this case, the sum is empty, and α is either an arc in T or is the exterior edge of a self-folded triangle in T° . If α is an arc in T , then $\mathbf{g}_\alpha = x_\alpha$, as desired. If α is the exterior edge of a self-folded triangle in T° , then $\mathbf{g}_\alpha = x_\beta x_{\beta'}$, where β and β' are the two tagged arcs in T associated to the self-folded triangle. This is correct in light of (6).

If P_α is not empty, then construct the tile cover α' in a marked surface $(\mathbf{S}', \mathbf{M}')$ with a triangulation T' . By construction, α' has distinct endpoints and does not cross itself, so it is a tagged arc in $(\mathbf{S}', \mathbf{M}')$. Also, $P_{\alpha'}$ is isomorphic (ignoring the weighting for a moment) to P_α . Furthermore, (T', α') is tidy, so there is an arc $\gamma' \in T'$ such that $\hat{y}_{\gamma'}$ occurs exactly once as a label in $P_{\alpha'}$, and Proposition 2.2 applies. Applying (4) to α' and multiplying through by $x_{\gamma'} \cdot \mathbf{g}_{\alpha'}$, we obtain

$$(7) \quad x_{\gamma'} \cdot \mathbf{g}_{\alpha'} \cdot F(P_{\alpha'}) = x_{\gamma'} \cdot \mathbf{g}_{\alpha'} \cdot F(P_{\text{blue}}) + x_{\gamma'} \cdot \mathbf{g}_{\alpha'} \cdot \hat{y}_{\text{orange}} \cdot F(P_{\text{red}}).$$

The labeled posets P_{blue} and P_{red} , their labels in the \hat{y} -monomials, and the monomial \hat{y}_{orange} all refer, of course, to the triangulation T' of $(\mathbf{S}', \mathbf{M}')$.

The quantity $F(P_{\text{blue}})$ in Equation (7) is the product of one or two polynomials $F(P_{\delta'})$ for tagged arcs δ' in $(\mathbf{S}', \mathbf{M}')$, and each of these posets $P_{\delta'}$ has strictly fewer elements than $P_{\alpha'}$, which has the same number of elements as P_α . By induction, the lambda length $\lambda(\delta')$ in $(\mathbf{S}', \mathbf{M}')$ is $\mathbf{g}_{\delta'} \cdot F(P_{\delta'})|_{y=1}$. Each δ' is a tagged arc, so this lambda length is a cluster variable $x_{\delta'}$. Similarly, $F(P_{\text{red}})$ is the product of one or two polynomials $F(P_{\delta'})$ and by induction, the lambda length of each δ' is $\mathbf{g}_{\delta'} \cdot F(P_{\delta'})|_{y=1}$, which is again a cluster variable $x_{\delta'}$.

Specializing all tropical variables to 1 and appealing to Proposition 2.3, the right side of (7) becomes the right side of the exchange relation exchanging $x_{\alpha'}$ and $x_{\gamma'}$. We conclude that the cluster variable $x_{\alpha'}$ is given by $\mathbf{g}_{\alpha'} \cdot F(P_{\alpha'})|_{y=1}$. This cluster variable is the lambda length $\lambda(\alpha')$.

We have seen that $\lambda(\alpha') = \mathbf{g}_{\alpha'} \cdot F(P_{\alpha'})|_{y=1}$, which shows that $\lambda(\alpha) = \mathbf{g}_\alpha \cdot F(P_\alpha)|_{y=1}$ by Lemmas 3.4, 3.5 and 3.6. \square

4. TROPICAL HYPERBOLIC GEOMETRY AND THE PRINCIPAL COEFFICIENTS CASE

We now prove the full version of Theorem 1.1, which can be rephrased as follows: the cluster variable x_α equals the \mathbf{g} -vector \mathbf{g}_α times the weighted sum $F(P_\alpha)$. The proof follows the same outline and uses some of the same results as in Section 3. We begin by explaining how cluster variables with principal coefficients can be realized in terms of laminations and certain hyperbolic lengths and corresponding tropical hyperbolic lengths. More details are found in [FT18a, Chapters 9–15].

4.1. Opened surfaces. As before, we are given a marked surface (\mathbf{S}, \mathbf{M}) and a tagged triangulation T with all arcs tagged plain, except possibly at some punctures p incident to exactly two tagged arcs of T , identical except for opposite taggings at p . The corresponding ordinary triangulation is T° . The tagged triangulation determines a multi-lamination \mathbf{L} of (\mathbf{S}, \mathbf{M}) . General choices of \mathbf{L} lead to arbitrary coefficients of geometric type, but for principal coefficients at T , the multi-lamination \mathbf{L} consists of elementary laminations L_γ for tagged arcs $\gamma \in T$, as defined in Section 2.1. We will also consider a multilamination \mathbf{A} consisting of elementary laminations for arcs in T° .

Recall the construction reviewed in Section 3.1: Once (\mathbf{S}, \mathbf{M}) and T are fixed, there is a bijection from the set of tuples of positive numbers indexed by T to the set of decorated hyperbolic structures on (\mathbf{S}, \mathbf{M}) —meaning hyperbolic structures with constant curvature -1 and distinguished horocycles about each marked point—with all boundary segments having lambda length 1. The tuples of positive numbers describe the lambda lengths x_α of the arcs $\alpha \in T$. Once a particular decorated hyperbolic structure is chosen (or equivalently, once the lambda lengths x_α are chosen), every tagged arc has a lambda length. The lambda length of a tagged arc γ , as a function of the lambda lengths x_α , is the coefficient-free cluster variable associated to γ .

The analogous construction with principal coefficients is summarized as follows: Once (\mathbf{S}, \mathbf{M}) and T are fixed (and thus \mathbf{L} is fixed to be the elementary laminations associated to T), we choose *two* tuples of positive numbers, each tuple indexed by T . For each tagged arc $\alpha \in T$, we choose a positive real number x_α that will eventually be its laminated lambda length. For each elementary lamination L_γ (for $\gamma \in T$), we choose a positive real weight y_γ that will play the role of a tropical variable. (In [FT18a], these tropical variables are denoted q_i .) These pairs of tuples determine an opening of \mathbf{S} (a surface in which some of the punctures are replaced by holes, meaning removed disks with an orientation), a hyperbolic metric with curvature -1 on the opened surface and

horocycles on all of the non-opened marked points, with boundary conditions on the metric that determine the lambda lengths of the boundary segments and the holes in terms of the y_γ . Once an opening, hyperbolic metric, and horocycles are chosen (or equivalently, once the laminated lambda lengths x_α and the weights y_γ are chosen), every tagged arc has a laminated lambda length. The laminated lambda length of an arc γ , as a function of the x_α and the y_γ , is the principal-coefficients cluster variable associated to γ . We now provide more details.

The **transverse measure** $l_{L_\gamma}(p)$ of a puncture p with respect to an elementary lamination L_γ is 1 for each endpoint of γ that is at p and tagged plain and -1 for each endpoint of γ that is at p and tagged notched, and otherwise zero. Since it is possible for γ to have both endpoints at p if they have the same tagging, these contributions can combine to make 2 or -2 but can never cancel. (In this definition of transverse measure of a puncture, we have combined [FT18a, (14.2)] with the definition of L_γ and the definition of a lift of a lamination that will be given below.) The **transverse measure** $l_{L_\gamma}(\beta)$ of a boundary segment β with respect to L_γ is the number of times (0, 1, or 2) that L_γ has an endpoint on β . Thus if q is the right endpoint of β (right as we stand on β looking towards the interior of \mathbf{S}), then in light of the definition of L_γ , the transverse measure of β is the number of endpoints of γ that are at q . The **tropical lambda length** is defined to be $c(s) = \prod_{\gamma \in T} y_\gamma^{l_{L_\gamma}(s)}$ for s a puncture or boundary segment. (Later, when we define tropical lambda lengths of tagged arcs, we use the notation $c_{\bar{\mathbf{L}}}$. There, $\bar{\mathbf{L}}$ stands for a particular lift of the multilamination \mathbf{L} to the opened surface. However, the tropical lambda lengths of punctures and boundary segments are independent of the lift, and we suppress the subscript $\bar{\mathbf{L}}$ here.)

To make an **opening** of (\mathbf{S}, \mathbf{M}) , we choose some subset \mathbf{P} of the punctures and open those by replacing each $p \in \mathbf{P}$ with a hole C_p (a new boundary component where an open disk is removed from \mathbf{S}) with marked point M_p on C_p . Each C_p for $p \in \mathbf{P}$ is given an orientation, clockwise or counterclockwise. The opening depends on the choice of the weights y_γ . Specifically, we use these weights to assign lambda lengths to holes and boundary segments in the opened surface, thus constraining the choice of which holes are opened, the orientations of opened holes, the hyperbolic metric, and the choice of horocycles at marked points on the boundary, as we now explain.

Assuming some opening of (\mathbf{S}, \mathbf{M}) , choose a finite-area hyperbolic metric σ with constant curvature -1 with a cusp at each *non-opened* marked point, with each boundary segment a geodesic, including the holes C_p . (We emphasize that there are no cusps at the points M_p for opened punctures p .) The **signed length** $l(p)$ of an opened puncture p is the hyperbolic length of C_p if C_p is oriented clockwise in the opened surface or the negative of the length of C_p if it is oriented counterclockwise. The signed length of an unopened puncture is 0. Assuming also some choice of horocycles at marked points on the boundary, we define the **signed length** $l(\bar{\alpha})$ of a boundary component between the horocycles at its endpoints as in Section 3.1. The **lambda length** of a puncture or boundary segment s is defined to be $\lambda(s) = e^{l(s)/2}$.

We constrain the choice of the hyperbolic metric σ by specifying lambda lengths for the punctures and boundary segments: Every puncture or boundary segment s had lambda length $\lambda(s)$ equal to the tropical lambda length $c(s)$. The tropical lambda lengths of punctures and boundary segments depend only on the weights y_γ . (For more general coefficients, they depend on the choice of a weighted multilamination, but we have fixed principal coefficients and thus fixed the multilamination.) Thus the assigned lambda lengths of punctures and boundary segments also depend only on the weights y_γ . In particular, the weights y_γ determine the opening of surface: If $c(p) = 1$ for some puncture p , then the puncture p is not opened. If $c(p) > 1$, then p is opened and C_p is oriented clockwise, and if $c(p) < 1$ then p is opened and C_p is oriented counterclockwise.

4.2. Laminated lambda lengths. We have seen how the weights y_λ determine an opening of (\mathbf{S}, \mathbf{M}) and place constraints on the hyperbolic metric and horocycles by specifying lambda lengths of the opened punctures and the boundary segments. We will eventually impose additional constraints coming from the x_γ . But first, we define laminated lambda lengths of tagged arcs.

At each M_p , for p an opened puncture, there is a **perpendicular horocyclic segment** h_p near C_p . This is a short segment of the horocycle that contains M_p and is perpendicular to every geodesic spiraling into C_p in the direction of the chosen orientation. We also define a

conjugate perpendicular horocycle \bar{h}_p near C_p : Let \bar{M}_p be the point on C_p a *signed* distance $2 \ln |\lambda(p) - \lambda(p)^{-1}|$ from M_p in the direction *against* the orientation of C_p . Let \bar{h}_p be the segment at \bar{M}_p that is perpendicular to every geodesic spiraling into C_p opposite of the chosen orientation.

Now let α be any tagged arc in (\mathbf{S}, \mathbf{M}) . We will need two different types of lifts of α to the opened surface. First, an arbitrary lift $\bar{\alpha}$ in a topological sense, meaning that $\bar{\alpha}$ maps to α under that map that collapses all opened punctures back down to punctures in (\mathbf{S}, \mathbf{M}) . Second, a corresponding geodesic α_σ in the opened surface: The geodesic α_σ is like a lift of α except that at any endpoint p of alpha that is opened, the geodesic spirals around C_p , with the direction of the chosen orientation of C_p if α is tagged plain at p or against the orientation of C_p if α is tagged notched there.

We will define the laminated lambda length of α as the ratio of the ordinary lambda length and tropical lambda length of the lift $\bar{\alpha}$. In particular, this ratio will not depend on the choice of lift.

We first define the lambda length of $\bar{\alpha}$, beginning with the case where, at any endpoints of α that are opened punctures, $\bar{\alpha}$ twists far enough about C_p in the direction that α_σ spirals. We choose a segment $[\alpha_\sigma]$ of α_σ that starts at the intersection of α_σ with a horocycle or perpendicular horocyclic segment at one endpoint of α (a conjugate horocycle or conjugate perpendicular horocyclic segment if α is notched at that endpoint), follows α_σ and ends at the intersection of α_σ with a horocycle or perpendicular horocyclic segment at the other endpoint (again conjugate if α is notched there). When α_σ spirals about an opened puncture at one or both endpoints, there are infinitely many choices of such a segment, but we choose it so that when we extend $[\alpha_\sigma]$ in a particular way, we obtain a curve that is isotopic to $\bar{\alpha}$: At an endpoint p of α that is not an opened puncture, we extend $[\alpha_\sigma]$ along α_σ all the way to p . At an endpoint p of α that is an opened puncture where α is tagged plain, we extend $[\alpha_\sigma]$ by following the perpendicular horocyclic segment h_p to M_p . At an endpoint p of α that is an opened puncture where α is tagged notched, we extend $[\alpha_\sigma]$ by following the conjugate perpendicular horocyclic segment \bar{h}_p to \bar{M}_p and then following C_p to M_p , in the direction that *agrees* with the orientation of C_p . The **signed length** $l([\alpha_\sigma])$ is the hyperbolic length of $[\alpha_\sigma]$, signed as in Section 3.1. (Thus $l([\alpha_\sigma])$ is negative when the endpoint of $[\alpha_\sigma]$ where α_σ crosses a horocycle h_p is closer to the other endpoint of α than it is to p . When p is an opened puncture, one can choose the lift $\bar{\alpha}$ so that $[\alpha_\sigma]$ spirals far enough about C_p to make $l([\alpha_\sigma])$ positive.) The **lambda length** $\lambda(\bar{\alpha})$ is $e^{l([\alpha_\sigma])/2}$.

If $\bar{\alpha}$ does not spiral far enough about some opened puncture, its lambda length is defined by the following rule: When we pass from $\bar{\alpha}$ to a different lift with one additional spiral in the clockwise direction, we multiply the lambda length by $\lambda(p)$. By adding clockwise and/or counterclockwise twists at one or both endpoints of $\bar{\alpha}$, we obtain a different lift that “twists far enough” as in the paragraph above, and then apply the rule to find the lambda length of $\bar{\alpha}$. (The well-definition of the lambda length, in light of this rule, is established as [FT18a, Lemma 10.7]. See [FT18a, Definition 10.6].)

We next define the tropical lambda length $c_{\mathbf{T}}(\bar{\alpha})$ of $\bar{\alpha}$. Following [FT18a, Section 14], we define the tropical lambda length for a lift of α to a “fully” opened surface, meaning a surface where *every* puncture of (\mathbf{S}, \mathbf{M}) is opened. We reuse the notation $\bar{\alpha}$ for this lift. As explained below in Remark 4.1, this is harmless because the tropical lambda length is independent of the details of how α is lifted to the fully opened surface near punctures that are not opened in the partially opened surface.

We arbitrarily choose, for each elementary lamination L_γ , a **lifted lamination** \bar{L}_γ in the (fully) opened surface that corresponds to L_γ except that anywhere L_γ spirals into a puncture p that is opened, the lift \bar{L}_γ instead ends on $C_p \setminus \{M_p\}$ and formally retains a “memory” of the spiral in the following sense: \bar{L}_γ formally includes the data of an orientation of C_p in the direction *opposite* to the direction that L_γ spirals. (Recall that the opening of the surface also includes an orientation of C_p for every puncture p that is opened. The orientations in the opening and orientations in the lifted laminations are independent, and we will be careful to distinguish them.) We write $\bar{\mathbf{L}}$ for the multilamination consisting of the chosen lift \bar{L}_γ of each L_γ . The laminated lambda lengths of

tagged arcs depend on this choice of liftings, but we fix the choice and we need not be concerned with the dependence. For details, see [FT18a, Remark 14.8].

To define the **tropical lambda length** $c_{\overline{\mathbf{L}}}(\overline{\alpha})$ of $\overline{\alpha}$, we first suppose that, at each endpoint p of α , the lift $\overline{\alpha}$ twists sufficiently many times around C_p in the direction given by \overline{L}_γ or in the opposite direction if α is tagged notched at the shared endpoint. The tropical lambda length is defined by way of the transverse measure. We assume that isotopy representatives of $\overline{\alpha}$ and \overline{L}_γ have been chosen to minimize intersections. The **transverse measure** $l_{\overline{L}_\gamma}(\overline{\alpha})$ of $\overline{\alpha}$ relative to a lifted elementary lamination \overline{L}_γ is the number of intersections between $\overline{\alpha}$ and \overline{L}_γ , plus an additional contribution $|l_{L_\gamma}(p)|$ for each endpoint p where α is tagged notched. (The additional contribution $|l_{L_\gamma}(p)|$ is 1 or 2; it is 2 if and only if γ has both endpoints at p .) The **tropical lambda length** is defined to be $c_{\overline{\mathbf{L}}}(\overline{\alpha}) = \prod_{\gamma \in T} y_\gamma^{l_{\overline{L}_\gamma}(\overline{\alpha})}$.

If $\overline{\alpha}$ does not twist sufficiently far about some opened puncture p , then the tropical lambda length $c_{\overline{\mathbf{L}}}(\overline{\alpha})$ is determined using the following rule: When we pass from $\overline{\alpha}$ to a different lift with one additional spiral in the clockwise direction (or counterclockwise if α is notched at p), we multiply the tropical lambda length by $c(p)$.

Remark 4.1. It is harmless to define the tropical lambda length of $\overline{\alpha}$ in terms of a lift to the fully opened surface. To see why, recall that a puncture p is un-opened if and only if $c(p) = 1$. Since lifts of α that differ only in the number of twists at p are related by multiples of $c(p)$, the tropical lambda length of $\overline{\alpha}$ is independent of how we lift $\overline{\alpha}$ further to the fully opened surface.

The **laminated lambda length** of α is $x_{\overline{\mathbf{L}}}(\alpha) = \frac{\lambda(\overline{\alpha})}{c_{\overline{\mathbf{L}}}(\overline{\alpha})}$. Because of the rules for how lambda lengths and tropical lambda lengths change when we change the lift of α and because $\lambda(p) = c(p)$ for every puncture p , the laminated lambda length is well defined (i.e. independent of the choice of lift $\overline{\alpha}$).

Once the tagged triangulation T and the lifts $\overline{\mathbf{L}}$ of the elementary laminations L_γ for $\gamma \in T$ are fixed, the choice of positive real weights y_γ and the choice of positive real numbers x_γ uniquely determines an opening of (\mathbf{S}, \mathbf{M}) , hyperbolic metric on the opened surface, and horocycles about the non-opened marked points such that $x_{\overline{\mathbf{L}}}(\gamma) = x_\gamma$ for every $\gamma \in T$ [FT18a, Corollary 15.5]. Thus, given T , $\overline{\mathbf{L}}$, the y_γ and the x_γ , every tagged arc in (\mathbf{S}, \mathbf{M}) has a well-defined (principal-coefficients) laminated lambda length. The laminated lambda lengths of tagged arcs, when viewed as functions of the x_γ and the y_γ , are the cluster variables in the cluster algebra with initial cluster $\{x_\gamma : \gamma \in T\}$, initial exchange matrix equal to the signed adjacency matrix of T , principal coefficients at T , and tropical variables y_γ for $\gamma \in T$. (This is part of [FT18a, Theorem 15.6].)

Recall from Section 3.1 (Equation (6)) that the lambda lengths of the tagged arcs in T determine the lambda lengths of the ordinary arcs in T° and vice versa. Precisely the same thing is true for laminated lambda lengths. Suppose β and γ are tagged arcs in T that coincide except that both are tagged plain at one endpoint q but they have different taggings at p , and suppose that δ is the ordinary arc in T° that is a loop at q , following β and γ around p . Then

$$(8) \quad x_{\overline{\mathbf{L}}}(\delta) = x_{\overline{\mathbf{L}}}(\beta)x_{\overline{\mathbf{L}}}(\gamma).$$

The ordinary arc δ is not a tagged arc, but it has a laminated lambda length, measured from h_q to h_q . Equation (8) is a combination of [FT18a, Lemma 10.14], which says that $\lambda(\delta) = \lambda(\beta)\lambda(\overline{\gamma})$ for appropriate lifts $\overline{\delta}$, $\overline{\beta}$, and $\overline{\gamma}$, and [FT18a, Lemma 14.10], which says that $c_{\overline{\mathbf{L}}}(\overline{\delta}) = c_{\overline{\mathbf{L}}}(\overline{\beta})c_{\overline{\mathbf{L}}}(\overline{\gamma})$.

4.3. Proof of the main theorem.

In this section, we prove Theorem 1.1.

Let α be a tagged arc in (\mathbf{S}, \mathbf{M}) . For the proof of Theorem 3.1, we constructed the tile cover α' of α in a surface $(\mathbf{S}', \mathbf{M}')$ with triangulation T' and lifted the hyperbolic structure of (\mathbf{S}, \mathbf{M}) to the tile cover, so that lambda lengths were preserved. For the proof of Theorem 1.1, we lift all of the data that defines the laminated lambda lengths to the tile cover. As before, when passing to the tile cover, we assume that when α has an endpoint that is the interior vertex of a self-folded triangle in T° , the tagging of α is plain at that endpoint.

Start with the tagged triangulation T of (\mathbf{S}, \mathbf{M}) and choose positive real numbers x_γ and y_γ for each tagged arc $\gamma \in T$. Let these choices determine an opening of (\mathbf{S}, \mathbf{M}) , a hyperbolic metric

on the opened surface, and a choice of horocycles about the non-opened marked points, as in Section 4.2. In particular, the laminated lambda length of each tagged arc $\gamma \in T$ is x_γ .

We lift these laminated lambda lengths to determine laminated lambda lengths of all tile edges in (\mathbf{S}, \mathbf{M}) . Equation (8) determines the laminated lambda lengths of all arcs in T° that are not in T : For every self-folded triangle in T° corresponding to tagged arcs β and γ in T , with β tagged notched and γ tagged plain at the interior vertex, the exterior edge of the self-folded triangle has laminated lambda length $x_\beta x_\gamma$. Thus we lift laminated lambda lengths of tagged arcs in T to laminated lambda lengths of tile edges in T' exactly as we lifted lambda lengths of tagged arcs in Section 3.2. (See Example 12.)

The weighted multilamination on (\mathbf{S}, \mathbf{M}) that defines an opening (and places some constraints on the metric) consists of elementary laminations L_γ with weights y_γ , for $\gamma \in T$. Since the tagged arcs $\gamma \in T$ don't necessarily all lift to the tile edges in T' (but rather the arcs in T° do), this weighted multilamination does not lift to the right weighted lamination on $(\mathbf{S}', \mathbf{M}')$. We will define an alternative weighted multilamination \mathbf{A} that defines the same opening and places the same constraints on the metric. Specifically, \mathbf{A} consists of the elementary lamination L_γ for each $\gamma \in T^\circ$, with weight is y_γ unless γ is an edge of a self-folded triangle in T° . For the edges of a self-folded triangle corresponding to tagged arcs β and γ in T , with β tagged notched and γ tagged plain at the interior vertex, the weights are as follows: Let $\delta \in T^\circ$ be the exterior edge of the self-folded triangle and identify the internal edge with γ . Then the weight of L_γ in \mathbf{A} is $\frac{y_\gamma}{y_\beta}$ and the weight of L_δ is y_β .

Example 14. We describe the weights in the alternative multilamination \mathbf{A} in the left picture of Figure 32. Arcs labeled 3 have weights $\frac{y_3}{y_2}$ and arcs labeled 11 have weights $\frac{y_{11}}{y_{10}}$. Any other arc labeled i has weight y_i .

The following lemma is immediate from the definition of tropical lambda lengths of boundary segments and punctures.

Lemma 4.2. *The weighted multilaminations \mathbf{L} and \mathbf{A} assign the same tropical lambda lengths to punctures and boundary segments as the principal coefficients weighted lamination.*

Lemma 4.2 means that the weighted multilamination \mathbf{A} defines the same opening of (\mathbf{S}, \mathbf{M}) as \mathbf{L} . Furthermore, the lambda lengths of punctures and boundary components are the same for the two weighted laminations. We will see that, for the right choice of lifting of the two multilaminations to the opened surface, the tropical lambda lengths of all tagged arcs are also the same.

Choose a lift $\bar{\mathbf{L}}$ of the principal-coefficients weighted multilamination \mathbf{L} to the opened surface so that, whenever two tagged arcs $\beta, \gamma \in T$ coincide except that β is notched at p and γ is tagged plain at p , the corresponding lifted elementary laminations \bar{L}_β and \bar{L}_γ differ only in that $\bar{\beta}$ twists one additional time about C_p in the clockwise direction (and also necessarily differ in the orientation of C_p , with \bar{L}_γ having a clockwise orientation). Choose a lift $\bar{\mathbf{A}}$ of the alternative weighted lamination \mathbf{A} that agrees with the lift of \mathbf{L} where the arcs in T° coincide with plain-tagged arcs in T and for every arc δ in T° that is the exterior edge of a self-folded triangle with interior edge γ , the lift of L_δ traces closely along the lift of L_γ , goes around the interior vertex, and traces closely back along the lift of γ .

Lemma 4.3. *If α is a tagged arc and $\bar{\alpha}$ is any lift of α , then $c_{\bar{\mathbf{L}}}(\bar{\alpha}) = c_{\bar{\mathbf{A}}}(\bar{\alpha})$.*

Proof. Suppose δ is the exterior edge of a self-folded triangle with interior edge γ and β is the tagged arc that agrees with γ but is tagged notched at the interior vertex p . We need to check that the (multiplicative) contributions of \bar{L}_β and \bar{L}_γ to $c_{\bar{\mathbf{L}}}(\bar{\alpha})$ are the same as the contributions of \bar{L}_γ and \bar{L}_δ to $c_{\bar{\mathbf{A}}}(\bar{\alpha})$.

Away from any endpoint of α that is at p , the lifts \bar{L}_β and \bar{L}_γ agree, and there is contribution $y_\beta y_\gamma$ to $c_{\bar{\mathbf{L}}}(\bar{\alpha})$ each time $\bar{\alpha}$ intersects \bar{L}_β and \bar{L}_γ . The corresponding contribution to $c_{\bar{\mathbf{A}}}(\bar{\alpha})$ is $y_\beta^2 \frac{y_\gamma}{y_\beta}$, because the weight of δ is y_β , the weight of γ is $\frac{y_\gamma}{y_\beta}$, and (away from p), \bar{L}_δ traces along \bar{L}_γ twice.

To deal with contributions to $c_{\bar{\mathbf{L}}}(\bar{\alpha})$ and $c_{\bar{\mathbf{A}}}(\bar{\alpha})$ at any endpoints of α that are at p , we first look at how these contributions change when $\bar{\alpha}$ changes. Since \bar{L}_β carries a counterclockwise orientation




FIGURE 41. An illustration of the proof of Lemma 4.3.

of C_p and \bar{L}_γ carries a clockwise orientation, whenever $\bar{\alpha}$ is given an additional clockwise twist at C_p , the tropical lambda length $c_{\bar{\mathbf{L}}}(\bar{\alpha})$ is multiplied by a factor of $\frac{y_\gamma}{y_\beta}$ (or $\frac{y_\beta}{y_\gamma}$ if α is notched at p). Since the weight of γ in \mathbf{A} is $\frac{y_\gamma}{y_\beta}$, when $\bar{\alpha}$ is given an additional clockwise twist, $c_{\bar{\mathbf{A}}}(\bar{\alpha})$ is also multiplied by a factor of $\frac{y_\gamma}{y_\beta}$ (or $\frac{y_\beta}{y_\gamma}$ if α is notched). Thus, to prove that the contributions to $c_{\bar{\mathbf{L}}}(\bar{\alpha})$ and $c_{\bar{\mathbf{A}}}(\bar{\alpha})$ at p are the same, we may make a convenient choice of the lift $\bar{\alpha}$. By matching the spirals of $\bar{\alpha}$ to those of \bar{L}_γ , we will assume that $\bar{\alpha}$ does not intersect \bar{L}_γ at p . (See Figure 41, where labels β , γ , and δ are shorthand for \bar{L}_β , \bar{L}_γ , and \bar{L}_δ . The left picture shows $\bar{\mathbf{L}}$ and the right picture shows $\bar{\mathbf{A}}$.)

For a moment, assume that the endpoint of α is tagged plain at p . Then \bar{L}_γ contributes a factor 1 to $c_{\bar{\mathbf{L}}}(\bar{\alpha})$. Since \bar{L}_β twists clockwise about C_p once more than \bar{L}_γ , a lift of α with one more clockwise twist would have contribution 1 from \bar{L}_β as well. Since we obtain $\bar{\alpha}$ by adding one counterclockwise twist to a lift with contribution 1 from \bar{L}_β and since \bar{L}_β carries a counterclockwise orientation at p , the contribution of \bar{L}_β to $c_{\bar{\mathbf{L}}}(\bar{\alpha})$ is y_β . Because \bar{L}_δ follows \bar{L}_γ to p , goes around p , and then follows \bar{L}_γ again, $\bar{\alpha}$ intersects \bar{L}_δ exactly once at p . Since also $\bar{\alpha}$ does not intersect \bar{L}_γ at p , the contribution to $c_{\bar{\mathbf{A}}}(\bar{\alpha})$ at p is again y_β .

If the endpoint of α is tagged notched at p , there are additional contributions of y_β^{-1} and y_γ to $c_{\bar{\mathbf{L}}}(\bar{\alpha})$ at p and there is an additional contribution of $\frac{y_\gamma}{y_\beta}$ to $c_{\bar{\mathbf{A}}}(\bar{\alpha})$ at p . \square

Remark 4.4. It may seem surprising that different weighted multilaminations give the same laminated lambda lengths. To avoid this non-uniqueness of multilaminations, one would need to replace laminations with the quasi-laminations of [Rea14, Section 4]. However, for the present purposes, the non-uniqueness is a benefit, because it allows us to lift the multilamination to the tile cover in the right way. Indeed, the fact that we lift the alternative multilamination rather than the principal coefficients lamination accounts for the fractional labels on some elements of P_α .

In light of Lemmas 4.2 and 4.3, we lift the alternative weighted multilamination \mathbf{A} to the tile cover. This is done by making the lamination correct in each tile. The result is a principal coefficients weighted multilamination on the tile cover where the weights have been specialized to match the weights on the alternative weighted multilamination in (\mathbf{S}, \mathbf{M}) .

Example 15. Continuing Example 14, the weights on the right picture of Figure 32 are $\frac{y_3}{y_2}$ for arcs labeled 3, $\frac{y_{11}}{y_{10}}$ for arcs labeled 11, and otherwise y_i for arcs labeled y_i .

Now that we know the weights on the arcs in T' , we can prove the following strengthening of Lemma 3.6.

Lemma 4.5. *If α' is the tile cover of α , then $F(P_\alpha) = F(P_{\alpha'})$.*

Proof. Elements of the posets have “weights”, and there are also “weights” y_γ for $\gamma \in T$. Both kinds of weights appear in this proof, but for clarity, we will use the word “weight” in this proof only for the weights on the posets and refer to the y_γ as “ y -variables”.

The two posets are isomorphic by construction, and we check that corresponding elements have the same weights. We checked in the proof of Lemma 3.6 that corresponding elements of the

posets have the same weight when all the y -variables are specialized to 1. Since each weight is a Laurent monomial in the x - and y -variables, we complete the proof by showing that corresponding elements of the posets have the same weight when all the x -variables are specialized to 1. We make that specialization for the remainder of the proof.

In P_α , an element corresponding to crossing an arc $\gamma \in T^\circ$ is weighted y_γ unless γ is the interior edge of a self-folded triangle, in which case it is weighted $\frac{y_\gamma}{y_\beta}$, where β is the exterior edge. Every element of $P_{\alpha'}$ is weighted by a single y -variable in the sense of the weighted lamination on $(\mathbf{S}', \mathbf{M}')$. Since we lifted the alternative weighted lamination to $(\mathbf{S}', \mathbf{M}')$, this is again y_γ unless γ is an interior edge, in which case it is $\frac{y_\gamma}{y_\beta}$, where β is the exterior edge. \square

The opening of (\mathbf{S}, \mathbf{M}) can be made by cutting corners from triangles in T° . We obtain the opening of $(\mathbf{S}', \mathbf{M}')$ by making corresponding cuts in the tiles and glueing the tiles together so that the removed corners become holes. At any endpoints of α that are opened, the neighborhood of the hole is the same in (\mathbf{S}, \mathbf{M}) as in $(\mathbf{S}', \mathbf{M}')$, except when an endpoint of α is the interior vertex of a self-folded triangle. In that case, the neighborhood of the hole in $(\mathbf{S}', \mathbf{M}')$ is a double cover of the neighborhood of the hole in (\mathbf{S}, \mathbf{M}) . We lift the hyperbolic metric to the tiles (minus removed corners) and also lift the horocycles, conjugate horocycles, perpendicular horocyclic segments, and conjugate perpendicular horocyclic segments. In the case where the neighborhood of the hole in $(\mathbf{S}', \mathbf{M}')$ is a double cover, we obtain two marked points (and two perpendicular horocyclic segments). Of the two marked points, we retain the one that is the endpoint of α' and ignore the other (and its perpendicular horocyclic segment).

Moving forward, we assume all the structure on $(\mathbf{S}', \mathbf{M}')$ that was lifted from (\mathbf{S}, \mathbf{M}) as described above. The following is the analog of Lemma 3.4 for the principal coefficients case.

Lemma 4.6. *If α' is the tile cover of α , then the laminated lambda length $\lambda(\alpha)$ in (\mathbf{S}, \mathbf{M}) equals the laminated lambda length $\lambda(\alpha')$ in $(\mathbf{S}', \mathbf{M}')$.*

Proof. Take any lift $\bar{\alpha}'$ of the tile cover α' that lies within the tiles in the opening of $(\mathbf{S}', \mathbf{M}')$. There is a projection of the opening of $(\mathbf{S}', \mathbf{M}')$ to the opening of (\mathbf{S}, \mathbf{M}) by sending each tile isometrically to the corresponding triangle, and we write $\bar{\alpha}$ for the projection of $\bar{\alpha}'$. As the notation suggests, $\bar{\alpha}$ is a lift of α to the opening of (\mathbf{S}, \mathbf{M}) .

The simple idea of the proof is as follows: The lambda length of $\bar{\alpha}'$ equals the lambda length of $\bar{\alpha}$ because the hyperbolic structure on the opening of $(\mathbf{S}', \mathbf{M}')$ is lifted, tile by tile, from the hyperbolic structure on the opening of (\mathbf{S}, \mathbf{M}) —everywhere where $\bar{\alpha}$ passes through the surface—and furthermore, all horocycles, conjugate horocycles, perpendicular horocyclic segments, and conjugate perpendicular horocyclic segments relevant to $\bar{\alpha}'$ are lifted from the corresponding horocycles, etc. relevant to $\bar{\alpha}$. Similarly, the weighted multilamination on $(\mathbf{S}', \mathbf{M}')$ is lifted from the alternative weighted multilamination on (\mathbf{S}, \mathbf{M}) , so the tropical lambda length of $\bar{\alpha}'$ equals the tropical lambda length of $\bar{\alpha}$ with respect to the alternative weighted multilamination, and thus by Lemma 4.3 equals the principal coefficients tropical lambda length of $\bar{\alpha}$. Since the laminated lambda length is the ratio of the lambda length over the tropical lambda length, it follows that the laminated lambda length of α equals the laminated lambda length of α' .

This simple argument is complete and correct in most cases, but fails when one or both endpoints of α is at the interior vertex p of a self-folded triangle in T° . The issue is that the neighborhood of C_p , which is vital to computing both lambda lengths and tropical lambda lengths, does not lift to the neighborhood of an opened puncture in the opening of $(\mathbf{S}', \mathbf{M}')$. Instead, the neighborhood of C_p is *doubly covered* by the neighborhood of an opened puncture in the opening of $(\mathbf{S}', \mathbf{M}')$. The argument is easily rescued in such cases. As in the proof of Lemma 3.4, a crucial point is that in such cases, α is tagged plain at p . Another crucial point is how lambda lengths and tropical lambda lengths change when the lifts are changed by inserting additional twists about C_p .

Let δ be the exterior edge of the self-folded triangle whose interior vertex is p , let γ be the interior edge, and let β be the tagged arc that agrees with γ but is tagged notched at p . The puncture p' in $(\mathbf{S}', \mathbf{M}')$ corresponding to this endpoint of α is contained in two tiles. The two edges incident to p' both have laminated lambda lengths x_γ . The weights on the corresponding

elementary laminations are both $\frac{y_\gamma}{y_\beta}$. Thus the tropical lambda length of p' is $c(p') = \frac{y_\gamma^2}{y_\beta^2}$, and this also equals $\lambda(p')$.

In (\mathbf{S}, \mathbf{M}) , on the other hand, there is one edge of T° incident to p , with lambda length x_γ . The associated elementary lamination has weight $\frac{y_\gamma}{y_\beta}$, so $c(p) = \frac{y_\gamma}{y_\beta} = \lambda(p)$. The fact that $c(p') = c(p)^2$ is consistent with the fact that the neighborhood of $C_{p'}$ doubly covers the neighborhood of C_p .

We can now easily see that the tropical lambda lengths of $\bar{\alpha}$ and $\bar{\alpha}'$ are still equal, despite the local double cover. The orientation of C_p in \bar{L}_γ is clockwise, as is the orientation of $C_{p'}$ in the two copies of \bar{L}_γ that were lifted to $(\mathbf{S}', \mathbf{M}')$. For a moment, assume that the lift $\bar{\alpha}'$ so that it twists far enough clockwise about $C_{p'}$ and so that $\bar{\alpha}$ also twists far enough about C_p . (We can do this at both endpoints of α if necessary. Since the requirement is that $\bar{\alpha}$ must twist far enough clockwise at both endpoints, this can be done even if the endpoints of α are both at p .) Now every contribution of \bar{L}_γ to the tropical lambda length of $\bar{\alpha}$ in (\mathbf{S}, \mathbf{M}) lifts to a corresponding contribution (with the same weight) to the tropical lambda length of $\bar{\alpha}'$ in $(\mathbf{S}', \mathbf{M}')$.

We see that, for a particular choice of $\bar{\alpha}'$, the tropical lambda lengths of $\bar{\alpha}$ and $\bar{\alpha}'$ are equal. When we change the lift of α' by adding a clockwise twist about $C_{p'}$, we multiply the tropical lambda length of the lift by $c(p') = y_\gamma^2/y_\beta^2$. This changes the corresponding lift of α by adding *two* clockwise twists about C_p , thus multiplying the tropical lambda length by $(y_\gamma/y_\beta)^2$. Thus for any particular choice of $\bar{\alpha}'$, the tropical lambda lengths of $\bar{\alpha}$ and $\bar{\alpha}'$ are equal.

We can similarly see that the lambda lengths of $\bar{\alpha}$ and $\bar{\alpha}'$ are equal. In the definition of these lambda lengths, we need the lifts to twist far enough in a particular direction. The direction around $C_{p'}$ and the direction around C_p are related by the projection from the tile cover to (\mathbf{S}, \mathbf{M}) . Thus we can take a lift of α' that twists far enough and projects to a lift of α that twists far enough. In this case, the segment $[\alpha'_{\sigma'}]$ projects to the segment $[\alpha_\sigma]$, where we have written σ' for the hyperbolic metric on $(\mathbf{S}', \mathbf{M}')$. Since σ' is the lift of σ , these two segments have the same signed length and therefore the two lifts have the same lambda length. To change these two lifts to $\bar{\alpha}'$ and $\bar{\alpha}$, we multiply them respectively by $\lambda(p')^k = y_\gamma^{2k}/y_\beta^{2k}$ and $\lambda(p)^{2k} = y_\gamma^{2k}/y_\beta^{2k}$ for some k . Thus $\lambda(\bar{\alpha}') = \lambda(\bar{\alpha})$. (We can make the same argument if both endpoints of α are interior vertices of a self-folded triangle, and if the endpoints of α coincide, the directions of “twisting far enough” are the same at both endpoints of α' .) \square

We now prove the main result of the paper. The proof parallels the proof of Theorem 3.1 and uses some of the same subsidiary results.

Proof of Theorem 1.1. Suppose α is a tagged arc in (\mathbf{S}, \mathbf{M}) . By Proposition 1.7, we can assume that when α has an endpoint that is the interior vertex of a self-folded triangle in T° , the tagging of α is plain at that endpoint. We need to show that

$$\lambda(\alpha) = \mathbf{g}_\alpha \cdot F(P_\alpha).$$

We argue by induction on the number of elements of P_α , more specifically taking this inductive hypothesis: the formula holds for any tagged arc α' in *any* triangulated marked surface such that $P_{\alpha'}$ has fewer elements than P_α . The base of the induction, when P_α is empty, is easy, and details are in the proof of Theorem 3.1.

If P_α is not empty, then construct the tile cover α' in a marked surface $(\mathbf{S}', \mathbf{M}')$ with a triangulation T' . Lift the opening, hyperbolic metric, weighted multilamination \mathbf{A} , and specified laminated lambda lengths to $(\mathbf{S}', \mathbf{M}')$ as explained above. By construction, $P_{\alpha'}$ is isomorphic with P_α as unweighted posets, and (T', α') is tidy. Thus there is an arc $\gamma' \in T'$ such that $\hat{y}_{\gamma'}$ occurs exactly once as a label in $P_{\alpha'}$. Thus, just as in the proof of Theorem 3.1 (Equation (7)), Proposition 2.2 says that

$$(9) \quad x_{\gamma'} \cdot \mathbf{g}_{\alpha'} \cdot F(P_{\alpha'}) = x_{\gamma'} \cdot \mathbf{g}_{\alpha'} \cdot F(P_{\text{blue}}) + x_{\gamma'} \cdot \mathbf{g}_{\alpha'} \cdot \hat{y}_{\text{orange}} \cdot F(P_{\text{red}}).$$

The quantity $F(P_{\text{blue}})$ in Equation (9) is the product of one or two polynomials $F(P_{\delta'})$ for tagged arcs δ' in $(\mathbf{S}', \mathbf{M}')$. Each of these posets $P_{\delta'}$ has strictly fewer elements than P_α . By induction, the laminated lambda length of δ' in $(\mathbf{S}', \mathbf{M}')$ is $\mathbf{g}_{\delta'} \cdot F(P_{\delta'})$, and this is the cluster variable associated to δ' in $(\mathbf{S}', \mathbf{M}')$. Similarly, $F(P_{\text{red}})$ is the product of one or two polynomials $F(P_{\delta'})$

and by induction, the laminate lambda length of each δ' is $\mathbf{g}_{\delta'} \cdot F(P_{\delta'})$, which is again a cluster variable $x_{\delta'}$. Proposition 2.3.1 (applied to α') says that the first term on the right side of Equation (9) is $\mathbf{g}_{\text{blue}} \cdot F(P_{\text{blue}})$, which is the product of the cluster variables for the (one or two) blue arcs. Writing y_{orange} for the product of the coefficient variables y_γ for orange arcs $\gamma \in T$, Proposition 2.3.2 says that the second term on the right side of Equation (9) is $y_{\text{orange}} \cdot \mathbf{g}_{\text{red}} \cdot F(P_{\text{red}})$, which is y_{orange} times the product of the cluster variables for the (one or two) red arcs. Thus the right side of (9) is the right side of the exchange relation exchanging the cluster variables for α' and γ . (See Remark 2.1.) We conclude that the cluster variable for α' in $(\mathbf{S}', \mathbf{M}')$ is $\mathbf{g}_{\alpha'} \cdot F(P_{\alpha'})$. This cluster variable is the laminated lambda length of α' in $(\mathbf{S}', \mathbf{M}')$. Lemma 4.6 now says that the laminated lambda length of α in (\mathbf{S}, \mathbf{M}) is $\mathbf{g}_{\alpha'} \cdot F(P_{\alpha'})$, which equals $\mathbf{g}_\alpha \cdot F(P_\alpha)$, by Lemmas 3.5 and 4.5. \square

5. LAMINATED LAMBDA LENGTHS OF TAGGED GEODESICS

Given a hyperbolic metric on $\mathbf{S} \setminus \mathbf{M}$, we will use the term *tagged geodesic* for any geodesic in \mathbf{S} that limits to points in \mathbf{M} at both ends and is tagged plain or notched at each of its endpoints that is a puncture. (The definition of a tagged arc requires that the arc not intersect itself, that it may not bound a once-punctured monogon, and that, if its two endpoints coincide, they must have the same tagging. For tagged geodesics, we drop all of these requirements, but we keep the requirement that tags at boundary points are plain.) The construction in [FT18a, Definition 15.3], reviewed in Sections 3 and 4, extends to the definition of a (laminated) lambda length x_α for any tagged geodesic α . In this section, we extend Theorems 1.1 and 3.1 to all tagged geodesics. We also give a combinatorial characterization of tagged geodesics.

5.1. Extending the theorems. The following theorem extends Theorem 1.1 by weakening the hypotheses on α .

Theorem 5.1. *The principal-coefficients laminated lambda length of a tagged geodesic α is*

$$(10) \quad x_\alpha = \mathbf{g}_\alpha \cdot \sum_I \prod_{e \in I} w(e),$$

where the sum is over all order ideals I in P_α .

The theorem associates to any tagged geodesic α both a \mathbf{g} -vector and an F -polynomial, in the same sense as for cluster variables. The \mathbf{g} -vector is \mathbf{g}_α and the F -polynomial is $F(P_\alpha)$.

If we set all the y_γ to 1, then $c(p) = 1$ for every puncture, so the opened surface has none of the punctures opened. Thus the lift $\bar{\mathbf{L}}$ is the same as \mathbf{L} and also every tagged geodesic α has $\bar{\alpha} = \alpha$, so lambda lengths in the (not) opened surface are the same as lambda lengths in the coefficient-free model. Also, $c_{\bar{\mathbf{L}}}(\bar{\alpha}) = 1$ for every α , so laminated lambda lengths are the same as ordinary lambda lengths. Thus the entire laminated lambda length construction collapses to the coefficient-free lambda length construction. We conclude that setting all y_γ to 1 in Theorem 5.1 implies the following theorem, which extends Theorem 3.1.

Theorem 5.2. *Suppose all boundary components of (\mathbf{S}, \mathbf{M}) have lambda length 1. The lambda length of a tagged geodesic α is*

$$(11) \quad \lambda(\alpha) = \mathbf{g}_\alpha \cdot \sum_I \prod_{e \in I} \tilde{w}(e),$$

where the sum is over all order ideals I in P_α .

We deduce Theorem 5.1 from Theorem 1.1 using a variation on the tile cover construction.

Proof of Theorem 5.1. Suppose α is a tagged geodesic. If α is a tagged arc, then we are done by Theorem 1.1, so assume α is not a tagged arc. In particular, α does not coincide, except for tagging, with an arc of T° .

We lift α to a tagged arc α' using a variation of the tile cover construction. As in Section 3.2, decompose α into segments, each having its endpoints on arcs of T° and otherwise not intersecting the arcs of T° . Construct a tile for each segment as before, except at any endpoint of α that is

the interior vertex of a self-folded triangle. At each such endpoint, keep the self-folded triangle instead of making a tile. Identify all of these triangles in the natural way.

Lift the opening, multilamination, hyperbolic structure, and specified lamination lengths to the tile cover as in Section 3.2, with the obvious simplifications due to keeping some self-folded triangles. Lemmas 3.5, 4.5 and 4.6 hold by essentially the same proofs. (For Lemma 3.5, the proof is simpler because the endpoints of α remain like the left pictures of Figure 40 instead of the right pictures. For Lemma 4.5, the proof is simpler because there are no double covers to account for.) Since α' is a tagged arc Equation (10) holds for α' by Theorem 1.1. By the analogs of these three lemmas, Equation (10) holds for α . \square

5.2. Combinatorial tagged geodesics. Theorem 5.1 applies to all tagged geodesics in (\mathbf{S}, \mathbf{M}) , but it is more convenient to apply the theorem to curves that are determined combinatorially, rather than by the metric. As a byproduct of the proof of the theorem, we explain how to do that.

A *combinatorial tagged geodesic* in (\mathbf{S}, \mathbf{M}) consists of two (not necessarily distinct) marked points p and q in \mathbf{M} , a tagging at each of the points that is a puncture, and a sequence of arcs of T° , either a single arc in T° with endpoints p and q or a sequence satisfying the following conditions: the sequence starts with the an arc opposite p in some triangle in T° ; the sequence ends with the arc opposite q in some triangle in T° ; every pair of adjacent entries in the sequence are *distinct* arcs in the *same* triangle of T° ; and the sequence of triangles determined by pairs of adjacent entries does not have the same triangle twice in a row. Combinatorial tagged geodesics are defined up to the symmetry of swapping endpoints and reversing the sequence.

A combinatorial tagged geodesic determines, up to isotopy, a tagged curve in \mathbf{S} connecting the two endpoints and crossing arcs of T° according to the sequence (or instead coinciding with the one arc in the sequence). In the other direction, if α is a tagged geodesic in (\mathbf{S}, \mathbf{M}) , then α determines (up to reversing) a sequence of arcs of T° in the obvious way: If α coincides with an arc in T° , then the sequence consists of that arc. Otherwise, the sequence consists of the arcs of T° crossed by α on the way from one endpoint of α to the other.

Proposition 5.3. *If α is a tagged geodesic in (\mathbf{S}, \mathbf{M}) , then the endpoints of α , its taggings, and the sequence determined by α constitute a combinatorial tagged geodesic in (\mathbf{S}, \mathbf{M}) . Conversely, given any combinatorial tagged geodesic in (\mathbf{S}, \mathbf{M}) , the curve it determines is isotopic to a unique tagged geodesic in (\mathbf{S}, \mathbf{M}) .*

Proof. Suppose α is a tagged geodesic in (\mathbf{S}, \mathbf{M}) . If α coincides with an arc in T° , there is nothing to check, so suppose not. Since the arcs in T° are also geodesics and since the hyperbolic metric on (\mathbf{S}, \mathbf{M}) has constant curvature -1 , α does not intersect the same arc in T° twice without intersecting some other arc in between. Thus the sequence of arcs crossed by α satisfies all the conditions of a combinatorial tagged geodesic.

Conversely, given a combinatorial tagged geodesic, let α be a tagged curve in \mathbf{S} determined by the combinatorial tagged geodesic. We construct a tile cover α' for α , ignoring the requirement that α is plain at any endpoint that is the interior point of a self-folded triangle. Since α' is isotopic to an arc in $(\mathbf{S}', \mathbf{M}')$, there is an isotopy representative of α' that is geodesic. For the reasons given in the previous paragraph, this geodesic representative visits the same sequence of arcs of T' as α' . The geodesic descends to a geodesic in (\mathbf{S}, \mathbf{M}) that is isotopic to α . \square

Acknowledgments. We thank Salvatore Stella for helpful conversations.

REFERENCES

- [Bir37] Garrett Birkhoff. Rings of sets. *Duke Math. J.*, 3(3):443–454, 1937.
- [CP03] Gabriel D. Carroll and Gregory Price. Two new combinatorial models for the ptolemy recurrence. Unpublished memo, 2003.
- [CP15] Cesar Ceballos and Vincent Pilaud. Cluster algebras of type D : pseudotriangulations approach. *Electron. J. Combin.*, 22(4):Paper 4.44, 27, 2015.
- [CS21] İlke Çanakçı and Sibylle Schroll. Lattice bijections for string modules, snake graphs and the weak Bruhat order. *Adv. in Appl. Math.*, 126: Paper No. 102094, 22, 2021.
- [FST08] Sergey Fomin, Michael Shapiro, and Dylan Thurston. Cluster algebras and triangulated surfaces. I. Cluster complexes. *Acta Math.*, 201(1):83–146, 2008.

- [FT17] Anna Felikson and Pavel Tumarkin. Bases for cluster algebras from orbifolds. *Adv. Math.*, 318:191–232, 2017.
- [FT18a] Sergey Fomin and Dylan Thurston. Cluster algebras and triangulated surfaces. Part II: Lambda lengths. *Mem. Amer. Math. Soc.*, 255(1223):v+97, 2018.
- [FT18b] Sergey Fomin and Dylan Thurston. Cluster algebras and triangulated surfaces Part II: Lambda lengths. *Mem. Amer. Math. Soc.*, 255(1223):v+97, 2018.
- [FZ03] Sergey Fomin and Andrei Zelevinsky. Y -systems and generalized associahedra. *Ann. of Math. (2)*, 158(3):977–1018, 2003.
- [MS10] Gregg Musiker and Ralf Schiffler. Cluster expansion formulas and perfect matchings. *J. Algebraic Combin.*, 32(2):187–209, 2010.
- [MSW11] Gregg Musiker, Ralf Schiffler, and Lauren Williams. Positivity for cluster algebras from surfaces. *Adv. Math.*, 227(6):2241–2308, 2011.
- [MSW13] Gregg Musiker, Ralf Schiffler, and Lauren Williams. Bases for cluster algebras from surfaces. *Compos. Math.*, 149(2):217–263, 2013.
- [OY22] Ezgi Kantarcı Oğuz and Emine Yıldırım. Cluster algebras and oriented posets. Preprint, [arXiv:2211.08011](https://arxiv.org/abs/2211.08011), 2022.
- [Pro02] James Propp. Lattice structure for orientations of graphs. Preprint, [arXiv:math/0209005](https://arxiv.org/abs/math/0209005), 2002.
- [Rea14] Nathan Reading. Universal geometric cluster algebras from surfaces. *Trans. Amer. Math. Soc.*, 366(12):6647–6685, 2014.
- [Sch08] Ralf Schiffler. A cluster expansion formula (A_n case). *Electron. J. Combin.*, 15(1): Research paper 64, 9, 2008.
- [Sta12] Richard P. Stanley. *Enumerative combinatorics. Volume 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2012.
- [Wen23] Daping Weng. F -polynomials of Donaldson-Thomas transformations. Preprint, [arXiv:2303.03466](https://arxiv.org/abs/2303.03466), 2023.
- [Wil20] Jon Wilson. Surface cluster algebra expansion formulae via loop graphs. Preprint, [arXiv:2006.13218](https://arxiv.org/abs/2006.13218), 2020.
- [Yur19a] Toshiya Yurikusa. Cluster expansion formulas in type A. *Algebr. Represent. Theory*, 22(1):1–19, 2019.
- [Yur19b] Toshiya Yurikusa. Combinatorial cluster expansion formulas from triangulated surfaces. *Electron. J. Combin.*, 26(2):Paper No. 2.33, 39, 2019.

CNRS & LIX, ÉCOLE POLYTECHNIQUE, PALAISEAU
 Email address: vincent.pilaud@lix.polytechnique.fr
 URL: <http://www.lix.polytechnique.fr/~pilaud/>

NORTH CAROLINA STATE UNIVERSITY
 Email address: reading@math.ncsu.edu
 URL: <https://nreadin.math.ncsu.edu>

UNIVERSITÄT ZU KÖLN, GERMANY AND NTNU, NORWAY
 Email address: schroll@math.uni-koeln.de
 URL: <https://sites.google.com/site/sibylleschroll/>