

Two New Combinatorial Models for the Ptolemy Recurrence

Gabriel D. Carroll
Gregory Price
Harvard University

April 14, 2003

Abstract

This is an abstract that has not been written yet. You might think of it as an abstract abstract.

1 Introduction

Let \mathcal{T} be an arbitrary triangulation of a convex n -gon. To each edge e of the triangulation, we associate a formal variable x_e . We now define multivariate rational functions $Q(v, w)$, for any distinct vertices v, w of the polygon, as follows:

- if v, w are the endpoints of an edge e of the triangulation (either a side or diagonal of the n -gon), then $Q(v, w) = Q(w, v) = x_e$;
- if t, u, v, w are distinct vertices of the polygon in cyclic order, then

$$Q(t, u)Q(v, w) + Q(t, w)Q(u, v) = Q(t, v)Q(w, u). \quad (1)$$

Thus, we associate rational functions with all the edges and diagonals of the polygon. If we have functions associated with the sides and one of the diagonals of a quadrilateral formed by four vertices of the polygon, then our recurrence allows us to associate a function with the other diagonal. It is easy to check that every diagonal of the original polygon can be reached by a succession of such operations, and it is also easy to check that the computation never requires division by zero (for example, just plug in 1 for every x_e , and then it follows by induction that all the functions obtained in this manner will take on positive values).

What is perhaps not so obvious is that each edge and diagonal will have just one function associated to it — that is, the rational function $Q(v, w)$ is uniquely determined, independent of the sequence of applications of (1) used to obtain it. Furthermore, these functions are all Laurent polynomials in the initial variables x_e (i.e. polynomials in x_e, x_e^{-1}),

in which every term has coefficient 1 and has each x_e raised to the power -1 or 1 (or zero).

(Diagram of a polygon, with thingies associated to all its edges and such)

At least partial results along these lines have been known for some time. The uniqueness of the $Q(v, w)$ follows e.g. from interpreting them as lengths of hyperbolic geodesics (see [6]). Indeed, (1) is strikingly reminiscent of Ptolemy’s theorem from Euclidean geometry, relating the sides and diagonals of a cyclic quadrilateral — hence, we refer to (1) as the “Ptolemy recurrence.” However, the Euclidean setting is not enough to prove uniqueness of the $Q(u, v)$ for arbitrary values of the initial variables x_e , since the sides and diagonals of a cyclic polygon are algebraically dependent. Using a slightly different but equivalent form of the recurrence, Prešič and Mikrinovič, in [8], also obtained uniqueness by showing that the only possible solutions to the recurrence are of the form

$$Q(r, s) = \pm \begin{vmatrix} f(r) & f(s) \\ g(r) & g(s) \end{vmatrix}$$

for suitable functions f, g , and it is easy to verify that their solution satisfies the Ptolemy recurrence for all quadruples of vertices. Yet another proof of uniqueness comes from interpreting $Q(v, w)$ as Plücker coordinates on the Grassmannian of all 2-dimensional subspaces of \mathbb{C}^n , which is defined precisely by the relations given by (1); see [5] for details.

That the rational functions $Q(v, w)$ are Laurent polynomials in the initial variables x_e , with coefficients in \mathbb{Z} — regardless of the triangulation used — is a more recent result. This fact follows from work in [4] and [5], where Fomin and Zelevinsky identified these functions as cluster variables in the cluster algebra of type A_{n-3} . In [4], Fomin and Zelevinsky conjectured that all coefficients in these Laurent polynomials are positive, though their methods do not suffice to prove positivity results. The article [3], by Conway and Coxeter, is also loosely related; it seems to be the earliest connection between any form of the recurrence (1) and triangulated polygons in the literature, and it effectively demonstrates that all the $Q(v, w)$ are positive integers when each x_e is set to 1 (although Conway and Coxeter were not thinking in terms of associating actual variables with each edge).

Our purpose here is to present two new, simple graph-theoretical models that are described by the Ptolemy recurrence. These two models, we hope, will more immediately explain why the $Q(v, w)$ are all Laurent polynomials, and at least one of them will provide another explanation of their uniqueness. In addition, the models will quickly prove that each Laurent polynomial has all coefficients equal to 1 and all exponents in the range $[-1, 1]$, generalizing Fomin and Zelevinsky’s conjecture.

2 The matchings model

We present the more complicated of our two models first, because it provides a new proof of the uniqueness property effectively from scratch.

Construct a bipartite graph G as follows: for each vertex v of the polygon, we define a vertex a_v of the graph; for each triangle Δ of the triangulation \mathcal{T} , we define a vertex a_Δ of the graph. Connect a_v to a_Δ by an edge precisely when v is a vertex of Δ . For any two distinct vertices v and w of the polygon, the graph $G - \{a_v, a_w\}$ (which we henceforth abbreviate $G_{v,w}$) has $n - 2$ vertices associated to vertices of the polygon and $n - 2$ vertices associated to triangles; thus it makes sense to consider its perfect matchings. If M is a matching of G , and e is an edge of the triangulation, define $w_M(e)$ to be the number of edges $a_v a_\Delta$ of G used in the matching such that v is an endpoint of e and e is a side of Δ . Define

$$m(M) = \prod_e x_e^{1-w_M(e)},$$

where the product is taken over all edges e of the triangulation. Finally, define

$$P(v, w) = \sum m(M),$$

where the sum ranges over all perfect matchings M of $G_{v,w}$.

(Diagram of G , some matching, and the corresponding monomial)

We claim that when $vw = e$ is an edge of the triangulation, $P(v, w) = x_e$, and furthermore that the polynomials P satisfy the Ptolemy recurrence (1) for all t, u, v, w . It then immediately follows that $Q(v, w) = P(v, w)$ for all v and w . The uniqueness of the $Q(v, w)$ thus follows. Furthermore, for any edge e and any matching M , $0 \leq w_M(e) \leq 2$. (The upper bound holds because the number of edges $a_v a_\Delta \in E(G)$ such that v is an endpoint of e and e is a side of Δ is either 2 or 4, depending on whether e is a side of 1 or 2 triangles in \mathcal{T} ; if there are 4 such edges, however, they form a cycle, so at most 2 of them can be used in the matching.) The statement that every variable appears in the polynomials $P(v, w)$ with only exponents ± 1 then follows. It is also immediate from our model that all coefficients of the polynomials are positive integers. It may not be obvious that the coefficients are all 1; our second model will make this more transparent.

Before we prove the claim, we need one definition: Call a vertex v of the triangulation a *hanging* vertex if it and its two neighbors form the vertices of a triangle of \mathcal{T} . Note that our triangulation must have at least two hanging vertices. (For example, there are n edges of the polygon, each of which belongs to one of $n - 2$ triangles, so we get at least 2 pairs of edges that belong to the same triangle. If $n > 3$, this means we have 2 triangles, each of which has two of its edges as edges of the n -gon, and each such triangle gives us a hanging vertex.)

Now, the first part of the claim — that when $P(v, w) = x_e$ when $vw = e$ is an edge of the triangulation — is proven by a straightforward induction on n . If $n = 3$, then there is only one matching M of $G_{v,w}$, and it is easy to check that $m(M) = x_e$.

(Diagram of the base case)

When $n > 3$, there exist at least two hanging vertices. These two vertices cannot be v and w , since then they could not be connected by an edge of the triangulation. Hence, we have a hanging vertex u distinct from v and w . Then u belongs to just one triangle Δ of \mathcal{T} . Let G' be the graph constructed from the $(n - 1)$ -gon whose vertices are the vertices of the original n -gon other than u , following the same construction used to build G from the n -gon; we can do this, since $\mathcal{T} - \{\Delta\}$ is a triangulation of this $(n - 1)$ -gon. Now, in any perfect matching M of $G_{v,w}$, a_u must be matched with a_Δ . Then the edges connecting a_Δ to the two vertices of the n -gon adjacent to u are not used in the matching. It follows that if we delete the vertices a_u, a_Δ from M , we get a matching M' of the graph $G'_{v,w} = G' - a_v, a_w$. Conversely, any matching M' of the graph $G'_{v,w}$ can be extended to a matching M of $G_{v,w}$ by including the edge $a_u a_\Delta$. Moreover, $m(M) = m(M')$; this follows from the evident identities $w_M(a_s a_t) = w_{M'}(a_s a_t)$ and $w_M(a_s a_u) = w_{M'}(a_s a_u) = 1$, where s, t are the vertices of the n -gon adjacent to u . So we have a bijection between the matchings M of $G_{v,w}$ and the matchings M' of $G'_{v,w}$, so by the induction hypothesis, there is just one of each, and $P(v, w) = m(M) = m(M') = x_e$.

(Diagram of the induction step)

The second part of the claim — that the polynomials $P(v, w)$ satisfy the Ptolemy recurrence — is virtually identical to Theorem 5.4 from Eric Kuo's paper [7]. However, because our weighting scheme is different from Kuo's (and for the sake of completeness), we include the proof here.

Given our vertices t, u, v, w in cyclic order, define a *doubled matching* to be a multiset of edges of G such that a_t, a_u, a_v, a_w each belong to exactly one of these edges, and every other vertex belongs to exactly two. Notice that the union (by multiplicity) of a matching of $G_{t,u}$ and a matching of $G_{v,w}$ is a doubled matching, as is the union of a matching of $G_{t,w}$ and a matching of $G_{u,v}$, or the union of a matching of $G_{t,v}$ and a matching of $G_{u,w}$. Moreover, if M is a doubled matching, then for each edge e of \mathcal{T} , we can define $w_M(e)$ as for ordinary matchings (again counting edges of M by multiplicity), and we can define $\hat{m}(M) = \prod_e x_e^{2-w_M(e)}$, where the product is over all edges e of the triangulation. It is immediately apparent that, if M_1 is a matching of $G_{t,u}$ and M_2 is a matching of $G_{v,w}$, then $\hat{m}(M_1 \cup M_2) = m(M_1)m(M_2)$, and similarly for the other two pairings of the vertices t, u, v, w .

Now, for any doubled matching M , let $N(M)$ be the number of ways of expressing M as the union of a matching of $G_{t,v}$ and a matching of $G_{u,w}$. The multiplicativity property above then implies that

$$\sum_M N(M) \hat{m}(M) = P(t, v)P(u, w),$$

where the sum ranges over all doubled matchings M . If we can show that $N(M)$ is also the total number of ways of expressing M either as the union of a matching of $G_{t,u}$ and a matching of $G_{v,w}$ or as the union of a matching of $G_{t,w}$ and one of $G_{u,v}$, then it will similarly follow that

$$\sum N(M)\hat{m}(M) = P(t, u)P(v, w) + P(t, w)P(u, v),$$

and we can then deduce the Ptolemy recurrence. Hence, we now turn to this task.

In M , viewed as a multigraph, every vertex of G has degree 2, except for a_t, a_u, a_v, a_w , which each have degree 1. It follows that M can be partitioned into double edges, cycles not using any of the vertices t, u, v, w , and two paths, whose vertices are a_t, a_u, a_v, a_w in some order. Moreover, we cannot have one path running from a_t to a_v and the other path running from a_u to a_w : from the planarity of the graph G and the ordering of the points around the n -gon, we see that these two paths would have to intersect at some vertex of G , but then this vertex would have degree at least 3 in M (4 if it is not among the endpoints a_t, a_u, a_v, a_w). Since every vertex has degree at most 2 in M , this is a contradiction. So our paths either run from a_t to a_u and a_v to a_w , or from a_t to a_w and a_u to a_v .

Suppose that we have a path from a_t to a_u and a path from a_v to a_w . Then we claim there are no decompositions of M as the union of a matching of $G_{t,u}$ and one of $G_{v,w}$. Indeed, the edges of either path must alternate between the two matchings. However, each path has even length (because G is bipartite), so (say) the path from a_t to a_u must have both its first edge and its last edge belonging to the matching of $G_{v,w}$. This is a contradiction.

On the other hand, if we wish to decompose M into a matching of $G_{t,w}$ and one of $G_{u,v}$, then we proceed as follows: Each doubled edge of M must occur in both matchings; each of the two paths can also be decomposed in a unique way — the edge containing (say) a_t must be used in the matching of $G_{u,v}$, and then successive edges along the path must alternate between the two matchings. The last edge, which contains a_u , is then assigned to the matching of $G_{t,w}$, as required. Similarly, the last path can be safely decomposed in a unique way. Finally, each cycle has even length (again by bipartiteness) and uses none of the vertices a_t, a_u, a_v, a_w , so we can assign the edges alternately to the two matchings, starting with either one. It is easy to check that this process does indeed give us perfect matchings of the two graphs $G_{t,w}, G_{u,v}$. The decomposition algorithm involved a two-way choice for each cycle in M ; thus, the number of expressions of M as the union of a matching of $G_{t,w}$ and one of $G_{u,v}$ is 2^C , where C is the number of cycles in M .

But the same argument shows that the number of expressions of M as the union of a matching of $G_{t,v}$ and a matching of $G_{u,w}$ is also 2^C — that is, $N(M) = 2^C$. Since we saw that there were 0 decompositions into a matching of $G_{t,u}$ and a matching of $G_{v,w}$, we see that $N(M)$ is indeed the total number of decompositions into these two types of matchings and into matchings of $G_{t,w}$ and $G_{u,v}$, as we sought to prove. Moreover, exactly the same argument, with the roles of t and v switched, shows that the same result occurs when M has paths from a_t to a_w and a_u to a_v . So the result holds for any M , and summing over all M gives

$$P(t, u)P(v, w) + P(t, w)P(u, v) = \sum N(M)\hat{m}(M) = P(t, v)P(u, w),$$

completing the proof of the claim.

We summarize our findings as follows:

Theorem 1 *Let $Q(v, w) = x_e$ when v, w are the endpoints of an edge e of the triangulation, and successively define values of $Q(v, w)$ for other pairs of vertices v, w via the recurrence (1). Then, for any distinct v and w , $Q(v, w) = \sum m(M)$, where M ranges over all perfect matchings of $G_{v,w}$, with m, G defined as above. This value is obtained regardless of the sequence of applications of (1) used to compute $Q(v, w)$.*

□

Next, before proceeding to our second model, we make a definition. For distinct vertices v, w , consider the triangles of \mathcal{T} that meet the line segment from v to w . We can see that these triangles form a triangulation of a polygon whose vertices are among the vertices of the n -gon. (This holds by a sort of induction — as we walk along the segment from v to w , at any point, the triangles so far form a triangulation of such a polygon, and this remains true when we cross an edge of the triangulation.) We call this polygon the *strip* $S_{v,w}$ from v to w . If vw is an edge of \mathcal{T} , we adopt the convention that the strip from v to w is just that edge.

(Diagram showing two vertices and the strip between them)

Because the strip is itself a polygon with v and w as vertices, and equipped with a triangulation induced by \mathcal{T} , we can compute $Q(t, u)$, where t, u are any two vertices of $S_{v,w}$, by the recurrence (1). It is clear that this gives the same result as running the recurrence on the original n -gon in its entirety, using the same sequence of quadruples of vertices. In particular, we can construct the graph H on the vertices and triangles of the strip $S_{v,w}$ (as in the construction of G above). Then the perfect matchings of $H_{v,w} = H - \{a_v, a_w\}$ are naturally in bijection with the perfect matchings of $G_{v,w}$, since both are in bijection with the terms of the Laurent polynomial $Q(v, w)$, which is the same regardless of whether we use the strip or the larger polygon to compute it. In fact, each matching of $G_{v,w}$ can be obtained by extending the corresponding matching of $H_{v,w}$, and the extension is the same for all matchings. We will not provide a detailed proof here, but the argument is essentially an induction by successive removals of hanging vertices, much as in the proof of the first part of Theorem 1 above. The key observation is simply that a triangle of \mathcal{T} with a hanging vertex distinct from v and w cannot be part of the strip: if the line segment from v to w intersected such a triangle, it would have to pass through two of its sides, but two sides are sides of the n -gon and so cannot meet such a segment.

The fact that $Q(v, w)$ is the same regardless of whether we calculate it in $S_{v,w}$ or in the larger n -gon gives rise to another simple observation:

Proposition 1 *If e is an edge of \mathcal{T} such that x_e appears in the Laurent polynomial $Q(v, w)$, then e is an edge of one of the triangles of $S_{v,w}$.*

□

Corollary 1 *If M is any matching of $G_{v,w}$, and e is an edge of \mathcal{T} not occurring in the strip $S_{v,w}$, then $w_M(e) = 1$.*

□

3 The trails model

Our next combinatorial model uses a graph whose vertices and edges are actually those of the triangulation \mathcal{T} . For convenience, we will presently assume the vertices v and w are fixed, and the triangulation is such that the strip from v to w is the entire n -gon (equivalently, v and w are the only hanging vertices). We may assume this, as otherwise we can chop off extra triangles without affecting the value of $Q(v, w)$, as we have seen.

We construct a graph whose vertices are those of the n -gon, and whose edges are precisely the edges of triangles in the triangulation \mathcal{T} . We refer to the edges of the graph that are sides of the n -gon as *exterior* edges, and those that are diagonals as *interior* edges. Direct the exterior edges so as to lead from v to w along the boundary of the polygon; leave the interior edges undirected. Now consider trails running from v to w in this graph, where the exterior edges must be followed in the direction indicated. (A *trail* is like a path, except that vertices — but not edges — may be used more than once. This follows standard terminology; see e.g. [1] or [2].) Such a trail will be called a *zigzag trail* if the 2nd, 4th, 6th, . . . edges are all interior edges, and there are no two distinct vertices t, u on the same side of line vw , with the edges leading from t to u , such that the trail visits u and later visits t .

If the edges $e_1, e_2, e_3, \dots, e_k$ form a zigzag trail T , then k is necessarily odd, since the last edge leads to the hanging vertex w and so is an exterior edge. Let

$$p(T) = x_{e_1} x_{e_2}^{-1} x_{e_3} x_{e_4}^{-1} x_{e_5} \cdots x_{e_k}.$$

We then claim:

Theorem 2 *If the strip $S_{v,w}$ is the entire n -gon, then*

$$Q(v, w) = \sum p(T),$$

where the sum is over all zigzag trails from v to w .

Unlike with the previous model, our proof now assumes that the $Q(v, w)$ are all unique, although it assumes nothing further.

Before we can prove this theorem, we first need a lemma on the structure of substrips of $S_{v,w}$.

Lemma 1 *Let s, t be the two vertices of the n -gon adjacent to w . Suppose that the vertices joined to t by interior edges are $s_0, s_1, \dots, s_k = s$, and that the direction of the exterior edges leads from s_i to s_{i+1} for each i . Denote $s_{k+1} = w$. Then, s_i, s_{i+1} are successive vertices of the n -gon and $s_i s_{i+1} t$ is a triangle of \mathcal{T} for each $i = 0, \dots, k$, and the strip $S_{v,t}$ is composed of those triangles of \mathcal{T} not of the form $s_i s_{i+1} t$.*

Proof: Since s_i, s_{i+1} are consecutive among the vertices connected to t by an interior edge, some triangle of \mathcal{T} must have $\angle s_i t s_{i+1}$ as one of its angles, so this triangle must be $s_i s_{i+1} t$, giving the second claim. Between the intersections of segment vw with $s_i t$ and with $s_{i+1} t$, it can intersect no other diagonal of \mathcal{T} (because such a diagonal would have to have t as one of its endpoints); therefore, every triangle of \mathcal{T} other than $s_i t s_{i+1}$ lies on the opposite side of line $s_i t$ from s_{i+1} , or on the opposite side of line $s_{i+1} t$ from s_i . It follows that all vertices of the n -gon lie in one of these two regions, so there can be none between s_i and s_{i+1} , proving the first claim.

For the third, imagine a variable point p moving along segment wt from w to t . As p moves, the collection of triangles of \mathcal{T} met by line vp can only change when this line encounters a vertex of such a triangle. However, any such vertex is a vertex of our n -gon, and the only such vertices the line ever meets are v, w , and t . Thus, any triangle of $S_{v,w}$ not in $S_{v,t}$ must have one of these three as a vertex. Now, it is easy to see that $S_{v,w}$ and $S_{v,t}$ each contain exactly one triangle of \mathcal{T} with v as a vertex (the same one), and $S_{v,w}$ contains just one triangle of \mathcal{T} with w as a vertex, namely swt . So all the triangles in $S_{v,w}$ not in $S_{v,t}$ must have t as a vertex. Let r be the vertex of the n -gon adjacent to t and distinct from w . Then, $S_{v,t}$ should have exactly one triangle of \mathcal{T} with t as a vertex, and since rts_0 is a triangle of \mathcal{T} (by the same argument used for the $s_i t s_{i+1}$ above) and v lies inside angle rts_0 , this triangle is the one. The result then follows. \square

(Diagram of the strip $S_{v,w}$, with $S_{v,t}$ shaded, to make the argument clear)

Proof of Theorem 2: We use an induction again, on the number of triangles in the strip. If there are 0 triangles, so that the strip consists of a single edge e , then there is only one zigzag trail T , and $p(T) = x_e$, giving the base case. Otherwise, let s, t be the two vertices of the n -gon adjacent to w . We have from (1) that

$$\begin{aligned} Q(v, w) &= (Q(v, s)Q(t, w) + Q(v, t)Q(w, s))/Q(s, t) \\ &= Q(v, s)x_{tw}/x_{st} + Q(v, t)x_{sw}/x_{st}. \end{aligned}$$

, and from the induction hypothesis, $Q(v, s)$ is the sum of $p(T)$ where T ranges over all zigzag trails in the strip $S_{v,s}$, and $Q(v, t)$ is the sum of $p(T)$ over all zigzag trails in the strip $S_{v,t}$. (We can indeed apply the induction hypothesis here — these strips omit at least the one triangle having w as a vertex and so contain fewer triangles than the original n -gon.)

(Diagram of the triangulation, with directed exterior edges and v, w, s, t labeled, and a zigzag trail shown)

Every zigzag trail from v to w has either sw or tw as its last edge. Consider any trail T ending in the edge sw . If the previous edge was not ts , then it was rs for some r on the same side of line vw as t (by the interior edge condition on zigzag trails); then t cannot

have been visited at any earlier time on the trail (by the condition on the ordering of vertices). Replace the final edge sw by st . We thus get a trail T' from v to t , and we claim T' is in fact a zigzag trail in $S_{v,t}$.

To see this, notice that t cannot belong to any interior edge other than st , for such an edge would intersect the interior edge rs , which is prohibited. Then, by the lemma, $S_{v,t}$ consists of the triangles of \mathcal{T} , except for the single triangle swt . Since T' never uses either of the edges sw, tw , all of its edges do indeed belong to the strip $S_{v,t}$. We need to check that every even-positioned edge of T' is an interior edge of this strip. However, each such edge is an interior edge of $S_{v,w}$, and all interior edges of $S_{v,w}$ are interior edges of $S_{v,t}$, except for the edge st . Since st is the last edge of T' (thus in an odd position) and cannot occur elsewhere in T' , the condition is verified. Similarly, we see that every exterior edge of $S_{v,t}$ is traversed in the required direction — this property is inherited from $S_{v,w}$ for every edge except st (the one new exterior edge), and we see that st is indeed traversed from s to t , as required. Finally, the condition on ordering of vertices in T' follows from the same condition for T . So, we have verified that T' is a zigzag trail in $S_{v,t}$.

(Diagram of this case)

Now, if the edge of T preceding sw was ts , then remove both of these edges to obtain a trail T' from v to t . We again claim T' is a zigzag trail in $S_{v,t}$, but the verification is slightly more involved. Let $s_0, s_1, \dots, s_k = s$ be the vertices of the n -gon connected to t by interior edges and $s_{k+1} = w$, as in the statement of the lemma. Then, from the lemma, we know that the edges of the induced triangulation of the strip $S_{v,t}$ are the edges of \mathcal{T} other than $s_i s_{i+1}$ and $s_{i+1} t$ ($i = 0, \dots, k$), and $s_0 t$ is the only interior edge of $S_{v,w}$ that becomes an exterior edge of $S_{v,t}$.

The edge of T preceding ts cannot be $s_i t$ for any $i > 0$: otherwise, the edge preceding this in turn would have to be $s_{i-1} s_i$, but this edge is an exterior edge of $S_{v,w}$. However, it would occur in the fourth-from-last position in T , where an interior edge should appear — contradiction. Therefore, the edge of T preceding ts must be $s_0 t$ or rt , where r is the vertex of the n -gon adjacent to t and distinct from w . In the latter case, the edge preceding rt must then be an interior edge, so the vertex visited just before r is on the opposite side of line vw from t , and it cannot be any of s_1, \dots, s_k , as none of these are connected to r by an edge. It now follows from the condition on ordering of vertices that none of s_1, \dots, s_{k-1} are ever used in T ; therefore, T' does not use any of the edges $s_i s_{i+1}$ or $s_{i+1} t$ ($i = 0, \dots, k$). Thus, all the edges of T' really are contained in the strip $S_{v,t}$.

As in the previous case, we also need to check that the one new exterior edge $s_0 t$, if it is traversed at all, is traversed in the correct direction and occurs in an odd position. However, by the foregoing argument, this edge can only occur if it immediately precedes ts , so it is in third-from-last position in T (and thus last position in T'), and it is traversed in this direction. Finally, the condition on ordering of vertices for T' follows from the same condition for T . Thus, T' is indeed a zigzag trail in $S_{v,t}$.

So we have turned any zigzag trail T from v to w , ending in the edge sw , into a zigzag trail T' from v to t in $S_{v,t}$. Moreover, this process is reversible: if T' ends in st , replace it

(Diagram of the second case)

by sw to obtain a zigzag trail T in $S_{v,w}$; otherwise, the edge st can never have been used, so add it and tw to the end of the trail to obtain our zigzag trail T . It is straightforward to check that this operation is the inverse of that previously described. Thus, we have a bijection between zigzag trails T from v to w ending in sw and zigzag trails T' from v to t . Also, it is easily checked $p(T) = p(T')x_{sw}/x_{st}$ for any corresponding trails T, T' .

By a similar argument, we have a bijection between zigzag trails T from v to w ending in tw and zigzag trails T' from v to s (in $S_{v,s}$), satisfying $p(T) = p(T')x_{tw}/x_{st}$. Since every zigzag trail from v to w in $S_{v,w}$ ends in one of the edges sw, tw , we can sum over all such trails to find

$$\sum p(T) = Q(v, t)x_{sw}/x_{st} + Q(v, s)x_{tw}/x_{st} = Q(v, w),$$

proving the theorem. \square

This model further demonstrates that $Q(v, w)$ is a Laurent polynomial in the edge variables, with each variable taking on exponents of ± 1 , as we already saw with the previous model. Furthermore, the condition on ordering of vertices makes it clear that there is at most one way to arrange a given set of edges into a zigzag trail, from which it follows that each coefficient of this Laurent polynomial is 1. We summarize these results formally:

Theorem 3 *In any Laurent polynomial $Q(v, w)$ obtained via the Ptolemy recurrence (1), every coefficient is 1, and every variable takes on only the exponents 1 and -1 .*

We might also remark that these polynomials are homogeneous of degree 1 (since each $p(T)$ has this degree), although this actually was already self-evident, since the recurrence is homogeneous of degree 2 in the $Q(v, w)$'s, and all the initial conditions are of degree 1.

4 The exterior edges

We continue to assume that v, w are fixed such that $S_{v,w}$ is the entire n -gon.

We now have two classes of objects — matchings of $G_{v,w}$ and zigzag trails in $S_{v,w}$ — each of which is in bijection with the terms of the Laurent polynomial $Q(v, w)$; therefore, they are in bijection with each other. We would like to describe this latter correspondence explicitly. In one direction, the map is obvious: given a matching M , we can simply read off from the monomial $m(M)$ which edges appear in the corresponding trail (namely, the edges e such that x_e appears in $m(M)$), and then there is just one way to arrange those edges into a zigzag trail. The map in the other direction is not as obvious. We could attempt an inductive construction of the matching from the zigzag trail, by our usual method of peeling off triangles with hanging vertices. However, we will not pursue this here. Instead, we will investigate a particular detail of the correspondence that leads into a further discovery about $Q(v, w)$.

Proposition 2 *Let t, u be vertices such that tu is an edge of the n -gon, and let s be the other vertex of the (unique) triangle Δ of \mathcal{T} having tu as an edge. Then, a_Δ is joined to a_s in a matching M if and only if the edge tu appears in the corresponding trail T .*

Proof: The edge $a_\Delta a_s$ appears in the matching precisely when neither of the edges $a_\Delta a_t, a_\Delta a_u$ does, which in turn is equivalent to $w_M(tu) = 0$, so that x_{tu} has exponent 1 in $m(M)$. Also, this edge appears in T if and only if it has exponent 1 in the $p(T)$ (it cannot have exponent -1 , since exterior edges can appear only in odd positions). So both statements are equivalent to x_{tu} having exponent 1 in $m(M) = p(T)$. \square

(Diagram of a zigzag trail and the associated matching, highlighting exterior edges and the corresponding matching edges)

So, given at least part of a zigzag trail, the corresponding part of the associated matching can be immediately determined (and vice versa). What is interesting is that this effectively determines the entire correspondence — that is, a zigzag trail is entirely determined by the exterior edges it contains (and, equivalently, a matching of $G_{v,w}$ is entirely determined by the edges of the form $a_\Delta a_s$, as above, that it contains). This statement also has an immediate algebraic interpretation, which is a strengthening of the statement about coefficients in Theorem 3.

Theorem 4 *A zigzag trail in $S_{v,w}$ is uniquely determined by the collection of exterior edges it traverses. Equivalently: If we set $x_e = 1$ for every interior edge e , and regard $Q(v,w)$ as a Laurent polynomial in the variables x_e for exterior edges e , then every coefficient is still 1.*

Proof: Divide the exterior edges into two “arcs” separated by the line vw ; let e_1, \dots, e_k be the exterior edges used by the trail T on one side of this line, arranged in order around the polygon from v to w , and let f_1, \dots, f_l be the exterior edges (if any) used by the trail on the other side of the line. The first edge of the trail must be an exterior edge; without loss of generality it is e_1 . Now, by the condition on ordering of vertices, e_1, \dots, e_k must appear in T in that order, and similarly f_1, \dots, f_l must appear in T in that order.

Now, every exterior edge in T appears in an odd position, so any two successive such edges must be separated by an odd number of interior edges. However, every interior edge has its two endpoints on opposite sides of line vw , so traversing an odd number of such edges entails crossing this line an odd number of times. Therefore, any two successive exterior edges in T are on opposite sides of vw .

But this means that $l = k - 1$ or k , and the exterior edges of T must be traversed in the order

$$e_1, f_1, e_2, f_2, \dots, e_k \text{ or } e_1, f_1, e_2, f_2, \dots, e_k, f_k,$$

accordingly. Moreover, if we consider just the interior edges of the triangulation \mathcal{T} , the subgraph formed by these edges is easily seen to be acyclic (if there were a cycle, some diagonal of the polygon would have to cross another), so there is at most one trail along

interior edges from the endpoint of one exterior edge e_i (resp. f_i) to the beginning point of the next edge f_i (resp. e_{i+1}). Putting all these trails together, we see that there is just one zigzag trail that uses the exterior edges e_1, f_1, \dots, e_k (or \dots, e_k, f_k) in that order, and no others. This proves the result. \square

(Maybe a diagram of external edges, showing the order in which they must be connected?)

We close this section with one more easy fact that should not go unnoticed.

Proposition 3 *If we set x_e to 1 whenever e is an interior edge, then $Q(v, w)$ is in fact a polynomial (not just a Laurent polynomial) in the remaining variables x_e , and it is linear in each such variable.*

Proof: As observed earlier, if e is any exterior edge, it can only appear in an odd position in any zigzag trail T , so it can only have exponent $+1$ in the monomial $p(T)$, if it appears at all. \square

5 Conclusion

We have developed two models that give us explicit bijections between classes of objects — one class of trails, another class of matchings — and the terms of the Laurent polynomials $Q(v, w)$ generated by the Ptolemy recurrence. In the process, we have obtained new information about the structure of these polynomials; we hope that these models also provide some new intuitions as to why the $Q(v, w)$ are Laurent polynomials in the first place, and why they are uniquely determined. It would be interesting to compare these models with preexisting incarnations of the Ptolemy recurrence, such as hyperbolic lengths or Plücker coordinates, and see if any further correspondences can be found. However, the prospects for fruitful comparisons here are somewhat doubtful, as these other models do not provide interpretations of the individual terms of the polynomials, only of the polynomials' numerical values.

The Ptolemy recurrence and its Laurent property arose in the work of Fomin and Zelevinsky ([5]) in connection with the study of finite-type cluster algebras; this particular recurrence corresponds to the cluster algebra of type A_n . This connection then suggests the question (which Dylan Thurston asks explicitly, [9]) of whether similar models exist for other finite-type cluster algebras, so that their corresponding recurrences can similarly be understood combinatorially. Such models would, we expect, provide similar insights into the coefficients and exponents of the Laurent polynomials arising from these cluster algebras. This seems to be a promising direction for future investigation.

6 Acknowledgments

We are grateful to the National Science Foundation and the National Security Agency for their financial support of our research, and to the mathematics departments of Harvard University and the University of Wisconsin-Madison for administering the funds. We also would like to thank the other members of the REACH (Research Experiences in Algebraic Combinatorics at Harvard) group who have helped in our work on this problem, and especially Dylan Thurston for providing background information and for helping to introduce the problem to us. Finally, we thank Jim Propp, the director of REACH, for providing guidance and suggestions throughout our research. And we thank the theorems for being true. Or maybe we don't.

References

- [1] Bollobás, *Graph Theory*.
- [2] Chartrand, *Introductory Graph Theory*.
- [3] Conway and Coxeter's frieze pattern article.
- [4] Fomin and Zelevinsky, "Cluster Algebras I: Foundations." (Must make sure the claims made about what these dudes said are accurate!)
- [5] Fomin and Zelevinsky, "Cluster Algebras II: Finite Type Classification."
- [6] Something about hyperbolic geometry, once I figure out what to put here.
- [7] E. Kuo, "Applications of Graphical Condensation for Enumerating Matchings and Tilings," <http://www.cs.berkeley.edu/~kuo/condensation.ps> .
- [8] Prešić and Mikrinović, "Sur une équation fonctionnelle d'ordre supérieure," *Publikacije Elektroteh. Fak. Unic. Beogr. Ser. Math.-Phys.* **70** (1962), or wherever else I find this proof of the formula in terms of determinant stuff.
- [9] Dylan Thurston, personal communication.