

Stokes posets and serpent nests

F. Chapoton*

Université de Strasbourg, CNRS, IRMA UMR 7501

received 23rd Feb. 2016, revised 8th Nov. 2016, accepted 16th Nov. 2016.

We study two different objects attached to an arbitrary quadrangulation of a regular polygon. The first one is a poset, closely related to the Stokes polytopes introduced by Baryshnikov. The second one is a set of some paths configurations inside the quadrangulation, satisfying some specific constraints. These objects provide a generalisation of the existing combinatorics of cluster algebras and nonnesting partitions of type \mathbb{A} .

Keywords: poset, quadrangulation, Tamari lattice

Introduction

The research leading to this article started with the desire to understand the apparition of associahedra in an article of Baryshnikov (2001) and in an article of Kapranov and Saito (1999). Associahedra, also known as Stasheff polytopes, are classical in algebraic topology, where they are used to define and study associativity up to homotopy, see Stasheff (1963). More recently, they have made their way into the theory of cluster algebras, in which they control the combinatorial aspects of the type \mathbb{A} cluster algebras, see Fomin and Zelevinsky (2003). It was not clearly obvious how to relate the contexts used by Baryshnikov and Kapranov-Saito to either associativity or cluster algebra theory. Moreover both articles were seeing associahedra as members of a larger family of polytopes, all closely tied together, whereas associahedra usually stand alone.

Another strand of research then came into play, centred on the Tamari lattices (see the reference Müller-Hoissen et al. (2012)). The Tamari lattices are closely related to associahedra and associativity, and can be defined by orienting the associativity relation from left-bracketing to right-bracketing. They are also well-understood in the context of cluster algebras, where they are just one among several Cambrian lattices of type \mathbb{A} , depending on the choice of an orientation of the Dynkin diagram, see Reading (2006); Reading and Speyer (2009). It is remarkable that all the Cambrian lattices attached to type \mathbb{A}_n share the same underlying graph, which is nothing else than the graph of vertices and edges of the associahedra. This is also the flip graph of triangulations and the mutation graph of cluster algebras of type \mathbb{A} .

The Tamari lattices have yet another, maybe less well-known, interpretation. They can be seen as a special case of the partial orders defined by Happel and Unger (2005) and Riedtmann and Schofield (1991) on the sets of tilting modules over quivers.

*L'auteur a bénéficié d'une aide de l'Agence Nationale de la Recherche (projet Carma, référence ANR-12-BS01-0017). Il remercie aussi l'Institut Mittag-Leffler pour son accueil chaleureux et ses excellentes conditions de travail.

One can consider the notion of module over the Tamari lattices, defined for example using their incidence algebras. It turns out that there is a very interesting subset of modules, indexed by the set of quadrangulations of a regular polygon. It would be too long to tell here how precisely these modules were stumbled upon, but it involved an anticyclic operad made with all Tamari lattices, and a mysterious map from the K_0 of Tamari lattices to rational functions, see Chapoton (2007). All this seems in fact to take place naturally in the Tamari lattices seen as partial orders on tilting modules.

Anyway, these modules indexed by quadrangulations can be seen as living in the derived categories of modules over the Tamari lattices. A natural question to ask was then: what are the morphisms between these modules? It appeared that they may be described using flips of quadrangulations, under the condition that a flip should not be repeated twice.

At this point, the connection was made with the work of Baryshnikov, which also involved quadrangulations and their flips. It turned out that the match was perfect. This is essentially the story behind the content of the article, which concentrates on the combinatorial side and leaves most of the representation theory aspects aside.

Let us now describe what is done, and then comment on the results.

In section 1, a poset is associated with every quadrangulation in a regular polygon. This is done using the compatibility relation between quadrangulations introduced by Baryshnikov. First, one defines directed graphs (digraphs) by orienting the allowed flips in a specific way. Then it takes some work to show that these digraphs are Hasse diagrams of posets, that we choose to call the Stokes posets. The main tool is an inductive description of the digraphs, relating a quadrangulation and a smaller one with a leaf square removed. One also shows that the associated undirected graphs are connected and regular.

In the articles of Baryshnikov and his collaborators on the subject (Baryshnikov (2001); Baryshnikov et al. (2012); Hickok et al. (2011)), no digraphs were considered, but instead polytopes were associated with quadrangulations. Their graphs of vertices and edges coincides by definition with the associated undirected graphs of our digraphs.

In the rest of the first section, one proposes a conjectural manner to define the same digraphs without having to use the compatibility condition. The proof that this works is sadly missing.

In section 2, one considers examples and properties of the Stokes posets. In particular, one proves that the Tamari lattice is recovered as a special case, for a very regular quadrangulation.

It is then shown that the posets attached to some quadrangulations can be written as Cartesian products of smaller posets of the same kind. This happens under the condition that there is a bridge in the quadrangulation. One then says a few words about the simplicial complexes attached to the quadrangulation, which can be thought of as the dual of Baryshnikov's Stokes polytopes. In particular, enumerative aspects of this simplicial complex can be encoded into an F -triangle.

Next, a conjecture is proposed about what happens when a quadrangulation is twisted along an edge, namely that the undirected graph, simplicial complex and F -triangle should not change.

Then section 3 is a rather sketchy proposal for a vast generalisation of the theory of Stokes posets to other root systems and finite Coxeter groups. This is stated first in the setting of exceptional sequences, and then using factorisation of Coxeter elements into reflections, or equivalently maximal chains in the noncrossing partition lattices.

In the next section 4, one introduces combinatorial configurations fitting inside quadrangulations. They are made of elementary pieces called *serpents*, and therefore named *serpent nests*. These combinatorial objects should be closely related to the posets introduced in the first section. For every quadrangulation Q , there should be as many elements in the poset attached to Q as there are serpent nests in Q . The

serpent nests should be thought of as playing the same role as nonnesting partitions play for classical root systems.

After their definition, the section contains various results about serpent nests, that sometimes are similar to some statements or conjectures about the Stokes posets made in the previous sections.

First, it is shown that serpent nests do factorise in presence of a bridge, just as Stokes posets do. Then an enumerative study is made, by introducing h -vectors and H -triangles that counts serpent nests according to their ranks. A duality is obtained, that implies a symmetry of the h -vector. The H -triangle is conjectured to be related to the F -triangle, by the same formula that work between clusters F -triangles and nonnesting partitions H -triangles for all root systems (see for example Thiel (2014) and the references therein).

It is then shown that, for the same regular quadrangulation already considered in the first section, the serpent nests are in a simple bijection with nonnesting partitions of type \mathbb{A} . Another interesting family of examples is considered, for which the number of serpent nests is given by a Lucas sequence. This gives essentially the (bridge-less) quadrangulations with the smallest possible numbers of serpent nests.

A paragraph is devoted to an interesting strategy to count serpent nests, namely cut quadrangulations in two parts along one edge, compute something related to open serpent nests for every half, and reconstitute the result from that. This seems also to be related to the Dynkin type \mathbb{B} and the usual folding procedure of simply-laced Dynkin diagrams.

Then section 5 deals with an algebraic structure that gathers all the quadrangulations together, namely a commutative algebra endowed with a derivation. Product is disjoint union or union along a bridge and the derivative is given by cutting at inner edges. This is analogue to the parabolic structure on quivers and Dynkin diagrams.

The last section 6 contains elementary enumerative results about various sorts of quadrangulations. The equivalence classes under twisting are also counted, as they should correspond to the distinct undirected graphs of Stokes posets and do describe quadrangulations with the same serpent nests. For some reason related to fans and not studied in this article, quadrangulations that do not contain a cross are also considered.

Let us now summarise the article and give further comments. Taking as input a quadrangulation, one defines a poset and a finite set. For some special quadrangulations, one recovers the Tamari lattices and the nonnesting partitions associated to the Dynkin type \mathbb{A} by the theory of cluster algebras and root systems. It is in fact expected that all the posets will be lattices, and also that other Cambrian lattices can be obtained for other ribbon quadrangulations.

It is therefore very tempting to make an analogy where quadrangulations play the role of quivers, and quadrangulations up to twist that of Dynkin diagrams. Then one gets analogues of clusters, positive clusters, exchange graphs, but only at a combinatorial level. It seems that most of the properties known in the combinatorial side of cluster theory do extend. It would be interesting to see if analogues of cluster variables could make sense.

As explained before, the article Kapranov and Saito (1999) was one source of motivation. The relationship with our results is not as direct as in the case of Baryshnikov, but should exist. It would probably be best seen using the noncrossing trees and exceptional sequences point of view explained in section 3.

Flips in quadrangulations have been considered also in the context of m -analogues of cluster theory for $m = 2$. The 2-Cambrian lattices introduced recently in Stump et al. (2015) are defined by allowing two consecutive flips inside an hexagon, but not three. The precise relationships with this article remain to be

explored.

As a side remark, one can note that very similar combinatorial objects as the ones involved here (quadrangulations and noncrossing trees) have appeared in the study of multiple zeta values, see Ecalle (2011); Dupont (2014). Whether there is a connection or not remains to be settled.

The name of Stokes polytopes was used by Baryshnikov because of the relation of his work with the Stokes phenomenon. In the article Iwaki and Nakanishi (2014), the Stokes phenomenon is also connected to cluster algebras. It is not clear if these connections are the same or at least could fit in a common framework.

Let us finish by saying that there are several aspects missing in the present article, the most important ones being the in-depth study of related fans and polytopes, only sketched in Baryshnikov (2001), and also the closely related question to check that the posets are indeed lattices. These questions are left for future work.

1 Stokes posets

Informally, a *quadrangulation* is a partition of the interior of a regular polygon into quadrilaterals. This is possible if and only if the regular polygon has an even number of vertices.

More formally, one considers a regular polygon with vertices numbered from 0 to $2n + 1$ in the negative (clockwise) direction. A quadrangulation is defined as a collection of pairs of (not consecutive) vertices of this polygon, such that the associated edges do not cross and define only quadrilateral regions.

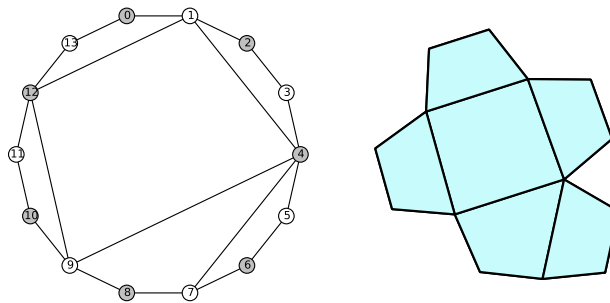


Figure 1: Two views of the same quadrangulation ($n = 6$)

It will sometimes be convenient to depict quadrangulations by deforming the quadrilaterals so that they become more regular and of similar area. The boundary is then no longer a regular polygon, but a simple polygonal closed curve, as in the right of Figure 1.

The number of distinct quadrangulations inside the polygon with $2n + 2$ vertices is also the number of ternary trees with $2n + 1$ leaves and is classically given by

$$\frac{1}{2n + 1} \binom{3n}{n}. \quad (1)$$

The bijection is the obvious planar duality, after a boundary edge has been chosen as root.

Let us fix some terminology. A quadrangulation in the regular polygon with $2n + 2$ vertices is made of $2n + 2$ boundary edges and $n - 1$ inner edges. The n quadrilateral regions will be called *squares*. The word *edge* will generally be used to designate both inner and boundary edges, unless the context make it clear that it must be of one kind only.

1.1 Compatibility

Y. Baryshnikov has introduced a compatibility condition between quadrangulations. Let us describe this precisely.

It will be convenient for us to consider a fixed quadrangulation Q drawn inside a regular polygon with $2n + 2$ vertices, and to define only compatibility with Q . Let us choose an alternating black and white colouring of the vertices of the regular polygon. This choice is arbitrary, and the compatibility condition will not depend on it.

The quadrangulation Q is seen as a (blue = dashed) theatre backdrop, on which one superimposes another (red = plain) regular polygon, slightly rotated in the negative (clockwise) direction by an angle of $\frac{2\pi}{4n+4}$, so that the vertices interlace. The black-and-white colouring of the vertices of the rotated polygon is inherited from the initial colouring.

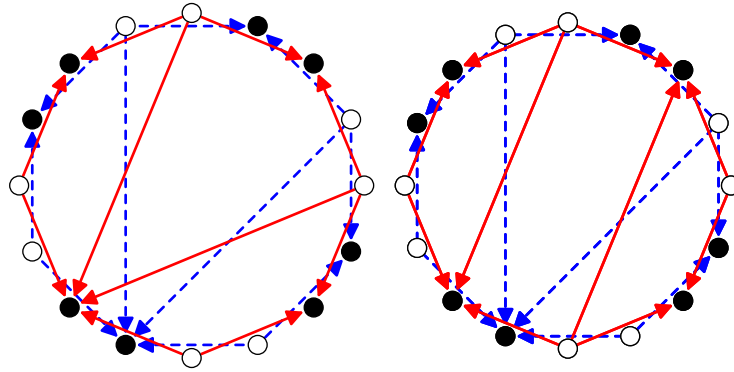


Figure 2: Backdrop Q (blue and dashed) and two Q -compatible Q' and Q'' (red and plain)

By convention, all the edges (red/plain or blue/dashed) are oriented from white vertices \circ to black vertices \bullet .

An edge i in the red polygon is said to be *compatible* with an edge j of Q if

- either they do not cross,
- or they do cross and the pair (red oriented edge i , blue oriented edge j) defines the positive orientation of the plane.

For clarity, let us express the second condition in another way. At their crossing point p , the two oriented edges define two oriented rays starting at p . One requires that the angle from the red ray to the blue ray is between 0 and π .

An edge i in the red polygon is said to be Q -compatible if it is compatible with all edges of Q . This is equivalent to say that i is compatible with all inner edges of Q , because the condition is always satisfied with respect to boundary edges of Q .

Moreover, compatibility with Q is also automatic for the boundary edges of the red polygon.

A quadrangulation Q' is Q -compatible if and only if all its edges are Q -compatible. This is equivalent to require that every inner edge of Q' is compatible with all inner edges of Q .

Note that Q -compatibility is unchanged under switching the choice of the black and white colouring of the vertices, as this changes the orientation of all edges.

It is clear that there is a finite number of Q -compatible quadrangulations. This number depends on Q . By rotating the blue backdrop quadrangulation Q by $\pm \frac{2\pi}{4n+4}$, one defines two quadrangulations that are easily seen to be compatible with Q . Abusing notation, let us call Q the quadrangulation obtained in this way by a small negative rotation, and let $\tau(Q)^{(i)}$ be the other one.

1.2 Flips

Proposition 1.1 *Let Q be a fixed quadrangulation. Let Q' be a Q -compatible quadrangulation. For every inner edge i of Q' , there exists exactly one other inner edge j such that $Q' \setminus \{i\} \cup \{j\}$ is a Q -compatible quadrangulation.*

Proof: This follows from Lemma 1.2 below. □

Lemma 1.2 *Let Q be a fixed quadrangulation. Consider an hexagon H made of six Q -compatible edges. Then there are exactly two Q -compatible inner edges in the interior of H .*

Proof: Let $a_0, b_0, a_1, b_1, a_2, b_2$ be the vertices of the hexagon H in clockwise order, with indices in $\mathbb{Z}/3\mathbb{Z}$. One can assume without restriction that the colours are white for a vertices and black for b vertices. It follows from the Q -compatibility of H that inner edges of Q can only enter H by crossing a inner edge $a_{i+1} \rightarrow b_i$ of H and can only exit by crossing a inner edge $a_i \rightarrow b_i$ of H . If there is no inner edge in Q entering H through $a_{i+1} \rightarrow b_i$ and leaving H through $a_i \rightarrow b_i$ for some i (call it a *long inner edge*), then Q contains an empty hexagon, which is absurd. Therefore there is at least a long inner edge in Q , say from $a_2 \rightarrow b_1$ to $a_0 \rightarrow b_0$. All long inner edges must then enter and exit H by crossing the same edges of H , for otherwise they would cross. Then among the three possible inner edges inside the hexagon H , one can check that $a_0 \rightarrow b_1$ and $a_2 \rightarrow b_0$ are Q -compatible, but that $a_1 \rightarrow b_2$ is not. □

1.3 Flip graphs

Proposition 1.1 implies that one can move between Q -compatible quadrangulations by removing any inner edge and inserting another one. These moves between quadrangulations will be called *flips*.

This defines a graph with vertices the Q -compatible quadrangulations and edges the flips. This is clearly a regular graph where every vertex has arity $n - 1$. Let us denote it by γ_Q and call it the *flip graph* of Q -compatible quadrangulations. Later on, one will show that this graph is connected.

Let us now describe how one can orient the edges of the flip graph γ_Q .

⁽ⁱ⁾ This notation is chosen to remind of the Auslander-Reiten translation.

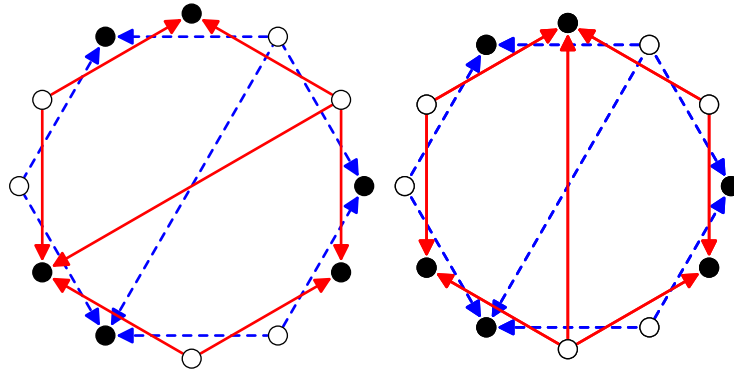


Figure 3: Simple example of flip, oriented from left **(id)** to right **(op)**

Consider a flip relating two quadrangulations Q' and Q'' . Only the inner edge inside the modified hexagon H is changed. The two possible inner edges i' and i'' can be distinguished as follows. Recall from the proof of Lemma 1.2 that there must be at least one *long inner edge* j in Q entering the hexagon and leaving it through the opposite side. The starting vertex of the inner edge of H is adjacent

(id) either to the starting vertex of the long inner edge j ,

(op) or to the ending vertex of the long inner edge j .

Let us orient the flips from the case of adjacent starting vertices **(id)** to the case of far-away starting vertices **(op)**.

This is illustrated in Figure 3, with the **(id)** case on the left and the **(op)** case on the right.

For an oriented flip from Q' to Q'' , one says that it is an *out-flip* for Q' and an *in-flip* for Q'' . For every Q -compatible quadrangulation Q' , the number of in-flips plus the number of out-flips is $n - 1$, as every inner edge corresponds either to an in-flip or to an out-flip.

This allows to define an orientation of the flip graph γ_Q . Let us call this oriented graph Γ_Q .

To study properties of the graph Γ_Q , one will proceed by induction on the *size* n of the quadrangulation Q , which is its number of squares.

In any quadrangulation with $n \geq 2$ squares, one can always find a boundary edge e of Q that bounds a square s with exactly one neighbour square⁽ⁱⁱⁱ⁾. Up to switching black and white, one can assume that the edge e goes from white to black in counterclockwise order along the boundary. Let s be the square bordering the chosen edge e and let f be the unique inner edge of the square s . This situation is illustrated in Figure 4. Let $Q \setminus s$ be the quadrangulation obtained from Q by removing s .

Let Q' be any Q -compatible quadrangulation. The underlying red polygon has two white vertices i and $i + 2$ (modulo $2n + 2$) that lie just before and just after the chosen blue boundary edge e . Recall that the vertices are numbered in clockwise order.

One can associate with Q' a $Q \setminus s$ -compatible quadrangulation $\theta_s(Q')$ as follows. Pictorially, one collapses the boundary section of the red polygon between white vertices i and $i + 2$ to a single white vertex

⁽ⁱⁱⁱ⁾ This is because a quadrangulation is a tree of squares, hence has at least one leaf.

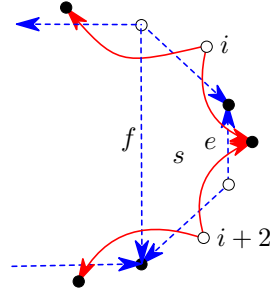


Figure 4: Sketch of configuration near boundary edge e

in a smaller red polygon. Edges of Q' then get identified if they happen to coincide after the collapsing. Note that there can be no inner edges in Q' incident to the black vertex $i+1$, as it would not be compatible with the edge f .

Let us describe now the fibres of the map θ_s . If there are k inner edges in $\theta_s(Q')$ starting at the collapsed vertex, then there must be $k+1$ inner edges in Q' starting either at i or $i+2$. There are $k+2$ choices for the distribution of these $k+1$ starting points between i and $i+2$. Therefore the fibre has $k+2$ elements.

Moreover, these $k+2$ elements are naturally totally ordered by oriented flips, as there is a unique sequence of $k+1$ oriented flips going from the case where all $k+1$ edges start at i to the case where all $k+1$ edges start at $i+2$. This gives a natural total order on every fibre of θ_s .

Oriented flips between Q -compatible quadrangulations are of two kinds.

Proposition 1.3 *Let $Q' \rightarrow Q''$ be a flip of Q -compatible quadrangulations. Then exactly one of the following holds:*

(inFib) $\theta_s(Q') = \theta_s(Q'')$ and the flip goes down one step in this fibre of θ_s ,

(outFib) $\theta_s(Q') \rightarrow \theta_s(Q'')$ is a flip between $Q \setminus s$ -compatible quadrangulations.

Proof: Let H the red hexagon that contains and defines the flip. Let f be the inner blue edge bounding the square s .

Assume first that f is a long inner edge in H . Then H must have two opposite sides made of an edge starting from i and an edge starting from $i+2$. There is at most one such hexagon where an out-flip is possible, and this flip does not change the image by θ_s , but moves down by one step in the total order inside the fibre. This is the situation **(inFib)**.

Otherwise, f is not a long inner edge in H . In this case, the hexagon H remains an hexagon in the reduced quadrangulation $Q \setminus s$, hence the flip induces a flip inside $Q \setminus s$. This is the situation **(outFib)**. \square

We will say that these are *in-fibre flips* and *out-fibre flips*.

Lemma 1.4 *For every sequence of two flips $Q' \xrightarrow{(\text{inFib})} Q'' \xrightarrow{(\text{outFib})} Q'''$, one can find a quadrangulation $Q^\#$ and flips*

$$Q' \xrightarrow{(\text{outFib})} Q^\# \xrightarrow{(\text{inFib})} \dots \xrightarrow{(\text{inFib})} Q''' . \quad (2)$$

Proof: If the two hexagons involved in the two consecutive flips do not overlap, the flips do just commute. One can simply exchange them.

Otherwise everything happens inside an octagon. There are essentially just two cases to consider: there can be either 4 or 5 ways to fill the octagon. In both cases, one can find the expected sequence of flips by inspection. It will involve 1 or 2 final in-fibre flips. \square

Proposition 1.5 *For every pair of quadrangulations Q' to Q'' related by a sequence of flips, one can find a sequence of flips from Q' to Q'' starting with out-fibre flips and ending with in-fibre flips.*

Proof: This follows from the previous lemma 1.4, by repeated application at the first possible place in the sequence. This rewriting preserves the number of out-fibre flips. This process is finite, because at every reduction step, the first out-fibre flip not already packed at the beginning get strictly closer to the beginning of the sequence. \square

Using these tools, one can now proceed to proofs by induction on n .

Proposition 1.6 *Every sequence of oriented flip moves is finite.*

Proof: By induction on the number of squares n . This is obvious for the case $n = 1$, where there is no flip at all.

Otherwise, let us consider an infinite sequence of flips. One can map it to a sequence made either of flips in the smaller quadrangulation $Q \setminus^s$ or a down-step in the total order of a fibre of θ_s . Because these fibres are finite, one can obtain an infinite sequence of flips in $Q \setminus^s$. This is absurd by induction. \square

It follows that there cannot be any oriented cycles in Γ_Q .

Theorem 1.7 *The flip graph Γ_Q is the Hasse diagram of a partial order.*

Proof: The proof is by induction on n . One has to show that the graph is transitively reduced.

This means that for a single flip $Q' \rightarrow Q''$, there should be no longer sequence of flips starting at Q' and ending at Q'' . Here longer means having at least 2 steps.

Assume first that the flip $Q' \rightarrow Q''$ is inside a fibre of θ_s . Then any flip exiting this fibre would never come back to this fibre by induction. So it is only possible to look for the longer sequence inside the fibre. But there cannot be any longer path between two consecutive elements of a total order.

Assume now that the flip $Q' \rightarrow Q''$ induces a flip $\theta_s(Q') \rightarrow \theta_s(Q'')$. Assume moreover that there is some longer sequence of flips from $Q' \rightarrow Q''$. By proposition 1.5, one can suppose that it starts with out-fibre flips and ends with in-fibre flips.

If the longer sequence starts with an out-fibre flip, this must be the flip $Q' \rightarrow Q''$, otherwise one cannot reach Q'' , by induction hypothesis. It must then stop, otherwise it will just stop at some lower point in the fibre. So in fact, it was not longer, which is absurd.

If the longer sequence starts with an in-fibre flip, it can only contain in-fibre flips, and therefore can never reach Q'' which is in a different fibre. This is a contradiction. \square

By convention, the partial order will be decreasing along the edges of the oriented graph Γ_Q .

Proposition 1.8 *The map θ_s is a morphism of posets.*

Proof: This is clear by the description of flips using θ_s in proposition 1.3. \square

Recall the definition of two special Q -compatible quadrangulations Q and $\tau(Q)$ at the end of subsection 1.1.

Proposition 1.9 *The unique Q -compatible quadrangulation with only out-flips is Q . The unique Q -compatible quadrangulation with only in-flips is $\tau(Q)$.*

Proof: Let us first note that Q has only out-flips. This follows from its definition by a small negative rotation. Similarly, $\tau(Q)$ has only in-flips, because it is defined by a small positive rotation.

It remains to show the converse statements. Let us prove that for Q and out-flips, the case of $\tau(Q)$ and in-flips being similar. The proof is illustrated in Figure 5.

Let Q' be a Q -compatible quadrangulation with only out-flips.

Let $s = (i, j, k, \ell)$ be a square of Q that is a leaf in the tree structure of Q , where the boundary edges are $i-j, j-k$ and $k-\ell$. One can assume without restriction that j is a black vertex and k a white vertex.

Let $\hat{j}, \hat{k}, \hat{\ell}$ be the boundary edges of Q' corresponding in the obvious way to the vertices j, k, ℓ . If the unique square of Q' containing \hat{j} and \hat{k} also contains $\hat{\ell}$, then one can remove the boundary edges $i-j, j-k$ and $k-\ell$ from Q and the boundary edges $\hat{j}, \hat{k}, \hat{\ell}$ from Q' and proceed by induction. It then follows that Q' is equal to Q .

Otherwise, there are inner edges of Q' starting at the common vertex of \hat{k} and $\hat{\ell}$. One of them is bounding the square of Q' containing \hat{j} and \hat{k} . Let H be the hexagon containing this edge in its interior. One can check that the inner edge of H is an in-flip for Q' , which is absurd.

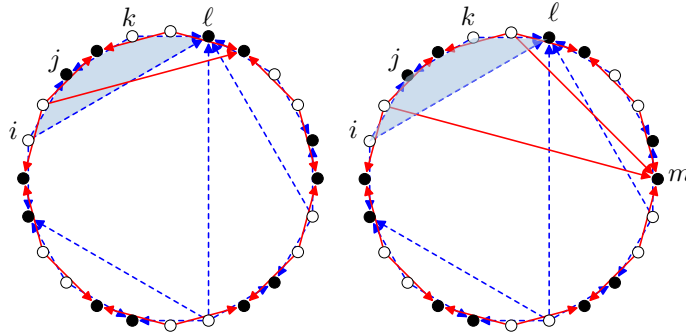


Figure 5: Induction using a leaf of Q

\square

Proposition 1.10 *The oriented flip graph Γ_Q is the Hasse diagram of a connected poset. The undirected flip graph γ_Q is connected.*

Proof: This follows from the existence of a unique minimal element $\tau(Q)$ (and a unique maximal element Q) in the poset, proved in proposition 1.9. \square

From now on, Γ_Q will denote both the poset and its Hasse diagram. Note that these posets are not graded, as one can see already with the pentagon that one can obtain starting with some of the quadrangulations of an octagon.

1.4 Another description ?

Let us now consider another potential way to define the directed graph Γ_Q . This is only conjectural for the moment.

Fix an initial quadrangulation Q . Consider the directed graph D with vertices all quadrangulations, and edges defined by counter-clockwise rotation inside hexagons (where the edge extremities are moved to adjacent vertices). This large directed graph D obviously contains Γ_Q .

One would like to find another way to recover Γ_Q , by keeping only some vertices and edges of D , without having to use the compatibility relation.

The idea is that the set of vertices of Γ_Q should have the following equivalent description. Let us say that a path in D is *repetition-free* if no quadrangulation occurs twice. A quadrangulation Q' is said to be *reachable from Q* if no repetition-free path in D from Q to Q' contains two consecutive flips inside the same hexagon. Then quadrangulations reachable from Q should be exactly the same as elements of Γ_Q .

If this holds, then one can build the set of reachable quadrangulations step by step, using for example the following algorithm.

- Start from the singleton set $S = \{Q\}$.
- Assume that some set S of reachable quadrangulations have been found. Then either
 1. there exists a quadrangulation Q' not in S , obtained from an element of S by a flip, reachable, and such that all quadrangulations on all repetition-free paths in D from Q to Q' are already in S ,
 2. or there is no such Q' and the algorithm stops.

To see what this algorithm does inside Γ_Q , let us fix a linear extension of Γ_Q , starting with Q and ending with τ_Q . Then the algorithm would succeed by adding the elements of Γ_Q according to this linear extension.

This would allow to define the graph Γ_Q without using the compatibility condition introduced by Baryshnikov, at the price of a somewhat greater complexity. This could be useful in more general contexts where no analogue of compatibility is known.

This description of Γ_Q would be summarised by the following sentence: do not rotate twice in the same hexagon.

2 Examples and properties of Stokes posets

2.1 Catalan and other examples

Let us first consider the special case of the quadrangulations C_n with inner edges

$$(0, 3), (0, 5), \dots, (0, 2n - 1) \tag{3}$$

inside the polygon with $2n + 2$ vertices. They have n squares. An example is depicted in Figure 6.

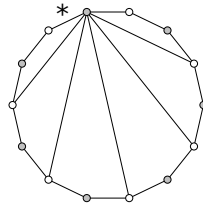


Figure 6: The quadrangulation C_6 of the Catalan family

Let us choose the border edge of the red polygon corresponding to the vertex 0 of the blue polygon as a root. Recall that choosing such a root allows to identify quadrangulations with rooted ternary trees, by planar duality. The root is marked by $*$ in the figure.

Proposition 2.1 *The C_n -compatible quadrangulations are in bijection with rooted ternary trees with no middle branch, and therefore also with rooted binary trees. The flips correspond to the left-to-right rotation moves of rooted binary trees. In-flips correspond to right-branches and out-flips to left branches.*

Proof: This description holds at the starting quadrangulation, which has only left branches. It remains to show that this correspondence is preserved under flips. Details are left to the reader. \square

Therefore, in this case, the number of Q -compatible quadrangulations is the Catalan number

$$\frac{1}{n+1} \binom{2n}{n}, \quad (4)$$

and the poset Γ_Q is the Tamari lattice (see for example Müller-Hoissen et al. (2012) for its usual definition using rotation of binary trees).

More generally, for *ribbon quadrangulations* (defined as quadrangulations where no square has neighbour squares on opposite sides), the posets Γ_Q are expected to match with the Cambrian lattices of type \mathbb{A} .

Figure 7 illustrates the directed flip graph Γ_Q for a quadrangulation with 4 squares, having 12 compatible quadrangulations. This is the simplest example of a Stokes poset that is really new, as the quadrangulation is not a ribbon one.

See Figure 8 for the numbers of Q -compatible quadrangulations for some (connected) quadrangulations with up to 6 squares.

2.2 Factorisation property

Let us now describe a factorisation property, that allows to restrict the attention to a special class of quadrangulations.

A square s in a quadrangulation is called a *bridge* if it has exactly two neighbour squares and these squares are attached to opposite edges of s . A quadrangulation is called *connected* if it does not contain any bridge. See Figure 9 for a quadrangulation with several bridges.

Let Q be a quadrangulation. Assume that s is a bridge in Q . Let Q_1 (resp. Q_2) be the quadrangulation obtained from Q by removing all squares from one side of s (resp. from the other side). Equivalently, Q_1 and Q_2 are obtained by cutting along one edge of s and keeping the part that contains s .

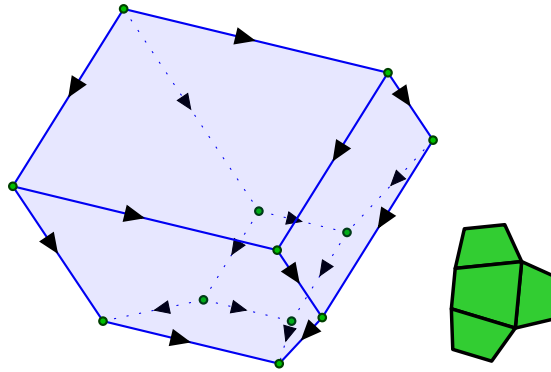


Figure 7: A quadrangulation Q and the flip graph of its 12 Q -compatible quadrangulations

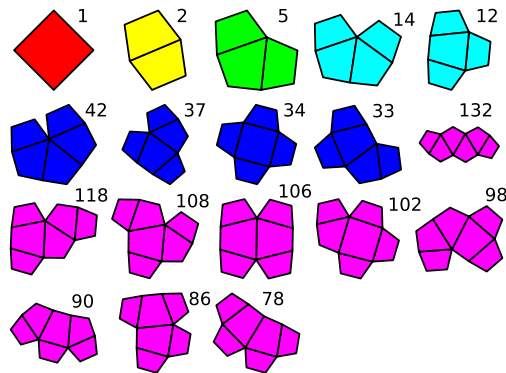


Figure 8: Quadrangulations with up to 6 squares and number of compatible quadrangulations.

Proposition 2.2 *The directed graph Γ_Q is the direct product of the directed graphs Γ_{Q_1} and Γ_{Q_2} .*

Proof: The hypothesis on Q says that in the polygon there are four vertices $i, i + 1$ and $j, j + 1$ that form the given square s of Q . Because the edges $(i, j + 1)$ and $(j, i + 1)$ are oriented in opposite directions, no Q -compatible edge can cross them both. This implies that Q -compatible quadrangulations can be described as pairs made of one Q_1 -compatible quadrangulation and one Q_2 -compatible quadrangulation. One can check that the flips occur in each factor independently. \square

2.3 Simplicial complex and F -triangle

Let Q be a fixed quadrangulation. Let G_Q be the simplicial complex whose simplices are the sets of non-crossing Q -compatible inner edges. Every Q -compatible quadrangulation correspond to a maximal simplex in G_Q .

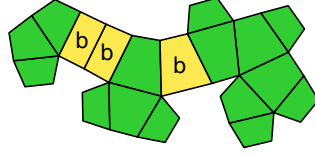


Figure 9: A quadrangulation with 3 bridges (marked with b).

This simplicial complex is pure, because every partial Q -compatible quadrangulation can always be completed into a Q -compatible quadrangulation. To see this, one can observe that a set of Q -compatible edges cuts the polygon into pieces, and each piece behave with respect to compatibility as a smaller polygon, with an underlying quadrangulation inherited from Q .

The simplicial complex G_Q is the flag complex, on the ground set of Q -compatible inner edges, for the compatibility given by being non-crossing.

Let I_Q^- be the set of inner edges in Q (seen as Q -compatible inner edges), and let Φ_Q^+ be the set of all other Q -compatible inner edges.

Let us define the F -triangle as the polynomial in two variables

$$F_Q(x, y) = \sum_{f \in G_Q} x^{\#f \cap \Phi_Q^+} y^{\#f \cap I_Q^-}. \quad (5)$$

When evaluated at (x, x) , this reduces to the usual f -vector of the simplicial complex G_Q .

There is some kind of parabolic structure in Q -compatible quadrangulations, akin to the classical theory of roots systems or Coxeter groups. The role of simple roots is played by the inner edges of Q .

Proposition 2.3 *There holds*

$$y \partial_y F_Q = y \sum_{e \in Q} F_{Q_e^-} F_{Q_e^+}, \quad (6)$$

where the sum runs over inner edges of Q , and Q_e^- , Q_e^+ are the quadrangulations obtained by cutting Q along e .

Proof: The left hand side is counting simplices in the simplicial complex G_Q , with a marked initial inner edge. These are clearly in bijection with the disjoint union, over inner edges of Q , of pairs of simplices in the simplicial complexes $G_{Q_e^+}$ and $G_{Q_e^-}$. \square

Let now s be a bridge in a quadrangulation Q , as defined in subsection 2.2.

Let Q_1 (resp. Q_2) be the quadrangulation obtained from Q by removing all squares from one side of s (resp. from the other side).

Proposition 2.4 *The simplicial complex Γ_Q is the direct product of the simplicial complexes Γ_{Q_1} and Γ_{Q_2} . The F -triangle of Q is the product of the F -triangles of Q_1 and Q_2 .*

Proof: This follows from the same basic properties of the compatibility relation as prop. 2.2. \square

In the article Baryshnikov (2001), all the simplicial complexes G_Q are described as the dual of simple polytopes. In particular, they are all spherical, and we will use here this result. It would certainly be interesting to study the fans and the polytopes that are involved.

The F -triangle has a nice symmetry.

Proposition 2.5 *There holds*

$$F_Q(-1-x, -1-y) = (-1)^n F_Q(x, y). \quad (7)$$

Proof: This is a simple consequence of the fact that G_Q are spherical simplicial complexes, see (Chapoton, 2004/05, Prop. 5) for the proof of this equation in a very similar context. \square

2.4 Twisting along an edge

Let Q be a quadrangulation, and e be an inner edge of Q . Let Q'_e and Q''_e be the two quadrangulations defined by cutting Q along e . Let $\overline{Q''_e}$ be the image of Q''_e by a reflection in the plane (the mirror image of Q''_e).

The *twist* of Q along e (on the Q''_e side) is the quadrangulation defined by gluing back Q'_e and $\overline{Q''_e}$ along their boundary edges corresponding to e .

Twisting twice on the same side gives back Q . Twisting successively on opposite sides give the mirror image of Q .

Let Q and Q' be quadrangulations related by twisting along one inner edge. One expects the following enumerative invariance.

Conjecture 2.6 *The F -triangles of G_Q and $G_{Q'}$ are equal.*

In particular, the number of Q -compatible quadrangulations should be the same as the number of Q' -compatibles ones.

Under a strong additional hypothesis, a more structural property seems to hold.

Conjecture 2.7 *Assume that every square in Q'_e has at most 2 neighbors in Q but is not a bridge. The undirected flip graphs γ_Q and $\gamma_{Q'}$ are isomorphic. The simplicial complexes G_Q and $G_{Q'}$ are isomorphic.*

It should be noted that the oriented flip-graphs are not the same under this hypothesis.

In an earlier version of this article, conjecture 2.7 was proposed without the additional hypothesis. But a minimal counterexample for this original conjecture (with 8 squares) has recently been found by Pilaud and Manneville.

Let us comment briefly on these two conjectures.

First they have been checked by computer for a lot of examples. They are both directly motivated by similar results known to hold for all Cambrian lattices of finite Weyl groups. Let us fix a finite Weyl group W and a Coxeter element c . All Coxeter elements in W can be connected through a sequence of conjugations by simple reflections.

In that case, the directed graph is the Hasse diagram of the Cambrian lattice of (W, c) . The undirected flip graphs does not change when changing the Coxeter element c because of the identification with the flip graph coming from mutations in the cluster algebra of type W .

For the conjecture about F -triangles, in the case of cluster complexes, the invariance under change of Coxeter element c is also known, and follows from the explicit relation with M -triangles or H -triangles, as explained just before section 4.1.

3 Noncrossing trees and exceptional sequences

In this section, which does not contain much details, some generalisations of the flip graphs are proposed, using exceptional sequences on Dynkin diagrams or equivalent objects for Coxeter groups. Beware that the correctness of this proposal depends on the unproven alternative description of Γ_Q in subsection 1.4.

Let us start by a simple combinatorial reformulation of the flip graphs of quadrangulations in terms of other combinatorial objects.

A *noncrossing tree* in the regular polygon with $n + 1$ vertices is a set of edges between vertices of the polygon, with the following properties

- edges do not cross pairwise,
- any two vertices are connected by a sequence of edges,
- there is no loop made of edges.

Note that boundary edges are allowed in the set, and a typical noncrossing tree will contain only some of them.

There is a simple bijection between quadrangulations of the $2n + 2$ -polygon and noncrossing trees in the $n + 1$ -polygon. Assume that vertices of the $2n + 2$ -polygon have been coloured black and white alternating. Then every quadrangle contains a unique black-black diagonal (one of its two diagonals). The collection of these diagonal edges can be seen to form a noncrossing tree. Conversely, by drawing a noncrossing tree using the black vertices of the $2n + 2$ -polygon, one can consider all black-white edges that do not cross edges of the noncrossing tree. This gives back the quadrangulation.

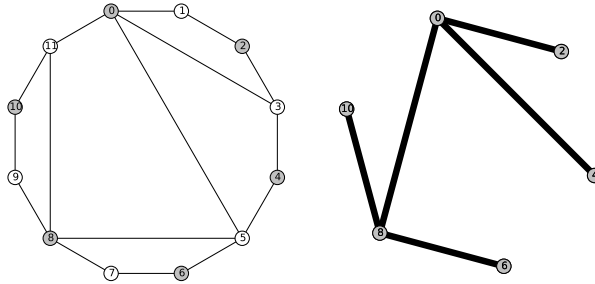


Figure 10: Bijection between quadrangulations and noncrossing trees

Passing through this bijection, one can translate the flips of quadrangulations into a simple operation on noncrossing trees. It is in fact enough to consider what happens for the flip inside an hexagon, as every flip will behave locally the same.


The result is as follows. A flip of noncrossing trees will change just one edge of the noncrossing tree. Assume that there are edges $i - j$ and $k - j$ in the noncrossing tree T , with $i < j < k$ in the clockwise cyclic order, and that there is no other edge incident to j between them inside the ambient polygon. Then the flip consists in replacing $i - j$ by $i - k$.

Flipping twice in the same hexagon of a quadrangulation translates into flipping again along two other sides of the same triangle (i, j, k) of a noncrossing tree.

Using the conjectural alternative description of subsection 1.4, one would therefore be able to restate the digraph Γ_Q in terms of flips of noncrossing trees, not flipping twice consecutively along two sides of the same triangle.

3.1 From noncrossing trees to exceptional sequences

For short, let \mathbb{A}_n denote in this section the equi-oriented quiver of type \mathbb{A}_n .



The noncrossing trees in the $n + 1$ -polygon have been related in Araya (2013) to exceptional sequences (up to shifts and reordering) in the derived category $D \bmod \mathbb{A}_n$.

The correspondence goes as follows. By numbering the vertices of the $n + 1$ -polygon from 0 to n clockwise, every inner edge gets a label (i, j) with $i < j$. Let the inner edge (i, j) correspond to the indecomposable module over \mathbb{A}_n with support $[i + 1, j]$.

Then Araya proved that a noncrossing tree is sent by this map to a collection of indecomposable modules that can be ordered into an exceptional sequence and, conversely all the underlying sets of indecomposable modules coming from exceptional sequences are obtained in this way. Let us call such a collection of indecomposable modules an *exceptional set*⁽ⁱⁱⁱ⁾.

Given the description above of the flips acting on noncrossing trees, one can readily check that they correspond, after the bijection to exceptional sets, to the action of the braid group on exceptional sequences (see for example Crawley-Boevey (1992) for the braid group action). More precisely, given an exceptional set, pick an ordering into an exceptional sequence, act by one generator s_i of the braid group to get another exceptional sequence, and then forget about the ordering. This describe what the flips are.

Then moving twice consecutively in a same triangle gets translated to acting twice by the same s_i at the same place.

This would allow to generalise the definition of the flip graphs to any finite Dynkin diagram, provided that the conjectural description of subsection 1.4 holds.

3.2 From exceptional sequences to chains of noncrossing partitions

In fact, it is even possible to get from here to the more general setting of finite Coxeter groups.

It has been known since the article Ingalls and Thomas (2009) (see also Igusa and Schiffler (2010)) that exceptional sequences in derived categories of representations of Dynkin quivers are closely related to the corresponding lattices of noncrossing partitions.

More precisely, let W be a finite crystallographic Coxeter group, and c be a Coxeter element in W . As is well-known, this data can be translated into a quiver Q with underlying graph the Dynkin diagram.

There is a bijection between maximal chains in the noncrossing partition lattice attached to c , and exceptional sequences of modules over Q . One can label each cover relation in the noncrossing partition lattice by a reflection^(iv). Using this labelling, maximal chains in the noncrossing partition lattice are identified with factorisations of the Coxeter element c as a product of n reflections, where n is the rank.

In this new language, the action of the braid group on exceptional sequences get translated to a similar action of the braid group on maximal chains of the noncrossing partition lattice^(v). When maximal

⁽ⁱⁱⁱ⁾ It could also be called an exceptional collection, but there is no universal agreement on the meaning of that.

^(iv) In the set of all reflections, not only simple ones.

^(v) Sometimes named the Hurwitz action.

chains are considered as factorisations of the Coxeter element into reflections, this is given by a simple conjugation action.

This interpretation would allow to extend the definition of the flip graphs to all finite Coxeter groups, where quadrangulations are replaced by maximal chains in the noncrossing partition lattice, up to reordering.

In these generalisations, some properties are lost. In particular, the flip graphs are no longer regular, as one can see already for some examples in type \mathbb{D}_4 .

4 Serpent nests

In this section, we turn to another aspect of quadrangulations, that should in fact be related to the previous one.

Let us start by some combinatorial definitions.

A *serpent*^(vi) in a quadrangulation Q is a set of two distinct squares s, s' such that the unique path from s to s' in Q is made of a sequence of right-angle turns. This means that following this path, one never enters and leaves a square through opposite sides. Abusing notation, one will also consider that the path is the serpent. By convention, the path starts and ends at the centre of the squares. See the leftmost part of Figure 12 for a visual example.

A *serpent nest* is a set of serpents inside a quadrangulation, satisfying an extra condition and modulo an equivalence relation:

- The **extra condition** is the following: two serpent extremities s_1 and s_2 can share the same square if and only if the serpents do not leave this square by the same side.
- The **equivalence relation** is the following: two sets of serpents are considered equivalent if the pattern of paths inside every square is the same in both.

Let us explain in more details the meaning of this equivalence relation. Inside every square, there can be 8 kinds of path segments: 4 kinds entering by an edge and stopping at the centre (these can appear at most once), and 4 kinds entering by an edge and leaving by an adjacent edge (these can appear any number of times). The pattern of paths just remembers how many paths segments of each kind there are.

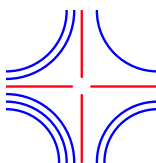


Figure 11: An example of the pattern of paths in a square

Another way to describe this equivalence relation is the following. Consider two serpents that cross the same edge. Cut both of them into two pieces along this edge, and glue them back after swapping them. This gives another set of serpents that also satisfies the extra condition. The equivalence relation is the closure of this kind of move.

^(vi) The word “serpent” is a synonym for “snake”, which is already used for something else in the context of cluster algebras.

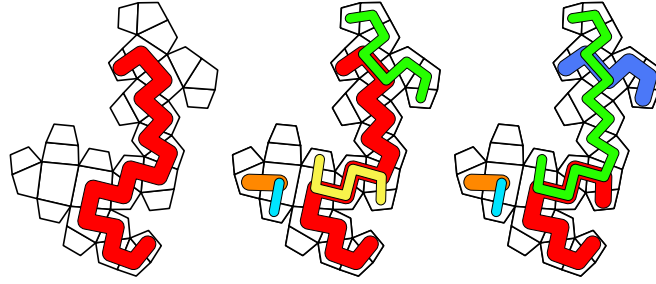


Figure 12: A serpent and two representatives of the same serpent nest of rank 5

The extra condition implies that every square can contain at most 4 serpent extremities. Therefore, there are only a finite number of different serpent nests in each quadrangulation.

Let SN_Q be the set of serpent nests in a quadrangulation Q .

Let us define the *rank* $rk(\sigma)$ of a serpent nest σ as the number of serpents in it. By a well-known principle, this is half the number of serpent extremities.

Let Q be a quadrangulation. Assume that s is a bridge in Q (as defined in subsection 2.2). Let Q_1 (resp. Q_2) be the quadrangulation obtained from Q by removing all squares from one side of s (resp. from the other side).

Theorem 4.1 *There is a bijection between serpent nests in Q and pairs of serpent nests in Q_1 and Q_2 .*

Proof: In fact, no serpent can cross the bridge, by definition, because they must make right-angle turns at every step. Therefore every serpent in Q is either in Q_1 or in Q_2 . The extra condition and equivalence relation are also compatible with this decomposition. \square

Let now Q and Q' be quadrangulations related by twisting along some inner edge e , as defined in 2.4. Every serpent in Q can be mapped to a serpent in Q' by just twisting it. This is clearly a bijection.

Theorem 4.2 *This bijection induces a bijection between serpent nests in Q and in Q' , that preserves the rank.*

Proof: The extra condition in the definition of serpent nests is obviously preserved. Compatibility with the equivalence relation is easy. \square

For a quadrangulation Q , let h_Q be the generating polynomial for the rank of serpent nests:

$$h_Q(x) = \sum_{\sigma \in SN_Q} x^{rk(\sigma)}. \tag{8}$$

Abusing notation, this is called the *h*-vector.

Let us now define a duality on serpent nests.

For this we need the following definition. A *short serpent* is a serpent whose extremities are adjacent squares. We say that a serpent nest σ contains a short serpent at the edge e if one set of serpents in the

equivalence class of σ contains a short serpent at this edge. Equivalently, there are serpent extremities in each square adjacent to e that both leave their respective squares through e .

Let σ be a serpent nest. The dual $\bar{\sigma}$ is another serpent nest, defined as follows.

Consider an edge e between two squares s and s' of Q . If both s and s' contains a serpent extremity leaving through e , remove them both. If neither s nor s' contains a serpent extremity leaving through e , add both. Otherwise (exactly one of s or s' has a serpent extremity leaving through e), do nothing.

This change has to be performed for every inner edge of Q .

This is clearly an involution on serpent nests. This amounts to remove all existing short serpents, and introduce some wherever possible.

Proposition 4.3 *The rank of the dual of a serpent nest σ is $n - 1 - \text{rk}(\sigma)$, where $n - 1$ is the number of inner edges of the quadrangulation.*

Proof: The quadrangulation has n squares and $n - 1$ inner edges.

Consider a serpent nest σ containing k serpents among which ℓ short serpents. There are $k - \ell$ long serpents, and therefore $2k - 2\ell$ inner edges where the duality does nothing (near the extremities of the long serpents). The duality changes something around $n - 1 - 2k + 2\ell$ inner edges. So in the dual $\bar{\sigma}$, there are $n - 1 - 2k + \ell$ short serpents. Therefore $\bar{\sigma}$ contains $n - 1 - k$ serpents. \square

This implies that the h -vector is palindromic:

$$h_Q(x) = x^{n-1} h_Q(1/x). \quad (9)$$

For the four examples of size 5 in Figure 8, one finds the following h -vectors

$$x^4 + 8x^3 + 15x^2 + 8x + 1, \quad (10)$$

$$x^4 + 8x^3 + 16x^2 + 8x + 1, \quad (11)$$

$$x^4 + 9x^3 + 17x^2 + 9x + 1, \quad (12)$$

$$x^4 + 10x^3 + 20x^2 + 10x + 1. \quad (13)$$

From the description of the duality, fixed points are serpent nests where each inner edge is adjacent to exactly one serpent extremity that crosses this edge.

It seems that the number of fixed points is (up to sign) the value of the h -vector at $x = -1$. This could be related to the Charney-Davis conjecture about h -vectors of simplicial spheres, see Reiner and Welker (2005); Charney and Davis (1995); Leung and Reiner (2002).

Let us now define a refinement of the h -vector.

A serpent in a serpent nest σ is called a *simple serpent* if two conditions hold. The first condition is that its extremities are adjacent squares. The second condition is that there is no other serpent crossing the same edge. This is a stronger restriction than just being a short serpent.

The h -triangle of a quadrangulation Q is defined as the following polynomial in two variables:

$$H_Q(x, y) = \sum_{\sigma \in \text{SN}_Q} x^{\text{rk}(\sigma)} y^{s(\sigma)}, \quad (14)$$

where $s(\sigma)$ is the number of simple serpents in the serpent nest σ .

Given an inner edge e in a quadrangulation Q , one can define two quadrangulations by cutting Q along this edge. Let us call them Q'_e and Q''_e . This will be called a *parabolic decomposition*.

The H -triangle has a simple property with respect to parabolic decomposition.

Proposition 4.4 *There holds*

$$y\partial_y H_Q = xy \sum_{e \in Q} H_{Q'_e} H_{Q''_e}. \quad (15)$$

Proof: The left hand side is counting pairs made of a serpent nest in Q with a marked simple serpent. Each term of the right hand side is counting pairs of serpent nests in Q'_e and Q''_e . The equality is given by the bijection that cuts along the unique edge crossed by the simple serpent. \square

There seems to be a close relation with the F -triangle defined in subsection 2.3. Let n be the number of inner edges.

Conjecture 4.5 *There holds the equation*

$$H_Q(x, y) = (x - 1)^n F_Q \left(\frac{1}{x - 1}, \frac{1 + (y - 1)x}{x - 1} \right). \quad (16)$$

This would imply in particular a relation between the f -vector and the h -vector, that is exactly the usual relation defining the h -vector from the f -vector for simplicial complexes.

Given prop. 2.5, this would imply the following symmetry.

Conjecture 4.6 *There holds*

$$H_Q(x, y) = x^n H_Q(1/x, 1 - x + xy). \quad (17)$$

The very same relationship (16) between F -triangles and H -triangles has first been conjectured in Chapoton (2004/05) in the context of Catalan combinatorics of finite Weyl groups and Coxeter groups. In this case, the F -triangle is related to cluster algebras (counting cones in the cluster fan according to the number of spanning positive and negative simple roots) and the H -triangle to nonnesting partitions (counting antichains in the root poset according to cardinality and number of simple roots). A similar formula was conjectured in Chapoton (2006) to relate the H -triangle to a third triangle, called the M -triangle, defined as a generating series for the Möbius numbers of the lattices of noncrossing partitions.

These conjectures have later been generalised in (Armstrong, 2009, Ch. 5) to the Fuss-Catalan combinatorics of Coxeter groups, where one more parameter m appears. The interested reader can find a lot of material on this subject in this reference.

This story is now completely proved, but maybe not fully understood, after the works of several authors, see Athanasiadis (2007); Krattenthaler (2006, 2005/07); Tzanaki (2008) and Thiel (2014).

In fact, all the statements proved about these triangles in the Catalan setting seems to extend to our current context, where the initial data is a quadrangulation instead of a Coxeter type. The M -triangle is missing, as there is no known analogue of the noncrossing partitions here. Note that the symmetry above gets a nice explication in the Catalan world by the self-duality of the lattices of noncrossing partitions.

4.1 Examples

Recall the quadrangulations C_n introduced in subsection 2.1, with inner edges

$$(0, 3), (0, 5), \dots, (0, 2n - 1) \quad (18)$$

inside the polygon with $2n + 2$ vertices, see Figure 6.

In this case, every pair of squares define a serpent and every serpent nest is uniquely determined by the set of serpent extremities. In one square, there can be at most two serpent extremities.

Recall that a nonnesting partition of the set $\{1, \dots, n-1\}$ is a set of segments (i, j) with $1 \leq i \leq j \leq n$, such that no segment contains another one.

By looking at serpent extremities as opening or closing parentheses, one can find a simple bijection between serpent nests and nonnesting partitions of $\{1, \dots, n-1\}$, that maps the number of serpents to the number of segments.

It follows that the number of serpent nests inside C_n is the Catalan number and the h -vector of C_n is given by the Narayana numbers.

Another nice family of examples is given by the quadrangulations of the general shape displayed in Figure 13, namely a linear string of $n + 1$ squares, with one square attached below each of them excepted the two extremes. Let us denote by L_n this quadrangulation, which has $2n$ squares. We call them the Lucas family, for reasons explained below. Let us also denote by K_n the quadrangulation obtained from L_n by removing the rightmost square in the linear string of $n + 1$ squares.

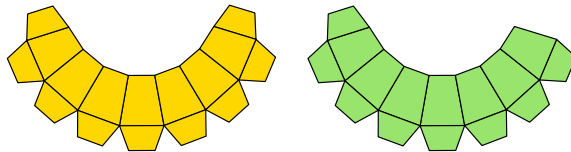


Figure 13: The quadrangulation L_8 of the Lucas family and the quadrangulation K_8

Let us denote by ℓ_n and k_n the h -vectors h_{L_n} and h_{K_n} . Recall that they are polynomials in x counting the serpent nests according to their rank.

Proposition 4.7 *The sequence ℓ_n satisfies the recurrence*

$$\ell_{n+2} = (1 + 4x + x^2)\ell_{n+1} + x(1 + x + x^2)\ell_n, \quad (19)$$

with initial conditions $\ell_0 = 0$ and $\ell_1 = 1 + x$.

Proof: One can obtain the following coupled recursions for ℓ_n and k_n by a combinatorial reasoning on serpent nests:

$$\ell_n = (1 + x)k_n + x\ell_{n-1} \quad \text{and} \quad k_n = (1 + x)\ell_{n-1} + xk_{n-1} + x(1 + x)\ell_{n-2}. \quad (20)$$

By elimination, one finds the required formula. \square With more care, a similar recursion can be found for the H -triangles.

It follows that the ℓ_n form a Lucas sequence, and therefore that there exist polynomials $(Z_n)_{n \geq 1}$ such that

$$\ell_n = \prod_{d|n} Z_d, \quad (21)$$

for all $n \geq 1$. Apparently, the Z_n have positive coefficients. There is no combinatorial explanation of this factorisation so far. One may speculate that this could be explained by some generalisations of noncrossing partitions lattices.

The first few numbers of serpent nests in L_n are

$$2, 12, 78, 504, 3258, 21060, 136134, 879984, 5688306, \dots \quad (22)$$

(twice sequence A090018 of `oeis.org`) and the first few values of Z_n at $x = 1$ are

$$2, 6, 39, 42, 1629, 45, 68067, 1746, 72927, 1881, 118841823, 1737, 4965759189, \dots \quad (23)$$

4.2 Counting serpent nests

A natural question to ask is, given a quadrangulation, how many serpent nests does it afford? Even better, one would like to know the h -vector that counts serpent nests according to their rank.

This does not seem to be easy. There is nevertheless an idea that allows to compute these numbers or polynomials for some large quadrangulations. It consists of

1. cutting a quadrangulation along one edge,
2. counting what could be called open serpent nests in both half-quadrangulations
3. and gluing the results back.

Let us call *open quadrangulation* a quadrangulation together with a distinguished boundary edge, that will be considered to be open. An *open serpent nest* inside an open quadrangulation is defined in the same way, except that now some of the serpents can end on the open edge. They are considered up to the same equivalence relation as serpent nests. Clearly, there is a finite number of such configurations.

The rank $\text{rk}(\sigma)$ of an open serpent nest σ is the number of (closed) serpents. Let $\text{op}(\sigma)$ be the number of open serpents in σ , namely those that end on the open edge.

Given an open quadrangulation Q , define its *open h -vector* as

$$h_Q(x, t) = \sum_{\sigma \in \text{SN}_Q} x^{\text{rk}(\sigma)} t^{\text{op}(\sigma)}. \quad (24)$$

One will glue back two open h -vectors to get an h -vector. The gluing is defined by the following rule. Let A and B be two polynomials in x and t . Define a polynomial $A \# B$ in x by

$$A \# B = \sum_{k \geq 0} A_k B_k x^k, \quad (25)$$

where A_k and B_k are the coefficients of t^k in A and B (when considered as polynomials in t with coefficients in $\mathbb{Z}[x]$).

Let Q be a quadrangulation, e be an inner edge of Q . Let Q'_e and Q''_e be the open quadrangulations obtained by cutting Q along the edge e .

Proposition 4.8 *The h -vector of Q is given by $h_Q = h_{Q'_e} \# h_{Q''_e}$.*

Proof: Cutting along e gives a bijection between serpent nests in Q and pairs of open serpent nests in Q'_e and Q''_e with the same number of open serpents. This implies the formula. \square

To give just a very simple example, consider the quadrangulation of an hexagon. Its h -vector is $1 + x$. On the other hand, for the open quadrangulation of a square, one find that the open h -vector is $1 + t$. One can recover $1 + x$ as $(1 + t) \# (1 + t)$.

This method allows for example to compute the number of serpents nests for quadrangulations made by gluing a large open Catalan-type quadrangulation (or a large open Lucas-type one) to any small open quadrangulation.

5 Parabolic algebras of quadrangulations and cross-trees

Consider the free commutative algebra \mathcal{F} generated by all quadrangulations (considered up to rotation). Let \mathcal{A} be its quotient by the relations

$$Q = Q_1 Q_2 \tag{26}$$

when there is a bridge s in a quadrangulation Q defining two quadrangulations Q_1 and Q_2 as in subsection 2.2.

Then \mathcal{A} is a graded commutative algebra, where the degree is the number of inner edges.

On this algebra, let us define an operator $\partial : \mathcal{A} \rightarrow \mathcal{A}$. It is given on a quadrangulation Q by

$$\partial Q = \sum_{e \in Q} Q'_e Q''_e \tag{27}$$

where the right hand side is the product of the two pieces obtained by cutting along the edge e . This extends uniquely as a derivation of the free commutative algebra \mathcal{F} , which then descends to the quotient algebra \mathcal{A} .

The operator ∂ is a derivation with respect to the product, and homogeneous of degree -1 .

Let Δ be the operator $\exp(\partial)$. This is a well-defined automorphism of \mathcal{A} . It does not preserve the grading, but does preserve an increasing filtration by the number of inner edges.

Recall the F -triangle of a quadrangulation was defined in 2.3 as a polynomial in x, y .

Proposition 5.1 *The F -triangle is a morphism of commutative algebra from \mathcal{A} to $\mathbb{Z}[x, y]$. It sends the derivation ∂ to ∂_y and Δ to the substitution operator $y \mapsto y + 1$.*

Proof: The fact that it is a morphism follows from Prop. 2.4. The derivation part follows from Prop. 2.3. The statement about Δ is a consequence of the statement about ∂ . \square

Let us call a *cross-tree* the equivalence class of a quadrangulation under the moves of twisting along an edge as defined in subsection 2.4.

Cross trees can also be considered as abstract sets of squares, glued together along their sides so as to form a tree-like shape, but where the two possible gluings along an edge are not distinguished.

The notion of bridge introduced in subsection 2.2 still make sense for cross-trees, as well as the associated notion of connectedness.

By the theorem 4.2, the set of serpent nests inside a quadrangulation really only depends on the cross-tree.

One can define an algebra similar to \mathcal{A} using cross-trees instead of quadrangulations. This is a quotient algebra of \mathcal{A} .

If conjecture 2.6 holds, the F -triangle only depends on the cross tree, and therefore is defined on this quotient algebra.

6 Enumerative aspects of quadrangulations

We shall compute here the number of connected quadrangulations up to rotation and the number of connected cross-trees, up to isomorphism.

This is done using the classical language of species, see Bergeron et al. (1998) for the basics of this fundamental theory.

Let X be the species of singletons. Let E_2 be the species of unordered pairs. Let C_4 be the species of oriented cycles of length 4.

6.1 Connected quadrangulations up to rotation

Let T be the species of all connected quadrangulations, T^\square the species of those pointed in a square, T^i the species of those pointed along an inner edge, and $T^{\square,i}$ the species of those pointed along a pair (square, adjacent inner edge).

Let T^b be the species of quadrangulations that can be obtained from a connected quadrangulation by cutting along an inner edge, keeping one of the two halves and marking the cut edge. The species T^b is an auxiliary one, in which it is not allowed that the square next to the marked edge has exactly one other neighbour opposite to the edge, but where it is allowed that this square has only two opposite neighbours.

One has the following combinatorial relations:

$$\begin{aligned} T^\square + T^i &= T + T^{\square,i}, \\ T^i &= E_2(T^b), \\ T^{\square,i} &= (T^b)^2, \\ T^\square &= X(1 + T^b + (T^b)^2 + (T^b)^3 + C_4(T^b)), \\ T^b &= X(1 + 2T^b + 3(T^b)^2 + (T^b)^3). \end{aligned}$$

The first equation is an instance of the dissymmetry equation for trees, see (Bergeron et al., 1998, Chap. 4) for the general principle behind it. Other equations are just obvious recursive descriptions. The last equation determines T^b and the other ones in turn determines the other species.

From these equations, one can compute the first few terms of the associated generating series:

$$\begin{aligned} T^b &: 1, 2, 7, 27, 114, 507, 2342, \dots \\ T &: 1, 1, 1, 4, 11, 42, 155, 659, 2810, \dots \\ T^i &: 0, 1, 2, 10, 41, 196, 924, \dots \\ T^\square &: 1, 1, 3, 12, 52, 231, 1079, \dots \\ T^{\square,i} &: 0, 1, 4, 18, 82, 385, 1848, \dots \end{aligned}$$

It is also interesting to consider the sub-species T^B of T^b where the square near the marked boundary edge does not have exactly two opposite neighbours. In some sense, this condition is a stronger form of connectedness. By removing one term in the equation defining T^b , one gets that

$$T^B = X(1 + 2T^b + 2(T^b)^2 + (T^b)^3), \quad (28)$$

which gives the first few terms

$$T^B : 1, 2, 6, 23, 96, 425, 1957, 9277, \dots$$

6.2 Connected cross-trees up to isomorphism

Let us compute here the number of connected cross-trees.

Let \mathbf{T} be the species of all connected cross-trees, \mathbf{T}^\square the species of those pointed in a square, \mathbf{T}^i the species of those pointed along an inner edge, and $\mathbf{T}^{\square,i}$ the species of those pointed along a pair (square, adjacent inner edge).

Let \mathbf{T}^b be the species of cross-trees that can be obtained from a connected cross-tree by cutting along an inner edge, keeping one of the two halves and marking the cut edge. The species \mathbf{T}^b is an auxiliary one, in which it is not allowed that the square next to the marked edge has exactly one other neighbour opposite to the edge, but where it is allowed that this square has only two opposite neighbours.

One has the following combinatorial relations:

$$\begin{aligned} \mathbf{T}^\square + \mathbf{T}^i &= \mathbf{T} + \mathbf{T}^{\square,i}, \\ \mathbf{T}^i &= E_2(\mathbf{T}^b), \\ \mathbf{T}^{\square,i} &= (\mathbf{T}^b)^2, \\ \mathbf{T}^\square &= X(1 + \mathbf{T}^b + E_2(\mathbf{T}^b) + \mathbf{T}^b E_2(\mathbf{T}^b) + E_2(E_2(\mathbf{T}^b))), \\ \mathbf{T}^b &= X(1 + \mathbf{T}^b + (\mathbf{T}^b)^2 + E_2(\mathbf{T}^b) + \mathbf{T}^b E_2(\mathbf{T}^b)). \end{aligned}$$

The first equation is again an instance of the dissymmetry equation for trees. Other equations are just obvious recursive descriptions. As before, the last equation determines \mathbf{T}^b and the other ones in turn determines the other species.

From these equations, one can compute the first few terms of the associated generating series for the number of isomorphism classes:

$$\begin{aligned} \mathbf{T}^b &: 1, 1, 3, 7, 20, 58, 178, 557, \dots \\ \mathbf{T} &: 1, 1, 1, 2, 4, 9, 21, 56, 153, \dots \\ \mathbf{T}^i &: 0, 1, 1, 4, 10, 33, 99, 324, \dots \\ \mathbf{T}^\square &: 1, 1, 2, 5, 14, 39, 120, \dots \\ \mathbf{T}^{\square,i} &: 0, 1, 2, 7, 20, 63, 198, \dots \end{aligned}$$

Just as for quadrangulations, it is also interesting to consider the sub-species \mathbf{T}^B of \mathbf{T}^b where the square near the marked boundary edge does not have exactly two opposite neighbours. By removing one term in the equation defining \mathbf{T}^b , one gets that

$$\mathbf{T}^B = X(1 + \mathbf{T}^b + (\mathbf{T}^b)^2 + \mathbf{T}^b E_2(\mathbf{T}^b)), \quad (29)$$

which gives the first few terms

$$\mathbf{T}^B : 1, 1, 2, 6, 16, 48, 145, 458, \dots$$

The first few terms of sequence \mathbf{T} are illustrated in Figure 8.

6.3 Connected quadrangulations with no cross

For reasons related to fans (that are not considered in this article), it is also interesting to count quadrangulations and cross-trees avoiding the cross (which means that no square has four neighbours). We will abuse notation by keeping the same names for these new species.

Let us start with quadrangulations. Using similar notations for various species of connected quadrangulations avoiding the cross, one has the following combinatorial relations:

$$\begin{aligned} T^\square + T^i &= T + T^{\square,i}, \\ T^i &= E_2(T^b), \\ T^{\square,i} &= (T^b)^2, \\ T^\square &= X(1 + T^b + (T^b)^2 + (T^b)^3), \\ T^b &= X(1 + 2T^b + 3(T^b)^2). \end{aligned}$$

The first equation is an instance of the dissymmetry equation. Other equations are just obvious recursive descriptions. The last equation determines T^b and the other ones in turn determines the other species.

From these equations, one can compute the first few terms of the associated generating series:

$$\begin{aligned} T^b &: 1, 2, 7, 26, 106, 452, 1999, \dots \\ T &: 1, 1, 1, 4, 10, 40, 141, \dots \\ T^i &: 0, 1, 2, 10, 40, 186, 846, \dots \\ T^\square &: 1, 1, 3, 12, 50, 219, 987, \dots \\ T^{\square,i} &: 0, 1, 4, 18, 80, 365, 1692, \dots \end{aligned}$$

6.4 Connected cross-trees with no cross

Using the same notations as before, one has the following combinatorial relations:

$$\begin{aligned} \mathbf{T}^\square + \mathbf{T}^i &= \mathbf{T} + \mathbf{T}^{\square,i}, \\ \mathbf{T}^i &= E_2(\mathbf{T}^b), \\ \mathbf{T}^{\square,i} &= (\mathbf{T}^b)^2, \\ \mathbf{T}^\square &= X(1 + \mathbf{T}^b + E_2(\mathbf{T}^b) + \mathbf{T}^b E_2(\mathbf{T}^b)), \\ \mathbf{T}^b &= X(1 + \mathbf{T}^b + (\mathbf{T}^b)^2 + E_2(\mathbf{T}^b)). \end{aligned}$$

The first equation is an instance of the dissymmetry equation for trees. Other equations are just obvious recursive descriptions. The last equation determines \mathbf{T}^b and the other ones in turn determines the other species.

From these equations, one can compute the first few terms of the associated generating series for the number of isomorphism classes:

$$\begin{aligned}\mathbf{T}^b &: 1, 1, 3, 6, 17, 44, 128, 365, 1091, \dots \\ \mathbf{T} &: 1, 1, 1, 2, 3, 8, 16, 42, 102, \dots \\ \mathbf{T}^i &: 0, 1, 1, 4, 9, 29, 79, 244, 727, \dots \\ \mathbf{T}^\square &: 1, 1, 2, 5, 12, 34, 95, 280, 829, \dots \\ \mathbf{T}^{\square,i} &: 0, 1, 2, 7, 18, 55, 158, 482, 1454, \dots\end{aligned}$$

The reader can check the first few terms of sequence \mathbf{T} in Figure 8.

References

- T. Araya. Exceptional sequences over path algebras of type A_n and non-crossing spanning trees. *Algebr. Represent. Theory*, 16(1):239–250, 2013. ISSN 1386-923X. doi: 10.1007/s10468-011-9304-4.
- D. Armstrong. Generalized noncrossing partitions and combinatorics of Coxeter groups. *Mem. Amer. Math. Soc.*, 202(949):x+159, 2009. ISSN 0065-9266. doi: 10.1090/S0065-9266-09-00565-1.
- C. A. Athanasiadis. On some enumerative aspects of generalized associahedra. *European J. Combin.*, 28(4):1208–1215, 2007. ISSN 0195-6698. doi: 10.1016/j.ejc.2006.02.002.
- Y. Baryshnikov. On Stokes sets. In *New developments in singularity theory (Cambridge, 2000)*, volume 21 of *NATO Sci. Ser. II Math. Phys. Chem.*, pages 65–86. Kluwer Acad. Publ., Dordrecht, 2001.
- Y. Baryshnikov, L. Hickok, N. Orlov, and S. Son. Stokes polyhedra for x-shaped polyminos. *DMTCS proc.*, 2012.
- F. Bergeron, G. Labelle, and P. Leroux. *Combinatorial species and tree-like structures*, volume 67 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1998. ISBN 0-521-57323-8. Translated from the 1994 French original by Margaret Readdy, With a foreword by Gian-Carlo Rota.
- F. Chapoton. Enumerative properties of generalized associahedra. *Sém. Lothar. Combin.*, 51:Art. B51b, 16, 2004/05. ISSN 1286-4889.
- F. Chapoton. Sur le nombre de réflexions pleines dans les groupes de Coxeter finis. *Bull. Belg. Math. Soc. Simon Stevin*, 13(4):585–596, 2006. ISSN 1370-1444.
- F. Chapoton. The anticyclic operad of moulds. *Int. Math. Res. Not. IMRN*, (20):Art. ID rnm078, 36, 2007. ISSN 1073-7928. doi: 10.1093/imrn/rnm078. URL <http://dx.doi.org/doi:10.1093/imrn/rnm078>.
- R. Charney and M. Davis. The Euler characteristic of a nonpositively curved, piecewise Euclidean manifold. *Pacific J. Math.*, 171(1):117–137, 1995. ISSN 0030-8730.

- W. Crawley-Boevey. Exceptional sequences of representations of quivers. In *Proceedings of the Sixth International Conference on Representations of Algebras (Ottawa, ON, 1992)*, volume 14 of *Carleton-Ottawa Math. Lecture Note Ser.*, page 7, Ottawa, ON, 1992. Carleton Univ.
- C. Dupont. The combinatorial Hopf algebra of motivic dissection polylogarithms. *Adv. Math.*, 264: 646–699, 2014. ISSN 0001-8708. doi: 10.1016/j.aim.2014.07.027. URL <http://dx.doi.org/doi/10.1016/j.aim.2014.07.027>.
- J. Ecalle. The flexion structure and dimorphy: flexion units, singulators, generators, and the enumeration of multizeta irreducibles. In *Asymptotics in dynamics, geometry and PDEs; generalized Borel summation. Vol. II*, volume 12 of *CRM Series*, pages 27–211. Ed. Norm., Pisa, 2011. With computational assistance from S. Carr.
- S. Fomin and A. Zelevinsky. Y -systems and generalized associahedra. *Ann. of Math. (2)*, 158(3):977–1018, 2003. ISSN 0003-486X. doi: 10.4007/annals.2003.158.977. URL <http://dx.doi.org/doi/10.4007/annals.2003.158.977>.
- D. Happel and L. Unger. On a partial order of tilting modules. *Algebr. Represent. Theory*, 8(2):147–156, 2005. ISSN 1386-923X. doi: 10.1007/s10468-005-3595-2.
- L. Hickok, N. Orlow, and S. Son. Investigating the Stokes Sets. Research Experiences for Graduate Students, University of Illinois, 09 2011.
- K. Igusa and R. Schiffler. Exceptional sequences and clusters. *J. Algebra*, 323(8):2183–2202, 2010. ISSN 0021-8693. doi: 10.1016/j.jalgebra.2010.02.003.
- C. Ingalls and H. Thomas. Noncrossing partitions and representations of quivers. *Compos. Math.*, 145(6):1533–1562, 2009. ISSN 0010-437X. doi: 10.1112/S0010437X09004023.
- K. Iwaki and T. Nakanishi. Exact WKB analysis and cluster algebras. *J. Phys. A*, 47(47):474009, 98, 2014. ISSN 1751-8113. doi: 10.1088/1751-8113/47/47/474009.
- M. M. Kapranov and M. Saito. Hidden Stasheff polytopes in algebraic K -theory and in the space of Morse functions. In *Higher homotopy structures in topology and mathematical physics (Poughkeepsie, NY, 1996)*, volume 227 of *Contemp. Math.*, pages 191–225. Amer. Math. Soc., Providence, RI, 1999.
- C. Krattenthaler. The M -triangle of generalised non-crossing partitions for the types E_7 and E_8 . *Sém. Lothar. Combin.*, 54:Art. B541, 34, 2005/07. ISSN 1286-4889.
- C. Krattenthaler. The F -triangle of the generalised cluster complex. In *Topics in discrete mathematics*, volume 26 of *Algorithms Combin.*, pages 93–126. Springer, Berlin, 2006. doi: 10.1007/3-540-33700-8_6.
- N. C. Leung and V. Reiner. The signature of a toric variety. *Duke Math. J.*, 111(2):253–286, 2002. ISSN 0012-7094. doi: 10.1215/S0012-7094-02-11123-5.

- F. Müller-Hoissen, J. M. Pallo, and J. Stasheff, editors. *Associahedra, Tamari lattices and related structures*, volume 299 of *Progress in Mathematical Physics*. Birkhäuser/Springer, Basel, 2012. ISBN 978-3-0348-0404-2; 978-3-0348-0405-9. doi: 10.1007/978-3-0348-0405-9. URL [http://dx.doi.org,docelec.univ-lyon1.fr/10.1007/978-3-0348-0405-9](http://dx.doi.org/doi.org,docelec.univ-lyon1.fr/10.1007/978-3-0348-0405-9). Tamari memorial Festschrift.
- N. Reading. Cambrian lattices. *Adv. Math.*, 205(2):313–353, 2006. ISSN 0001-8708. doi: 10.1016/j.aim.2005.07.010. URL <http://dx.doi.org/10.1016/j.aim.2005.07.010>.
- N. Reading and D. E. Speyer. Cambrian fans. *J. Eur. Math. Soc. (JEMS)*, 11(2):407–447, 2009. ISSN 1435-9855. doi: 10.4171/JEMS/155. URL <http://dx.doi.org/10.4171/JEMS/155>.
- V. Reiner and V. Welker. On the Charney-Davis and Neggers-Stanley conjectures. *J. Combin. Theory Ser. A*, 109(2):247–280, 2005. ISSN 0097-3165. doi: 10.1016/j.jcta.2004.09.003.
- C. Riedtmann and A. Schofield. On a simplicial complex associated with tilting modules. *Comment. Math. Helv.*, 66(1):70–78, 1991. ISSN 0010-2571. doi: 10.1007/BF02566636.
- J. D. Stasheff. Homotopy associativity of H -spaces. I, II. *Trans. Amer. Math. Soc.* 108 (1963), 275-292; *ibid.*, 108:293–312, 1963. ISSN 0002-9947.
- C. Stump, H. Thomas, and N. Williams. Cataland: Why the Fuss? *ArXiv e-prints*, Mar. 2015.
- M. Thiel. On the H -triangle of generalised nonnesting partitions. *European J. Combin.*, 39:244–255, 2014. ISSN 0195-6698. doi: 10.1016/j.ejc.2014.01.010.
- E. Tzanaki. Faces of generalized cluster complexes and noncrossing partitions. *SIAM J. Discrete Math.*, 22(1):15–30, 2008. ISSN 0895-4801. doi: 10.1137/060665683.