Improved Mixing for the Convex Polygon Triangulation Flip Walk

David Eppstein ✉

Department of Computer Science, University of California, Irvine, CA, USA

Daniel Frishberg ✉

Department of Computer Science, University of California, Irvine, CA, USA

Abstract

We prove that the well-studied *triangulation flip walk* on a convex point set mixes in time $O(n^3 \log^3 n)$, the first progress since McShine and Tetali's $O(n^5 \log n)$ bound in 1997. In the process we give lower and upper bounds of respectively $\Omega(1/(\sqrt{n} \log n))$ and $O(1/\sqrt{n})$ – asymptotically tight up to an $O(\log n)$ factor – for the *expansion* of the *associahedron* graph K_n. The upper bound recovers Molloy, Reed, and Steiger's $\Omega(n^{3/2})$ bound on the mixing time of the walk. To obtain these results, we introduce a framework consisting of a set of sufficient conditions under which a given Markov chain mixes rapidly. This framework is a purely combinatorial analogue that in some circumstances gives better results than the *projection-restriction* technique of Jerrum, Son, Tetali, and Vigoda. In particular, in addition to the result for triangulations, we show quasipolynomial mixing for the *k-angulation* flip walk on a convex point set, for fixed $k \geq 4$.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases associahedron, mixing time, mcmc, Markov chains, triangulations, quadrangulations, k-angulations, multicommodity flow, projection-restriction

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.56

Category Track A: Algorithms, Complexity and Games

Related Version This paper includes results from two of our arXiv preprints:
Previous Version: `https://arxiv.org/abs/2207.09972v1`
Previous Version: `https://arxiv.org/abs/2111.03898`
Full Version: `https://arxiv.org/abs/2207.09972`

1 Introduction and background

The study of *mixing times* – the art and science of proving upper and lower bounds on the efficiency of Markov chain Monte Carlo sampling methods – is a well-established area of research, of interest for combinatorial sampling problems, spin systems in statistical physics, probability, and the study of subset systems. Work in this area brings together techniques from spectral graph theory, combinatorics, and probability, and dates back decades; for a comprehensive survey of classic methods, results, and open questions see the canonical text by Levin, Wilmer, and Peres [27]. Recent breakthroughs [1, 2, 3, 8, 9, 10, 24, 26] – incorporating techniques from the theory of abstract simplicial complexes – have led to a recent slew of results for the mixing times of combinatorial chains for sampling independent sets, matchings, Ising model configurations, and a number of other structures in graphs, injecting renewed energy into an already active area.

We focus on a class of *geometric* sampling problems that has received considerable attention from the counting and sampling [4, 22] and mixing time [29, 31, 35, 6] research communities over the last few decades, but for which tight bounds have been elusive: sampling *triangulations*. A triangulation is a maximal set of non-crossing edges connecting pairs of points (see Figure 1) in a given n-point set. Every pair of triangles sharing an edge forms a quadrilateral. A triangulation *flip* consists of removing such an edge, and replacing it with the only other possible diagonal within the same quadrilateral. Flips give a natural Markov chain (the *flip walk*): one selects a uniformly random diagonal from a given triangulation and (if possible) flips the diagonal.

McShine and Tetali gave a classic result in a 1997 paper [29], showing that in the special case of a convex two-dimensional point set (a convex n-gon), the flip walk *mixes* (converges to approximately uniform) in time $O(n^5 \log n)$, improving on the best-known prior (and first polynomial) upper bound, $O(n^{25})$, by Molloy, Reed, and Steiger [31]. McShine and Tetali applied a Markov chain comparison technique due to Diaconis and Saloff-Coste [12] and to Randall and Tetali [32] to obtain their bound, using a bijection between triangulations and a structure known as *Dyck paths.* They noted that they could not improve on this bound using this bijection. Furthermore, they believed that an earlier *lower* bound of $\Omega(n^{3/2})$, also by Molloy, Reed, and Steiger [31], should be tight. We show the following result (see Section 3 for the precise definition of mixing time):

▶ **Theorem 1.** *The triangulation flip walk on the convex $n+2$-point set mixes in time $O(n^3 \log^3 n)$.*

Prior to the present paper, no progress had been made either on upper or lower bounds for this chain in 25 years – even as new polynomial upper bounds and exponential lower bounds were given for other geometric chains, from lattice point set triangulations [35, 6] to quadrangulations of planar maps [7], and despite many breakthroughs using the newer techniques for other problems.

In addition to this specific result, we give a general decomposition theorem – which we will state as Theorem 13 once we have built up enough preliminaries, for bounding mixing times by recursively decomposing the state space of a Markov chain. This theorem is a purely combinatorial alternative to the spectral result of Jerrum, Son, Tetali, and Vigoda [21].

1.1 Decomposition framework

To prove our result, we develop a general decomposition framework that applies to a broad class of Markov chains, as an alternative to prior work by Jerrum, Son, Tetali, and Vigoda [21] that used spectral methods. We obtain our new mixing result for triangulations, then generalize our technique to obtain the first nontrivial mixing result for k-angulations. In a companion paper [15] we further generalize this work to obtain the first rapid mixing bounds for Markov chains for sampling independent sets, dominating sets, and *b-edge covers* (generalizing edge covers) in graphs of bounded treewidth, and for maximal independent sets, b-matchings, and *maximal b-matchings* in graphs of bounded treewidth and degree. In that work we also strengthen existing results [18, 14] for proper q-colorings in graphs of bounded treewidth and degree.

The key observation that unifies these chains is that, when viewing their state spaces as graphs (exponentially large graphs relative to the input), they all admit a recursive decomposition satisfying key properties. First, each such graph, called a "flip graph," can be partitioned into a small number of induced subgraphs, where each subgraph is a *Cartesian product* of smaller graphs that are structurally similar to the original graph – and thus can be partitioned again into even smaller product graphs. Second, at each level of recursion, pairs of subgraphs are connected by large matchings. Intuitively, we can "slice" a flip graph into subgraphs that are well connected to each other, then "peel" apart the subgraphs using their Cartesian product structure, and repeat the process recursively. Each recursive level of slicing cuts through many edges (the large matchings), and indeed the peeling also disconnects many mutually well-connected subgraphs from one another. Prior work exists applying this "slicing" and "peeling" paradigm – albeit with spectral methods instead of purely combinatorial methods – using Jerrum, Son, Tetali, and Vigoda's decomposition theorem (Theorem 14) for combinatorial chains [21, 18, 14]. One of our contributions is to unify these applications, along with the geometric chains, into a sufficient set of conditions under which one can apply the existing decomposition theorem: Lemma 15.

A more substantial technical contribution is our Theorem 13, a combinatorial analogue to Jerrum, Son, Tetali, and Vigoda's Theorem 14. One can use our theorem in place of theirs and, in some cases, obtain better mixing bounds. In particular, in the case of triangulations, we obtain polynomial mixing via an adaptation of our (combinatorial) technique (Lemma 19) – and it is not clear how to adapt the existing spectral methods to get even a polynomial bound. In the case of k-angulations, our theorem gives a bound that has better dependence on the parameter k.

1.2 Paper organization

In the remainder of this section we will define the Markov chains we are analyzing and summarize our main results. Then, in Section 2, we will give intuition for the decomposition by describing its application to triangulations. In Section 4 we will present our general decomposition meta-theorems, and compare our contribution to prior work by Jerrum, Son, Tetali, and Vigoda [21]. In particular, we will discuss why our purely combinatorial machinery is needed for obtaining new bounds in the case of triangulations. In the full version of our paper [16] we will prove a general result that gives a coarse bound on triangulation mixing; we will then improve this bound to near tightness in the full paper version, and give a matching upper bound (up to logarithmic factors) in the full version. Also in the full version, we show that general k-angulations admit a decomposition satisfying a relaxation (Lemma 18) of our general theorem that implies quasipolynomial-time mixing. We analyze the particular quasipolynomial bound we obtain, and show that our combinatorial technique (Theorem 13) gives a better dependence on k than one would obtain with the prior decomposition theorem. In the full version of the paper we prove our general combinatorial decomposition theorem, Theorem 13. In the full version we prove a theorem about lattice triangulations, and fill in a few remaining proof details.

1.3 Triangulations of convex point sets and lattice point sets

Let P_n be the regular polygon with n vertices. Every triangulation t of P_{n+2} has $n-1$ diagonals, and every diagonal can be *flipped*: every diagonal D belongs to two triangles forming a convex quadrilateral, so D can be removed and replaced with the diagonal D' lying in the same quadrilateral and crossing D. The set of all triangulations of P_{n+2}, for $n \geq 1$, is the vertex set of a graph that we denote K_n (this notation is standard), whose edges are the flips between adjacent triangulations. The graph K_n is known to be realizable as the 1-skeleton of an $n-1$-dimensional polytope [28] called the *associahedron* (we also use this name for the graph itself). It is also known to be isomorphic to the rotation graph on the set of all binary plane trees with $n+1$ leaves [34], and equivalently the set of all parenthesizations of an algebraic expression with $n+1$ terms, with "flips" defined as applications of the associative property of multiplication.

The structure of this graph depends only on the convexity and the number of vertices of the polygon, and not on its precise geometry. That is, P_{n+2} need not be regular for K_n to be well defined.

McShine and Tetali [29] showed that the *mixing time* (see Section 3) of the uniform random walk on $K_{3,n+2}$ is $O(n^5 \log n)$, following Molloy, Reed, and Steiger's [31] lower bound of $\Omega(n^{3/2})$. These bounds together can be shown, using standard inequalities [33], to imply that the *expansion* of $K_{3,n+2}$ is $\Omega(1/(n^4 \log n))$ and $O(n^{1/4})$. It is easy to generalize triangulations to *k-angulations* of a convex polygon $P_{(k-2)n+2}$, and to generalize the definition of a flip between triangulations to a flip between k-angulations: a k-angulation is a maximal division of the polygon into k-gons, and a flip consists of taking a pair of k-gons that share a diagonal, removing that diagonal, and replacing it with one of the other diagonals in the resulting $2k-2$-gon. One can then define the *k-angulation flip walk* on the k-angulations of $P_{(k-2)n+2}$. An analogous graph to the associahedron is defined over the triangulations of the integer lattice (grid) point set with n rows of points and n columns. Substantial prior work has been done on bounds for the number of triangulations in this graph ([4, 22]), as well as characterizing the mixing time of random walks on the graph, when the walks are weighted by a function of the lengths of the edges in a triangulation ([6, 5]).

1.4 Convex triangulation flip walk and mixing time

Consider the following random walk on the triangulations of the convex $n+2$-gon:

for $t = 1, 2, \ldots$ **do**
  Begin with an arbitrary triangulation t.
  Flip a fair coin.
  If the result is tails, do nothing.
  Else, select a diagonal in t uniformly at random, and flip the diagonal.
end for

(The "do nothing" step is a standard MCMC step that enforces a technical condition known as *laziness*, required for the arguments that bound mixing time.) At any given time step, this walk induces a probability distribution π over the triangulations of the $n+2$-gon. Standard spectral graph theory shows that π converges to the uniform distribution in the limit. Formally, what McShine and Tetali showed [29] is that the number of steps before π is within *total variation distance* $1/4$ of the uniform distribution is bounded by $O(n^5 \log n)$ – in other words, that the *mixing time* is $O(n^5 \log n)$. Any polynomial bound means the walk *mixes rapidly.* We formally define total variation distance:

The *total variation distance* between two probability distributions μ and ν over the same set Ω is defined as

$$d(\mu, \nu) = \frac{1}{2} \sum_{S \in \Omega} |\pi(S) - \pi^*(S)|.$$

Consider a Markov chain with state space Ω. Given a starting state $S \in \Omega$, the chain induces a probability distribution π_t at each time step t. Under certain mild conditions, all of which are satisfied by the k-angulation flip walk, this distribution is known to converge in the limit to a *stationary* distribution π^*, which for the k-angulation flip walk is the uniform distribution on the k-angulations of the convex polygon. The *mixing time* is defined as follows: Given an arbitrary $\varepsilon > 0$, the *mixing time*, $\tau(\varepsilon)$, of a Markov chain with state space Ω and stationary distribution π^* is the minimum time t such that, regardless of starting state, we always have

$$d(\pi_t, \pi^*) < \varepsilon.$$

Suppose that the chain belongs to a family of chains, whose size is parameterized by a value n. (It may be that Ω is exponential in n.) If $\tau(\varepsilon)$ is upper bounded by a function that is polynomial in $\log(1/\varepsilon)$ and in n, say that the chain is *rapidly mixing.* It is common to omit the parameter ε, assuming its value to be the arbitrary constant $1/4$.

1.5 Main results

We show the following result for the expansion of the associahedron:

▶ **Theorem 2.** *The expansion of the associahedron* $K_{3,n+2}$ *is* $\Omega(1/(\sqrt{n}\log n))$ *and* $O(1/\sqrt{n})$.

We will prove the lower bound in the full paper version [16] using the *multicommodity flow*-based machinery we introduce in Section 4, after giving intuition in Section 2. Combining this result with the connection between flows and mixing [33] – with some additional effort in the full version – gives our new $O(n^3\log^3 n)$ bound (Theorem 1) for triangulation mixing.

Although the expansion lower bound is more interesting for the sake of rapid mixing, the upper bound in Theorem 2 – which we prove in the full version – recovers Molloy, Reed, and Steiger's $\Omega(n^{3/2})$ mixing lower bound [31]. It is also the first result showing that the associahedron has combinatorial expansion $o(1)$. By contrast, Anari, Liu, Oveis Gharan, and Vinzant recently proved [3, 2], settling a conjecture of Mihail and Vazirani [30], that matroids have expansion one. (Mihail and Vazirani in fact conjectured that all graphs realizable as the 1-skeleton of a *0-1 polytope* have expansion one.) Although the set of convex n-gon triangulations is not a matroid, it is an important subset system – and this work shows that it does not have expansion one. More generally, we give the following quasipolynomial bound for k-angulations:

▶ **Theorem 3.** *For every fixed* $k \geq 3$*, the* k*-angulation flip walk on the convex* $(k-2)n+2$*-point set mixes in time* $n^{O(k\log n)}$.

In the full version of the paper [16], we give a lower bound on the *treewidth* of the $n \times n$ integer lattice point set triangulation flip graph:

▶ **Theorem 4.** *The treewidth of the triangulation flip graph* F_n *on the* $n \times n$ *integer lattice point set is* $\Omega(N^{1-o(1)})$*, where* $N = |V(F_n)|$.

2 Decomposing the convex point set triangulation flip graph

2.1 Bounding mixing via expansion

We have a Markov chain that is in fact a random walk on the associahedron K_n. We wish to bound the mixing time of this walk. It turns out that one way to do this is by lower-bounding the *expansion* of the same graph K_n. Intuitively, expansion concerns the extent to which "bottlenecks" exist in a graph. More precisely, it measures the "sparsest" cut – the minimum ratio of the number of edges in a cut divided by the number of vertices on the smaller side of the cut:

The *edge expansion* (or simply *expansion*), $h(G)$, of a graph $G = (V, E)$ is the quantity

$$\min_{S\subseteq V:|S|\leq|V|/2} |\partial S|/|S|,$$

where $\partial S = \{(s,t) | s \in S, t \notin S\}$ is the set of edges across the $(S, V \setminus S)$ cut. It is known [20, 33] that a lower bound on edge expansion leads to an upper bound on mixing:

▶ **Lemma 5.** *The mixing time of the Markov chain whose transition matrix is the normalized adjacency matrix of a* Δ*-regular graph* G *is*

$$O\left(\frac{\Delta^2\log(|V(G)|)}{(h(G))^2}\right).$$

One can do better [13, 33] if the paths in a *multicommodity flow* are not too long (Section 3).

2.2 "Slicing and peeling"

We would like to show that there are many edges in every cut, relative to the number of vertices on one side of the cut. We partition the triangulations $V(K_n)$ into n equivalence classes, each inducing a subgraph of K_n. We show that many edges exist between each pair of the subgraphs. Thus the partitioning "slices" through many edges. After the partitioning, we show that each of the induced subgraphs has large expansion. To do so, we show that each such subgraph decomposes into many copies of a smaller flip graph K_i, $i < n$. This inductive structure lets us assume that K_i has large expansion – then show that the copies of the smaller flip graph are all well connected to one another. We call this "peeling," because one must peel the many K_i copies from one another – removing many edges – to isolate each copy. Molloy, Reed, and Steiger [31] obtained their $O(n^{25})$ mixing *upper* bound via a different decomposition, namely using the *central* triangle, via a non-flow-based method. That decomposition is the one we use for our quasipolynomial bound for general k-angulations in the full paper version. However, we use a different decomposition here, one with a structure that lets us obtain a nearly tight bound, via a multicommodity flow construction. We formalize the slicing step now:

Fix a "special" edge e^* of the convex $n+2$-gon P_{n+2}. For each triangle T having e^* as one of its edges, define the *oriented class* $\mathcal{C}^*(T)$ to be the set of triangulations of P_{n+2} that include T as one of their triangles. Let $\mathcal{T}_n$ be the set of all such triangles; let $\mathcal{S}_n$ be the set of all classes $\{\mathcal{C}^*(T) | T \in \mathcal{T}_n\}$.

Orient P_{n+2} so that e^* is on the bottom. Then say that T (respectively $\mathcal{C}^*(T)$) is to the *left* of T' (respectively $\mathcal{C}^*(T')$) if the topmost vertex of T lies counterclockwise around P_{n+2} from the topmost vertex of T'. Say that T' lies to the *right* of T. Write $T < T'$ and $T' > T$.

See Figure 1.

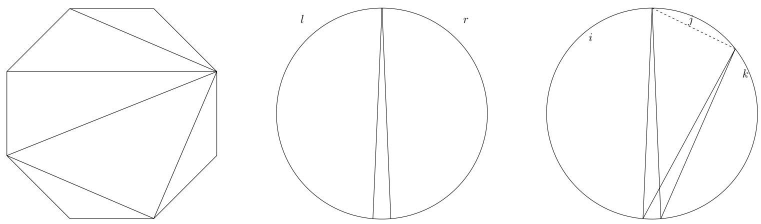

Figure 1 Left: A triangulation of the regular octagon. Center: a class $\mathcal{C}^*(T) \in \mathcal{S}_n$, represented schematically by the triangle T that induces it. We depict the regular $n+2$-gon as a circle (which it approximates as $n \to \infty$), for ease of illustration. Each triangulation $t \in \mathcal{C}^*(T)$ consists of T (the triangle shown), and an arbitrary triangulation of the two polygons on either side of T. Notice that $\mathcal{C}^*(T) \cong K_l \square K_r$, where T partitions the $n+2$-gon into an l-gon and an r-gon. Right: the matching $\mathcal{E}^*(T, T')$ between classes $\mathcal{C}^*(T) \cong K_i \square K_{j+k}$ and $\mathcal{C}^*(T') \cong K_{i+j} \square K_k$, is in bijection with the triangulations in $K_i \square K_j \square K_k$ (induced by the quadrilateral containing T and T'). Therefore, $|\mathcal{E}^*(T, T')| = C_i C_j C_k$.

We make observations about the structure of each class as an induced subgraph of K_n.

The *Cartesian product* graph $G \square H$ of graphs G and H has vertices $V(G) \times V(H)$ and edges

$$\{((u,v),(u',v)) | (u,u') \in E(G), v \in V(H)\}$$

$$\cup \{((u,v),(u,v')) | (v,v') \in E(H), u \in V(G)\}.$$

Given a vertex $w = (u,v) \in V(G) \times V(H)$, call u the *projection* of w onto G, and similarly call v the projection of w onto H. (Applying the obvious associativity of the Cartesian product operator, one can naturally define the product $G_1 \square G_2 \square \cdots \square G_k = \square_{i=1}^k G_i$.)

We can now characterize the structure of each class as an induced subgraph of K_n:

▶ **Lemma 6.** *Each class $\mathcal{C}^*(T)$ is isomorphic to a Cartesian product of two associahedron graphs K_l and K_r, with $l + r = n - 1$.*

Proof. Each triangle T partitions the $n + 2$-gon into two smaller convex polygons with side lengths $l + 1$ and $r + 1$, such that $l + r = n - 1$. Thus each triangulation in $\mathcal{C}^*(T)$ can be identified with a tuple of triangulations of these smaller polygons. The Cartesian product structure then follows from the fact that every flip between two triangulations in $\mathcal{C}^*(T)$ can be identified with a flip in one of the smaller polygons. ◀

Lemma 6 will be central to the peeling step. For the slicing step, building on the idea in Lemma 6 will help us characterize the edge sets between classes:

Given classes $\mathcal{C}^*(T), \mathcal{C}^*(T') \in \mathcal{S}_n$, denote by $\mathcal{E}^*(T, T')$ the set of edges (flips) between $\mathcal{C}^*(T)$ and $\mathcal{C}^*(T')$. Let $\mathcal{B}^*_{n,T'}(T)$ and $\mathcal{B}^*_{n,T}(T')$ be the *boundary sets* – the sets of endpoints of edges in $\mathcal{E}^*(T, T')$ – that lie respectively in $\mathcal{C}^*(T)$ and $\mathcal{C}^*(T')$.

▶ **Lemma 7.** *For each pair of classes $\mathcal{C}^*(T)$ and $\mathcal{C}^*(T')$, the boundary set $\mathcal{B}^*_{n,T'}(T)$ induces a subgraph of $\mathcal{C}^*(T)$ isomorphic to a Cartesian product of the form $K_i \square K_j \square K_k$, for some $i + j + k = n - 2$.*

Proof. Each flip between triangulations in adjacent classes $\mathcal{C}^*(T)$ involves flipping a diagonal of T to transform the triangulation $t \in \mathcal{C}^*(T)$ into triangulation $t' \in \mathcal{C}^*(T')$. Whenever this is possible, there must exist a quadrilateral Q, sharing two sides with T (the sides that are not flipped), such that both t and t' contain Q. Furthermore, every $t \in \mathcal{C}^*(T)$ containing Q has a flip to a distinct $t' \in \mathcal{C}^*(T')$. The set of all such boundary vertices $t \in \mathcal{C}^*(T)$ can be identified with the Cartesian product described because Q partitions P_{n+2} into three smaller polygons, so that each triangulation in $\mathcal{B}^*_{n,T'}(T)$ consists of a tuple of triangulations in each of these smaller polygons, and such that every flip between triangulations in $\mathcal{B}^*_{n,T'}(T)$ consists of a flip in one of these smaller polygons. ◀

▶ **Lemma 8.** *The set $\mathcal{E}^*(T, T')$ of edges between each pair of classes $\mathcal{C}^*(T)$ and $\mathcal{C}^*(T')$ is a nonempty matching. Furthermore, this edge set is in bijection with the vertices of a Cartesian product $K_i \square K_j \square K_k, i + j + k = n - 2$.*

Proof. The claim follows from the reasoning in Lemma 7 and from the observation that each triangulation in $\mathcal{B}^*_{n,T'}(T)$ has exactly one flip (namely, flipping a side of the triangle T) to a neighbor in $\mathcal{B}^*_{n,T}(T')$. ◀

Lemma 8 characterizes the structure of the edge sets (namely matchings) between classes; we would also like to know the sizes of the matchings. We will use the following formula:

Let C_n be the nth *Catalan number*, defined as $C_n = \frac{1}{n+1}\binom{2n}{n}$.

▶ **Lemma 9** ([25, 19])**.** *The number of vertices in the associahedron K_n is C_n, and this number grows as $\frac{1}{\sqrt{\pi} \cdot n^{3/2}} \cdot 2^{2n}$.*

We will prove the following in the full version of our paper [16]:

▶ **Lemma 10.** *For every $T, T' \in \mathcal{T}_n$,*

$$|\mathcal{E}^*(T, T')| \geq \frac{|\mathcal{C}^*(T)||\mathcal{C}^*(T')|}{C_n}.$$

Lemma 10 – which states that the number of edges between a pair of classes is at least equal to the product of the cardinalities of the classes, divided by the total number of vertices in the graph $|V(K_n)| = C_n$ – is crucial to this paper. To explain why this is, we will need to present our multicommodity flow construction (in the full version of the paper [16]). We will give intuition in Section 4. For now, it suffices to say that Lemma 10 implies that there are many edges between a given pair of classes, justifying (intuitively) the slicing step. For the peeling step, we need the fact that Cartesian graph products preserve the well-connectedness of the graphs in the product [17]:

▶ **Lemma 11.** *Given graphs $G_1, G_2, \ldots, G_k$, Cartesian product $G_1 \square G_2 \square \cdots \square G_k$ satisfies*

$$h(G_1 \square G_2 \square \cdots \square G_k) \geq \frac{1}{2} \min_i h(G_i).$$

Lemma 6 says that each of the classes $\mathcal{C}^*(T) \in \mathcal{S}_n$ is a Cartesian graph product of associahedron graphs K_l, K_r, $l < n, r < n$, allowing us to "peel" (decompose) $\mathcal{C}^*(T)$ into graphs that can then be recursively sliced into classes and peeled. Lemma 11 implies that the peeling must disconnect many edges, as it involves splitting a Cartesian product graph into many subgraphs (copies of K_l).

We will make all of this intuition rigorous in the full paper version by constructing our flow. The choice of paths through which to route flow will closely trace the edges in this recursive "slicing and peeling" decomposition. We will then show that, with this choice of paths, the resulting *congestion* – the maximum amount of flow carried along an edge – is bounded by a suitable polynomial factor. This will provide a lower bound on the expansion.

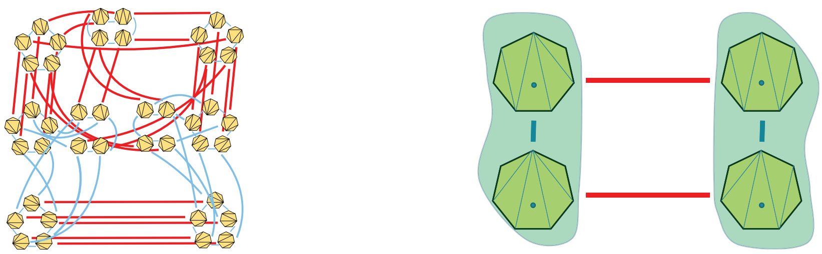

Figure 2 Left: The associahedron graph K_5, with each vertex representing a triangulation of the regular heptagon. Flips are shown with edges (in blue and red). The vertex set $V(K_n)$ is partitioned into a set $\mathcal{S}_n$ of five equivalence classes (of varying sizes). Within each class, all triangulations share the same triangle containing the bottom edge e^*. Flips (edges) between triangulations in the same class are shown in blue. Flips between triangulations in different classes are shown in red. To "slice" K_5 into its subgraphs, one must cut through these red matchings. **Right**: A class $\mathcal{C}^*(T)$ from the graph K_5 on the left-hand side, viewed as an induced subgraph of K_5. The identifying triangle T is marked with a blue dot. This subgraph is isomorphic to a Cartesian product of two K_2 graphs; each copy of K_2 induced by fixing the rightmost diagonal is outlined in green. "Peeling" apart this product requires disconnecting the two red edges connecting the K_2 copies.

3 Bounding expansion via multicommodity flows

The way we will lower-bound expansion is by using *multicommodity flows* [33, 23]. A *multicommodity flow* ϕ in a graph $G = (V, E)$ is a collection of functions $\{f_{st} : A \to \mathbb{R} \mid s, t \in V\}$, where $A = \bigcup_{\{u,v\} \in E} \{(u, v), (v, u)\}$, combined with a *demand function* $D : V \times V \to \mathbb{R}$.

Each f_{st} is a flow sending $D(s,t)$ units of a commodity from vertex s to vertex t through the edges of G. We consider the capacities of all edges to be infinite. Let $f_{st}(u,v)$ be the amount of flow sent by f_{st} across the arc (u,v). (It may be that $f_{st}(u,v) \neq f_{st}(v,u)$.) Let

$$f(u,v) = \frac{1}{|V|} \sum_{s,t \in V \times V} f_{st}(u,v),$$

and let $\rho = \max_{(u,v)\in A} f(u,v)$. Call ρ the *congestion*. Unless we specify otherwise, we will mean by "multicommodity flow" a *uniform multicommodity flow*, i.e. one in which $D(s,t) = 1$ for all s,t. The following is well established and enables the use of multicommodity flows as a powerful lower-bounding technique for expansion:

▶ **Lemma 12.** *Given a uniform multicommodity flow f in a graph $G = (V,E)$ with congestion ρ, the expansion $h(G)$ is at least $1/(2\rho)$.*

Lemma 12, combined with Lemma 5, gives an automatic upper bound on mixing time given a multicommodity flow with an upper bound on congestion – but with a quadratic loss. As we will discuss in the full paper version, one can do better if the paths used in the flow are short [13, 33].

4 Our framework

In addition to the new mixing bounds for triangulations and for general k-angulations, we make general technical contributions, in the form of three meta-theorems, which we present in this section. Our first general technical contribution, Theorem 13, provides a recursive mechanism for analyzing the expansion of a flip graph in terms of the expansion of its subgraphs. Equivalently, viewing the random walk on such a flip graph as a Markov chain, this theorem provides a mechanism for analyzing the mixing time of a chain, in terms of the mixing times of smaller *restriction* chains into which one decomposes the original chain – and analyzing a *projection* chain over these smaller chains. We obtain, in certain circumstances such as the k-angulation walk, better mixing time bounds than one obtains applying similar prior decomposition theorems – which used a different underlying machinery.

The second theorem, Lemma 15, observes and formalizes a set of conditions satisfied by a number of chains (equivalently, flip graphs) under which one can apply either our Theorem 13, or prior decomposition techniques, to obtain rapid mixing reuslts. Depending on the chain, one may then obtain better results either by applying Theorem 13, or by applying the prior techniques. Lemma 15 does not require using our Theorem 13; instead, one can use the spectral gap or log-Sobolev constant as the underlying techincal machinery using Jerrum, Son, Tetali, and Vigoda's Theorem 14. Prior work exists applying these techniques (using Theorem 14) to sampling q-colorings [18] in bounded-treewidth graphs and independent sets in regular trees [21], as well as probabilistic graphical models in machine learning [11] satisfying certain conditions. Lemma 15 amounts to an observation unifying these applications. We apply this observation to general k-angulations, noting that they satisfy a relaxation of this theorem (Lemma 18), giving a quasipolynomial bound. This bound will come from incurring a polynomial loss over logarithmic recursion depth.

The third theorem, Lemma 19, adapts the machinery in Theorem 13 to eliminate this multiplicative loss altogether, assuming that a chain satisfies certain properties. One such key property is the existence large matchings in Lemma 10 in Section 2. Another property, which we will discuss further after presenting Lemma 19, is that the *boundary sets* – the vertices in one class (equivalently, states in a restriction chain) having neighbors in another class –

are well connected to the rest of the first class. When these properties are satisfied, one can apply our flow machinery to overcome the multiplicative loss and obtain a polynomial bound. However, the improvement relies on observations about congestion that do not obviously translate to the spectral setting.

4.1 Markov chain decomposition via multicommodity flow

In this section we state our first general theorem. To place our contribution in context with prior work, we cast our flip graphs in the language of Markov chains. As we discussed in Section 1.4, any Markov chain satisfying certain mild conditions has a *stationary* distribution π^* (which in the case of our triangulation walks is uniform). We can view such a chain as a random walk on a graph $\mathcal{M}$ (an unweighted graph in the case of the chains we consider, which have uniform distributions and regular transition probabilities). In the case of convex polygon triangulations, we have $\mathcal{M} = K_n$.

The flip graph $\mathcal{M}$ has vertex set Ω and (up to normalization by degree) adjacency matrix P – and we abuse notation, identifying the Markov chain $\mathcal{M}$ with this graph. When π^* is not uniform, it is easy to generalize the flip graph to a *weighted* graph, with each vertex (state) t assigned weight $\pi(t)$, and each transition (edge) (t, t') assigned weight $\pi(t)P(t, t') = \pi(t')P(t', t)$. We assume here that this latter equality holds, a condition on the chain $\mathcal{M}$ known as *reversibility*. We then replace a uniform multicommodity flow with one where $D(t, t') = \pi(t)\pi(t')$ (up to normalization factors).

Consider a Markov chain $\mathcal{M}$ with finite state space Ω and probability transition matrix P, and stationary distribution π. Consider a partition of the states of Ω into classes $\Omega_1, \Omega_2, \dots, \Omega_k$. Let the *restriction* chain, for $i = 1, \dots, k$, be the chain with state space Ω_i, probability distribution π_i, with $\pi_i(x) = \pi(x)/(\sum_{y\in\Omega_i} \pi(y))$, for $x \in \Omega_i$, and transition probabilities $P_i(x, y) = P(x, y)/(\sum_{z\in\Omega_i} P(x, z))$. Let the *projection* chain be the chain with state space $\bar{\Omega} = \{1, 2, \dots, k\}$, stationary distribution $\bar{\pi}$, with $\bar{\pi}(i) = \sum_{x\in\Omega_i} \pi(i)$, and transition probabilities $\bar{P}(i, j) = \sum_{x\in\Omega_i, y\in\Omega_j} P(x, y)$.

▶ **Theorem 13.** *Let $\mathcal{M}$ be a reversible Markov chain with finite state space Ω probability transition matrix P, and stationary distribution π^*. Suppose $\mathcal{M}$ is connected (irreducible). Suppose $\mathcal{M}$ can be decomposed into a collection of restriction chains $(\Omega_1, P_1), (\Omega_2, P_2), \dots, (\Omega_k, P_k)$, and a projection chain $(\bar{\Omega}, \bar{P})$. Suppose each restriction chain admits a multicommodity flow (or canonical paths) construction with congestion at most $\rho_{\max}$. Suppose also that there exists a multicommodity flow construction in the projection chain with congestion at most $\bar{\rho}$. Then there exists a multicommodity flow construction in $\mathcal{M}$ (viewed as a weighted graph in the natural way) with congestion*

$$(1 + 2\bar{\rho}\gamma\Delta)\rho_{\max},$$

where $\gamma = \max_{i\in[k]} \max_{x\in\Omega_i} \sum_{y\notin\Omega_i} P(x, y)$, and Δ is the degree of $\mathcal{M}$.

We give a full proof in the full version of the paper. Jerrum, Son, Tetali, and Vigoda [21] presented an analogous (and classic) decomposition theorem, which we restate below as Theorem 14, and which has become a standard tool in mixing time analysis. The key difference between our theorem and theirs is that our theorem uses multicommodity flows, while their theorem uses the so-called *spectral gap* – another parameter that can use to bound the mixing time of a chain. Often, the spectral gap gives tighter mixing bounds than combinatorial methods. Their Theorem 14 gave bounds analogous to our Theorem 13, but with the multicommodity flow congestion replaced with the *spectral gap* of a chain

– and with a 3γ term in place of our 2γ. (They also gave an analogous version for the *log-Sobolev* constant – yet another parameter for bounding mixing times.) The spectral gap of a chain $\mathcal{M} = (\Omega, P)$, which we denote λ, is the difference between the two largest eigenvalues of the transition matrix P (which we can view as the normalized adjacency matrix of the corresponding weighted graph). The key point is that while on the one hand the mixing time τ satisfies $\tau \le \lambda^{-1} \log |\Omega|$, the bound on mixing using expansion in Lemma 5 comes from passing through the spectral gap: $\lambda \ge \frac{(h(\mathcal{M}))^2}{2\Delta^2}$, where Δ is the degree of the flip graph and $h(\mathcal{M})$ is the expansion of $\mathcal{M}$. The quadratic loss in passing from expansion to mixing is not incurred when bounding the spectral gap directly, so one can obtain better bounds via the spectral gap. Jerrum, Son, Tetali, and Vigoda gave a mechanism for doing precisely this:

▶ **Theorem 14** ([21]). *Let $\mathcal{M}$ be a reversible Markov chain with finite state space Ω probability transition matrix P, and stationary distribution π^*. Suppose $\mathcal{M}$ is connected (irreducible). Suppose $\mathcal{M}$ can be decomposed into a collection of restriction chains $(\Omega_1, P_1), (\Omega_2, P_2), \ldots, (\Omega_k, P_k)$, and a projection chain $(\bar{\Omega}, \bar{P})$. Suppose each restriction chain has spectral gap at least $\lambda_{\min}$. Suppose also that the projection chain has spectral gap at least $\bar{\lambda}$. Then $\mathcal{M}$ has gap at least*

$$\min\left\{\frac{\lambda_{\min}}{3}, \frac{\bar{\lambda}\lambda_{\min}}{3\gamma + \bar{\lambda}}\right\},$$

where γ is as in Theorem 13.

Our Theorem 13 has a simple, purely combinatorial proof (in the full paper version), and fills a gap in the literature by showing that such a construction can be used in place of the spectral machinery from the earlier technique. We also obtain a tighter bound on expansion than would result from a black-box application of Theorem 14. The cost to our improvement is in passing from expansion to mixing via the spectral gap. Nonetheless, we will show that in the case of triangulations, our Theorem 13 can be adapted to give a new mixing bound whereas, by contrast, it is not clear how to obtain even a polynomial bound adapting Jerrum, Son, Tetali, and Vigoda's spectral machinery. We will also show that for general k-angulations, one can, with our technique, use a combinatorial insight to eliminate the γ factor in our decomposition in favor of a Δ^{-1} factor (for k-angulations we have $\gamma = k/\Delta$) – whereas it is not clear how to do so with the spectral decomposition.

4.2 General pattern for bounding projection chain congestion

Our second decomposition theorem, which we will apply to general k-angulations, states that if one can recursively decompose a chain into restriction chains in a particular fashion, and if the projection chain is well connected, then Theorem 13 gives an expansion bound:

▶ **Lemma 15.** *Let $\mathcal{F} = \{\mathcal{M}_1, \mathcal{M}_2, \ldots\}$ be a family of connected graphs, parameterized by a value n. Suppose that every graph $\mathcal{M}_n = (\mathcal{V}_n, \mathcal{E}_n) \in \mathcal{F}$, for $n \ge 2$, can be partitioned into a set $\mathcal{S}_n$ of classes satisfying the following conditions:*

1. *Each class in $\mathcal{S}_n$ is isomorphic to a Cartesian product of one or more graphs $\mathcal{C}(T) \cong \mathcal{M}_{i_1} \square \cdots \mathcal{M}_{i_k}$, where for each such graph $\mathcal{M}_{i_j} \in \mathcal{F}$, $i_j \le n/2$.*
2. *The number of classes is $O(1)$.*
3. *For every pair of classes $\mathcal{C}(T), \mathcal{C}(T') \in \mathcal{S}_n$ that share an edge, the number of edges between the two classes is $\Omega(1)$ times the size of each of the two classes.*
4. *The ratio of the sizes of any two classes is $\Theta(1)$.*

Suppose further that $|\mathcal{V}_1| = 1$. Then the expansion of $\mathcal{M}_n$ is $\Omega(n^{-O(1)})$.

Lemma 15 is easy to prove given Theorem 13. An analogue in terms of spectral gap is easy to prove given Theorem 14. Furthermore, as we will prove in the full paper version, a precise statement of the bounds given by Lemma 15 is as follows:

▶ **Lemma 16.** *Suppose a flip graph $\mathcal{M}_n = (\mathcal{V}_n, \mathcal{E}_n)$ belongs to a family $\mathcal{F}$ of graphs satisfying the conditions of Lemma 15. Suppose further that every graph $\mathcal{M}_k = (\mathcal{V}_k, \mathcal{E}_k) \in \mathcal{F}$, $k < n$, satisfies*

$$|\mathcal{V}_k|/|\mathcal{E}_{k,\min}| \leq f(k),$$

for some function $f(k)$, where $\mathcal{E}_{k,\min}$ is the smallest edge set between adjacent classes $\mathcal{C}(T), \mathcal{C}(T') \in \mathcal{S}_k$, where $\mathcal{S}_k$ is as in Lemma 15. Then the expansion of $\mathcal{M}_n$ is

$$\Omega(1/(2f(n))^{\log n})),$$

where γ is as in Theorem 13, and Δ is the degree of $\mathcal{M}_n$.

Proof. Constructing an arbitrary multicommodity flow (or set of canonical paths) in the projection graph at each inductive step gives the result claimed. The term $|\mathcal{V}_k|/|\mathcal{E}_{k,\min}|$ bounds the (normalized) congestion in any such flow because the total amount of flow exchanged by all pairs of vertices (states) combined is $|\mathcal{V}_k|^2$, and the minimum weight of an edge in the projection graph is $|\mathcal{E}_{k,\min}|$.

Notice that we do not incur a $\gamma\Delta$ term here, because even if a state (vertex) in $\Omega_i \subseteq \mathcal{V}_k$ has neighbors $x \in \Omega_j, y \in \Omega_l$, z still only receives no more than $|\mathcal{V}_k|^2/\mathcal{E}_{k,\min}\}$ flow across the edges (z, x) and (z, y) combined. ◀

▶ Remark 17. The $\gamma\Delta$ factor in Theorem 13, which does not appear in Lemma 16, does appear in a straightforward appliation of Jerrum, Son, Tetali, and Vigoda's Theorem 14.

We will show that k-angulations (with fixed $k \geq 4$) satisfy a relaxation of Lemma 15:

▶ **Lemma 18.** *Suppose a family $\mathcal{F}$ of graphs satisfies the conditions of Lemma 15, with the $\Omega(1)$, $O(1)$, and $\Theta(1)$ factors in Conditions 3, 2, and 4 respectively replaced by $\Omega(n^{-O(1)})$, $O(n^{O(1)})$, and $\Theta(n^{O(1)})$. Then for every $\mathcal{M}_n \in \mathcal{F}$, the expansion of $\mathcal{M}_n$ is $\Omega(n^{-O(\log n)})$.*

Lemma 15 enables us to relate a number of chains admitting a certain decomposition process in a black-box fashion, unifying prior work applying Theorem 14 separately to individual chains. Marc Heinrich [18] presented a similar but less general construction for the Glauber dynamics on q-colorings in bounded-treewidth graphs; other precursors exist, including for the hardcore model on certain trees [21] and a general argument for a class of graphical models [11]. In the companion paper [15] we mentioned in Section 1, we apply Lemma 15 to chains for sampling independent sets and dominating sets in bounded-treewidth graphs, as well as chains on q-colorings, maximal independent sets, and several other structures, in graphs whose treewidth and degree are bounded.

4.3 Eliminating inductive loss: nearly tight conductance for triangulations

We now give the meta-theorem that we will apply to triangulations. Lemma 15 – using either Theorem 13 or Theorem 14 – gives a merely quasipolynomial bound when applied straightforwardly to k-angulations, including the case of triangulations – simply because the $f(n)$ term in Lemma 16 is $\omega(1)$ and thus the overall congestion is $\omega(1)^{\log n}$ (not polynomial). However, it turns out that the large matchings given by Lemma 10 between pairs of classes in the case of triangulations (but not general k-angulations), combined with some additional structure in the triangulation flip walk, satisfy an alternative set of conditions that suffice for rapid mixing. The conditions are:

▶ **Lemma 19.** *Let $\mathcal{F} = \{\mathcal{M}_1, \mathcal{M}_2, \dots\}$ be an infinite family of connected graphs, parameterized by a value n. Suppose that for every graph $\mathcal{M}_n = (\mathcal{V}_n, \mathcal{E}_n) \in \mathcal{F}$, for $n \geq 2$, the vertex set $\mathcal{V}_n$ can be partitioned into a set $\mathcal{S}_n$ of classes inducing subgraphs of $\mathcal{M}_n$ that satisfy the following conditions:*

1. *Each subgraph is isomorphic to a Cartesian product of one or more graphs $\mathcal{C}(T) \cong \mathcal{M}_{i_1} \square \cdots \mathcal{M}_{i_k}$, where for each such graph $\mathcal{M}_{i_j} \in \mathcal{F}$, $i_j < n$.*
2. *The number of classes is $n^{O(1)}$.*
3. *For every pair of classes $\mathcal{C}(T), \mathcal{C}(T') \in \mathcal{S}_n$, the set of edges between the subgraphs induced by the two classes is a matching of size at least $\frac{|\mathcal{C}(T)||\mathcal{C}(T')|}{|\mathcal{V}_n|}$.*
4. *Given a pair of classes $\mathcal{C}(T), \mathcal{C}(T') \in \mathcal{S}_n$, there exists a graph $\mathcal{M}_i$ in the Cartesian product $\mathcal{C}(T)$, and a class $\mathcal{C}(U) \in \mathcal{S}_i$ within the graph $\mathcal{M}_i$, such that the set of vertices in $\mathcal{C}(T)$ having a neighbor in $\mathcal{C}(T')$ is precisely the set of vertices in $\mathcal{C}(T)$ whose projection onto $\mathcal{M}_i$ lies in $\mathcal{C}(U)$. Furthermore, no class $\mathcal{C}(U)$ within $\mathcal{M}_i$ is the projection of more than one such boundary.*

Suppose further that $|\mathcal{V}_1| = 1$. Then the expansion of $\mathcal{M}_n$ is $\Omega(1/(\kappa(n)n))$, where $\kappa(n) = \max_{1 \leq i \leq n} |\mathcal{C}(S_i)|$ is the maximum number of classes in any $\mathcal{M}_i, i \leq n$.

Unlike Lemma 15, this lemma requires a purely combinatorial construction; it is not clear how to apply spectral methods to obtain even a polynomial bound. Condition 4 is crucial. To give more intuition for this condition, we state and prove the following fact about the triangulation flip graph (visualized in Figure 3):

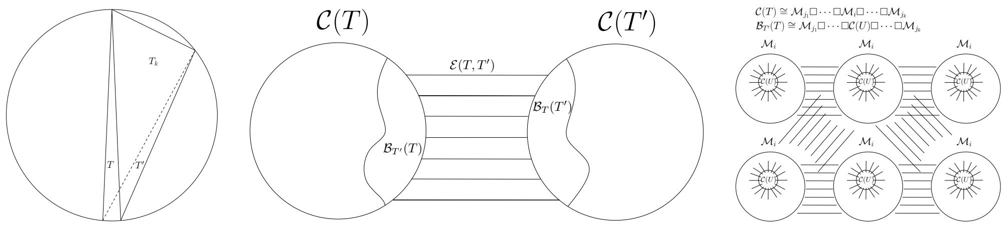

Figure 3 Left: (Lemma 20) The set of edges $\mathcal{E}^*(T,T')$ has $K_i \square \mathcal{C}^*(T_k)$ as its set of boundary vertices in $\mathcal{C}^*(T)$. **Center**: An illustration of Condition 3 in Lemma 19, showing a large matching $\mathcal{E}(T,T')$ between two classes (subgraphs) $\mathcal{C}(T)$ and $\mathcal{C}(T')$. **Right**: An illustration of Conditions 1 and 4 in Lemma 19: $\mathcal{C}(T)$ as a Cartesian product of smaller graphs $\mathcal{M}_{j_1}, \dots, \mathcal{M}_i, \dots, \mathcal{M}_{j_k}$ in the family $\mathcal{F}$. The schematic view shows this Cartesian product as a collection of copies of $\mathcal{M}_i$, connected via perfect matchings between pairs of the copies – with the pairs to connect determined by the structure of the Cartesian product. The boundary $\mathcal{B}_{T'}(T)$ (center) is isomorphic to a class $\mathcal{C}(U)$ (right) within $\mathcal{M}_i$, a graph in the product. Within each copy of $\mathcal{M}_i$, many edges connect $\mathcal{C}(U)$ to the rest of $\mathcal{M}_i$.

▶ **Lemma 20.** *Given $T, T' \in \mathcal{T}_n$, suppose T' lies to the right of T. Then the subgraph of $\mathcal{C}^*(T)$ induced by $\mathcal{B}^*_{n,T'}(T)$ is isomorphic to a Cartesian product $K_l \square \mathcal{C}^*(T_k)$, where $l + r = n - 1$, and where T_k has as an edge the right diagonal of T, and as the vertex opposite this edge the topmost vertex of T'. A symmetric fact holds for $\mathcal{B}^*_{n,T}(T')$.*

Proof. Every triangulation in $\mathcal{B}^*_{n,T'}(T)$ (i) includes the triangle T and (ii) is a single flip away from including the triangle T'. As we observed in the proof of Lemma 7, this implies that $\mathcal{B}^*_{n,T'}(T)$ consists of the set of triangulations in $\mathcal{C}^*(T)$ containing a quadrilateral Q.

Specifically, Q shares two sides with T: one of these is e^*, and the other is the left side of T. One of the other two sides of Q is the right side of $\mathcal{C}^*(T')$. Combining this side with the "top" side of Q and with the right side of T, one obtains the triangle T_k, proving the claim. ◀

Lemma 20 implies that there are many edges between the boundary set $\mathcal{B}^*_{n,T'}(T)$ and the rest of $\mathcal{C}^*(T)$: $\mathcal{C}^*(T) \cong K_l \square K_r$, where K_l and K_r are smaller associahedron graphs, so $\mathcal{C}^*(T)$ is a collection of copies of K_r, with pairs of copies connected by perfect matchings. Each K_r copy can itself be decomposed into a set $\mathcal{S}_r$ of classes, one of which, namely $\mathcal{C}^*(T_k)$, is the intersection of $\mathcal{B}^*_{n,T'}(T)$ with the K_r copy. Applying Condition 3 to the K_r copy implies that there are many edges between boundary vertices in $\mathcal{C}^*(T_k)$ to other subgraphs (classes) in the K_r copy. That is, the boundary set $\mathcal{B}^*_{n,T'}(T)$ is well connected to the rest of $\mathcal{C}^*(T)$.

Figure 3 visualizes this situation in general terms for the framework. We have now proven:

▶ **Lemma 21.** *The associahedron graph K_n, along with the oriented partition we have defined, satisfies the conditions of Lemma 19.*

Proof. The connectedness of K_n is known [29]. Conditions 1 and 3 follow from Lemma 6, Lemma 8, and Lemma 10. Concerning the boundary sets, Condition 4 follows from Lemma 20 and from the discussion leading to this lemma. ◀

Together with Lemma 5 and the easy fact that K_n is a $\Theta(n)$-regular graph, Lemma 21 implies rapid mixing, pending the proof of Lemma 19 – which we prove in the full paper version [16].

4.4 Intuition for the flow construction for triangulations

We will prove Lemma 19 in the full paper version, from which a coarse expansion lower bound for triangulations – and a corresponding coarse (but polynomial) upper bound for mixing – will be immediate by Lemma 21. We give some intuition now for the flow construction we will give in the proof of Lemma 19, and in particular for the centrality of Condition 3 and Condition 4 (corresponding respectively to Lemma 10 and Lemma 20 for triangulations). Consider the case of triangulations, for concreteness. Every $t \in \mathcal{C}^*(T), t' \in \mathcal{C}^*(T')$ must exchange a unit of flow. This means that a total of $|\mathcal{C}^*(T)||\mathcal{C}^*(T')|$ flow must be sent across the matching $\mathcal{E}^*(T,T')$. To minimize congestion, it will be optimal to equally distribute this flow across all of the boundary matching edges. We can decompose the overall problem of routing flow from each $t \in \mathcal{C}^*(T)$ to each $t' \in \mathcal{C}^*(T')$ into three subproblems: (i) *concentrating* flow from every triangulation in $\mathcal{C}^*(T)$ within the boundary set $\mathcal{B}^*_{n,T'}(T)$, (ii) routing flow across the matching edges $\mathcal{E}^*(T,T')$, i.e. from $\mathcal{B}^*_{n,T'}(T) \subseteq \mathcal{C}^*(T)$ to $\mathcal{B}^*_{n,T}(T') \subseteq \mathcal{C}^*(T')$, and (iii) *distributing* flow from the boundary $\mathcal{B}^*_{n,T}(T')$ to each $t' \in \mathcal{C}^*(T')$. Now, the amount of flow that must be concentrated from $\mathcal{C}^*(T)$ at *each* boundary triangulation $u \in \mathcal{B}^*_{n,T'}(T)$ (and symmetrically distributed from *each* $v \in \mathcal{B}^*_{n,T}(T')$ throughout $\mathcal{C}^*(T')$) is equal to

$$\frac{|\mathcal{C}^*(T)||\mathcal{C}^*(T')|}{|\mathcal{B}^*_{n,T'}(T)|} = \frac{|\mathcal{C}^*(T)||\mathcal{C}^*(T')|}{|\mathcal{B}^*_{n,T}(T')|} = \frac{|\mathcal{C}^*(T)||\mathcal{C}^*(T')|}{|\mathcal{E}^*(T,T')|} \leq C_n,$$

where we have used the equality $|\mathcal{B}^*_{n,T'}(T)| = |\mathcal{B}^*_{n,T}(T')| = |\mathcal{E}^*(T,T')|$ by Lemma 7 and Lemma 8, and where the inequality follows from Lemma 10. As a result, in the "concentration" and "distribution" subproblems (i) and (iii), at most C_n flow is concentrated at or distributed from any given triangulation (Figure 4). This bound yields a recursive structure: the concentration (respectively distribution) subproblem decomposes into a flow problem

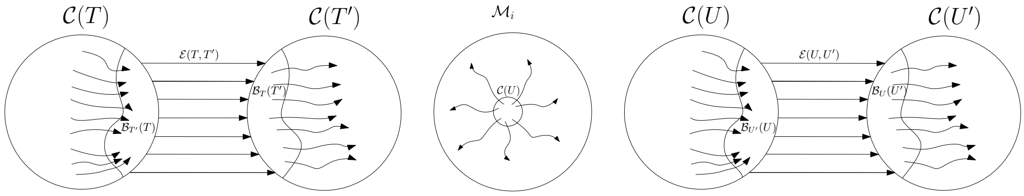

Figure 4 Left: The problem of sending flow from each $t \in \mathcal{C}^*(T)$ to each $t' \in \mathcal{C}^*(T')$, decomposed into subproblems: (i) *concentrating* flow within $\mathcal{B}^*_{n,T'}(T)$, (ii) *transmitting* the flow across the boundary matching $\mathcal{E}^*(T,T')$, and (iii) *distributing* the flow from $\mathcal{B}^*_{n,T}(T')$ throughout $\mathcal{C}^*(T')$. **Center**: Within each copy of $\mathcal{M}_i$ in the product $\mathcal{C}^*(T') \cong \mathcal{M}_{j_1} \square \cdots \square \mathcal{M}_i \square \cdots \square \mathcal{M}_{j_k}$, the distribution problem in Figure 4 induces the problem of distributing flow from a class $\mathcal{C}^*(U)$ – namely the projection of $\mathcal{B}^*_{n,T}(T')$ onto $\mathcal{M}_i$ – throughout the rest of $\mathcal{M}_i$. **Right**: The problem in the center figure induces subproblems in which $\mathcal{C}^*(U) \subseteq \mathcal{M}_i$ must send flow to each $\mathcal{C}^*(U') \subseteq \mathcal{M}_i$. These subproblems are of the same form as the original $\mathcal{C}^*(T), \mathcal{C}^*(T')$ problem (left), and can be solved recursively. The large matchings $\mathcal{E}^*(T,T'), \mathcal{E}^*(U,U')$ guaranteed by Condition 3 prevent any recursive congestion increase.

within $\mathcal{C}^*(T)$ (respectively $\mathcal{C}^*(T')$), in which, by the inequality, each triangulation has C_n total units of flow it must receive (or send). We will then apply Condition 4, observing (see Figure 4) that the concentration (symmetrically) distribution of this flow can be done entirely between pairs of classes $\mathcal{C}^*(U), \mathcal{C}^*(U')$ within copies of a smaller flip graph $\mathcal{M}_i$ in the Cartesian product $\mathcal{C}^*(T') \cong \mathcal{M}_{j_1} \square \cdots \square \mathcal{M}_i \square \cdots \square \mathcal{M}_{j_k}$.

The $\mathcal{C}^*(U), \mathcal{C}^*(U')$ subproblem is of the same form as the original $\mathcal{C}^*(T), \mathcal{C}^*(T')$ problem (Figure 4), and we will show that the C_n bound on the flow (normalizing to congestion one) across the $\mathcal{E}^*(T,T')$ edges will induce the same C_n bound across the $\mathcal{E}^*(U,U')$ edges in the induced subproblem. We further decompose the $\mathcal{C}^*(U), \mathcal{C}^*(U')$ problem into concentration, transmission, and distribution subproblems without any gain in overall congestion. To see this, view the initial flow problem in K_n as though every triangulation $t \in V(K_n)$ is initially "charged" with $|V(K_n)| = C_n$ total units of flow to distribute throughout K_n. Similarly, in the induced distribution subproblem within each copy of $\mathcal{M}_i = K_i$ in the product $\mathcal{C}^*(T')$, each vertex on the boundary $\mathcal{B}^*_{n,T'}(T)$ is initially "charged" with C_n total units to distribute throughout K_i. Just as the original problem in K_n results in each $\mathcal{E}^*(T,T')$ carrying at most C_n flow across each edge, similarly (we will show in the full paper version) the induced problem in K_i results in each $\mathcal{E}^*(U,U')$ carrying at most C_n flow across each edge. This preservation of the bound C_n under the recursion avoids any congestion increase.

One must be cautious, due to the linear recursion depth, not to accrue even a constant-factor loss in the recursive step (the coefficient 2 in Theorem 13). In Theorem 13, it turns out that this loss comes from routing *outbound* flow within a class $\mathcal{C}^*(T)$ – flow that must be sent to other classes – and then also routing *inbound* flow. The combination of these steps involves two "recursive invocations" of a uniform multicommodity flow that is inductively assumed to exist within $\mathcal{C}^*(T)$. We will show in the full paper version that one can avoid the second "invocation" with an initial "shuffling" step: a uniform flow within $\mathcal{C}^*(T)$ in which each triangulation $t \in \mathcal{C}^*(T)$ distributes all of its outbound flow evenly throughout $\mathcal{C}^*(T)$.

It is here that Jerrum, Son, Tetali, and Vigoda's spectral Theorem 14 breaks down, giving a 3-factor loss at each recursion level, due to applying the Cauchy-Schwarz inequality to a *Dirichlet form* that is decomposed into expressions over the restriction chains. Although Jerrum, Son, Tetali, and Vigoda gave circumstances for mitigating or eliminating their multiplicative loss, this chain does not satisfy those conditions in an obvious way.

References

1 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1319–1330, 2020. doi:10.1109/FOCS46700.2020.00125.

2 Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: High-dimensional walks and an FPRAS for counting bases of a matroid. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3313276.3316385.

3 Nima Anari, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials, entropy, and a deterministic approximation algorithm for counting bases of matroids. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 35–46, 2018. doi:10.1109/FOCS.2018.00013.

4 Emile E. Anclin. An upper bound for the number of planar lattice triangulations. *Journal of Combinatorial Theory, Series A*, 103(2):383–386, 2003. doi:10.1016/S0097-3165(03)00097-9.

5 Pietro Caputo, Fabio Martinelli, Alistair Sinclair, and Alexandre Stauffer. Dynamics of lattice triangulations on thin rectangles. *Electronic Journal of Probability*, 21, May 2015. doi:10.1214/16-EJP4321.

6 Pietro Caputo, Fabio Martinelli, Alistair Sinclair, and Alexandre Stauffer. Random lattice triangulations: structure and algorithms. *Annals of Applied Probability*, 25:1650–1685, 2015.

7 Alessandra Caraceni and Alexandre Stauffer. Polynomial mixing time of edge flips on quadrangulations. *Probability Theory and Related Fields*, 176(1):35–76, February 2020. doi:10.1007/s00440-019-00913-5.

8 Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Rapid mixing for colorings via spectral independence. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1548–1557, 2021. doi:10.1137/1.9781611976465.94.

9 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of Glauber dynamics up to uniqueness via contraction, 2020. arXiv:2004.09083.

10 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of glauber dynamics: Entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1537–1550, 2021.

11 Christopher M De Sa, Ce Zhang, Kunle Olukotun, and Christopher Ré. Rapidly Mixing Gibbs Sampling for a Class of Factor Graphs Using Hierarchy Width. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL: https://proceedings.neurips.cc/paper/2015/file/b29eed44276144e4e8103a661f9a78b7-Paper.pdf.

12 Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible markov chains. *The Annals of Applied Probability*, 3(3):696–730, 1993.

13 Persi Diaconis and Daniel Stroock. Geometric Bounds for Eigenvalues of Markov Chains. *The Annals of Applied Probability*, 1(1):36–61, 1991. doi:10.1214/aoap/1177005980.

14 Martin Dyer, Leslie Ann Goldberg, and Mark Jerrum. Matrix norms and rapid mixing for spin systems. *The Annals of Applied Probability*, 19(1):71–107, 2009. URL: http://www.jstor.org/stable/30243572.

15 David Eppstein and Daniel Frishberg. Rapid mixing of the hardcore Glauber dynamics and other Markov chains in bounded-treewidth graphs. *CoRR*, 2021. doi:10.48550/arXiv.2111.03898.

16 David Eppstein and Daniel Frishberg. Improved mixing for the convex polygon triangulation flip walk. *CoRR*, abs/2207.09972, 2022. doi:10.48550/arXiv.2207.09972.

17 F. Graham and P. Tetali. Isoperimetric inequalities for cartesian products of graphs. *Comb. Probab. Comput.*, 7:141–148, 1998.

18 Marc Heinrich. Glauber dynamics for colourings of chordal graphs and graphs of bounded treewidth, 2020. `arXiv:2010.16158`.

19 Peter J. Hilton and Jean J. Pedersen. Catalan numbers, their generalization, and their uses. *The Mathematical Intelligencer*, 13:64–75, 1991.

20 Mark Jerrum and Alistair Sinclair. Conductance and the rapid mixing property for Markov chains: The approximation of permanent resolved. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, 1988. `doi:10.1145/62212.62234`.

21 Mark Jerrum, Jung-Bae Son, Prasad Tetali, and Eric Vigoda. Elementary bounds on Poincaré and log-Sobolev constants for decomposable Markov chains. *The Annals of Applied Probability*, 14(4):1741–1765, 2004. URL: `http://www.jstor.org/stable/4140446`.

22 V. Kaibel and G. Ziegler. Counting lattice triangulations. *arXiv: Combinatorics*, 2002.

23 Volker Kaibel. On the expansion of graphs of 0/1-polytopes. In *The Sharpest Cut: The Impact of Manfred Padberg and His Work*, pages 199–216. SIAM, 2004.

24 Tali Kaufman and Izhar Oppenheim. High order random walks: Beyond spectral gap. *Combinatorica*, 40(2):245–281, 2020.

25 David A. Klarner. Correspondences between plane trees and binary sequences. *Journal of Combinatorial Theory*, 9(4):401–411, 1970. `doi:10.1016/S0021-9800(70)80093-X`.

26 Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. *arXiv e-prints*, 2020. `arXiv:2001.02827`.

27 David A Levin, Yuval Peres, and Elizabeth Wilmer. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

28 J. Loday. The multiple facets of the associahedron. In *Proc. 2005 Academy Coll. Series*, 2005.

29 Lisa McShine and P. Tetali. On the mixing time of the triangulation walk and other catalan structures. In *Randomization Methods in Algorithm Design*, 1997.

30 Milena Mihail and Umesh Vazirani. On the expansion of 0-1 polytopes. *Journal of Combinatorial Theory, Series B*, 1989.

31 Michael Molloy, Bruce Reed, and William Steiger. On the mixing rate of the triangulation walk. In *Randomization Methods in Algorithm Design*, 1997.

32 Dana Randall and Prasad Tetali. Analyzing glauber dynamics by comparison of markov chains. In *Latin American Symposium on Theoretical Informatics*, pages 292–304. Springer, 1998.

33 Alistair Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1(4):351–370, 1992. `doi:10.1017/S0963548300000390`.

34 Daniel D Sleator, Robert E Tarjan, and William P Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Journal of the American Mathematical Society*, 1(3):647–681, 1988.

35 Alexandre Stauffer. A Lyapunov function for Glauber dynamics on lattice triangulations. *Probability Theory and Related Fields*, 169:469–521, 2015.