COMBINATORIAL DIAMETERS AND AUTOMORPHISMS OF GELFAND-TSETLIN POLYTOPES

YIBO GAO, BEN KRAKOFF, AND LISA YANG

ABSTRACT. For a partition $\lambda = (\lambda_1, \dots, \lambda_n)$, one can construct a Gelfand-Tsetlin polytope GT_λ associated to λ. For all GT_λ, we give a formula for the diameter of the 1-skeleton and exactly describe the combinatorial automorphism group $\mathrm{Aut}(\mathrm{GT}_\lambda)$. Letting m be the number of distinct λ_i, we give an alternate proof of the formula in [GKT13] counting the number of vertices for $m \le 3$, and we describe a general approach for $m \ge 4$. In a special case, we can combinatorially describe the face poset, and use this to compute the f-polynomial. We conclude with some observations and a re-derivation of a known formula.

1. Introduction and Main Results

Gelfand-Tsetlin polytopes are constructed according to a partition λ and are of combinatorial interest for many reasons. Their integer points are in bijection with semi-standard Young Tableaux of shape λ, which have applications in representation theory. In this paper, we describe the diameter of the one skeleton of GT polytopes and the combinatorial automorphism group of GT polytopes. Our main theorems are stated below. We also examine some of the combinatorial properties of the GT polytopes including the enumeration of vertices and the special case when $\lambda = 12^{n-2}3$. First, we formally define GT polytopes and introduce some quick notation.

A partition λ of s is a sequence of weakly increasing positive integers $\lambda = (\lambda_1 \le \lambda_2 \le \lambda_3 \le \dots \le \lambda_n)$ such that $\sum_{i=1}^n \lambda_i = s$. We will often use multiplicative notation for λ and write $\lambda = (\lambda_1^{a_1}, \lambda_2^{a_2}, \dots, \lambda_m^{a_m})$ for $a_1, \dots, a_m \in \mathbb{Z}_{\ge 0}$ to denote a partition with a_1 copies of λ_1, a_2 copies of λ_2, and so forth. We may omit writing the term $\lambda_i^{a_i}$ if $a_i = 0$.

Definition 1.1 (GT Polytope)**.** Given a partition $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$, the Gelfand-Tsetlin Polytope GT_λ is the set of points $\vec{x} = (x_{i,j})_{1 \le j \le i \le n} \in \mathbb{R}^{n(n+1)/2}$ with $x_{i,i} = \lambda_i$ satisfying the following inequalities:

(1) $x_{i-1,j} \le x_{i,j} \le x_{i+1,j}$,
(2) $x_{i,j-1} \le x_{i,j} \le x_{i,j+1}$.

Suppose for some $i < j$ that $\lambda_i = \lambda_j$. Then for every $i' \ge i$ and $j' \le j$, we are forced to have $x_{i',j'} = \lambda_i$. Whenever such a situation occurs, we say that $x_{i',j'}$ is *fixed*. In general, GT_λ will be a polytope in $\mathbb{R}^d$ where d is at most $n(n-1)/2$.

Date: August 2, 2016.

These constraints can be visualized in a triangular array as shown in Figure 1. The polytope GT_λ corresponds to all sets of real numbers that can fill this triangular array with columns and rows weakly increasing.

$$\begin{array}{ccccccccc}
\lambda_1 & & & & & & & & \\
\rotatebox{90}{$\leq$} & & & & & & & & \\
x_{2,1} & \leq & \lambda_2 & & & & & & \\
\rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & & & & & & \\
x_{3,1} & \leq & x_{3,2} & \leq & \lambda_3 & & & & \\
\rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & & \rotatebox{90}{$\leq$} & & & & \\
x_{4,1} & \leq & x_{4,2} & \leq & x_{4,3} & \leq & \lambda_4 & & \\
\vdots & & \vdots & & \vdots & & \ddots & & \\
x_{n,1} & \leq & x_{n,2} & \leq & \dots & \leq & x_{n,n-1} & \leq & \lambda_n
\end{array}$$

FIGURE 1. Inequality constraints of GT polytopes.

We specify some notation and introduce two ways to model the faces of GT_λ: GT tilings and ladder diagrams.

Definition 1.2 (Notation)**.** We adopt the following conventions:

- n denotes the length of λ.
- m denotes the number of distinct values of λ.
- d denotes the dimension of GT_λ. It is easy to see $d = \binom{n}{2} - \sum_{i=1}^{m} \binom{a_i}{2}$.
- If $\lambda = (\lambda_1, \dots, \lambda_n)$, define its reversal $\lambda' := (\lambda_n, \dots, \lambda_1)$.
- $\mathcal{F}(\mathrm{GT}_\lambda)$ denotes the face poset of GT_λ ordered by inclusion.
- $\mathcal{I}_n = \{(i,j) : 1 \leq j \leq i \leq n\}$ denotes the triangular grid with shape shown in Figure 1.

For any polytope P, the 1-skeleton of P is a graph obtained by taking the vertices and edges of P. The diameter of a connected graph G is the minimum number of edges it takes to connect any two vertices. For GT polytopes, this combinatorial diameter is surprisingly low, and we have an exact formula for it's value given a partition λ.

Theorem (3.5)**.** *For any GT polytope GT_λ, diam$(GT_\lambda) = 2m - \delta_{1,a_1} - \delta_{1,a_m}$*

We also completely describe the combinatorial automorphism group of GT_λ, and discuss when such automorphisms are can be thought of as affine transformations or as purely combinatorial functions on the faces.

Theorem (4.18)**.** *Suppose $\lambda = (1^{a_1}, 2^{a_2})$ and $a_1, a_2 \geq 2$. If $a_1 = a_2 = 2$, then*

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2.$$

Otherwise,

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}}$$

Theorem (4.19)**.** *Suppose* $\lambda = 1^{a_1}\dots m^{a_m}$ *and* $m \geq 3$. *Let* $t = 1$ *if* $\lambda = \lambda'$ *and let* $t = 0$ *otherwise. Let* j *be the number of pairs* $a_k, a_{k+1} \geq 2$. *Then*

$$Aut(GT_\lambda) \cong \mathbb{Z}_2^t \ltimes_\varphi (S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{j+1}).$$

We conclude with some remarks on enumerative questions concerning these polytopes, such as counting vertices or computing f-vectors. We show that any GT polytope decomposes as a Minkowski sum of GT polytopes of part size 2, which themselves are notable because they are isomorphic to order polytopes of the product of two chains.

2. Preliminaries

In this section, we give two ways to combinatorially describe $\mathcal{F}(\mathrm{GT}_\lambda)$, both of which have distinct advantages in describing the facial structure of GT_λ.

2.1. GT Tilings.

It will be useful to think of the triangular array in Figure 1 as a triangular grid $\mathcal{I}_n$, whose squares are partitioned into tiles.

Definition 2.1 (GT Tilings, Definition 4.11 in [McA06])**.** Given a partition λ of length n and the triangular grid $\mathcal{I}_n = \{(i,j) : 1 \leq j \leq i \leq n\}$, a *GT tiling* is a partition $\mathcal{T}$ of $\mathcal{I}_n$ into disjoint nonempty sets, called *tiles*, such that each tile T is connected (squares are adjacent if they share an edge) and that if $(i_1, j_1), (i_2, j_2) \in T$, then $(i,j) \in T$ if $i_1 \leq i \leq i_2$, $j_1 \leq j \leq j_2$ and $i \geq j$.

These tilings form a poset ordered by refinement i.e. $\mathcal{T}_1 \leq \mathcal{T}_2$ if $\mathcal{T}_2$ is a refinement of $\mathcal{T}_1$. Note that this poset is also graded by number of tiles.

Theorem 2.2 (Theorem 4.14 in [McA06])**.** *Let* $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$. *Let* T_λ *denote the set of tilings such that* (i,i) *and* (j,j) *are in the same tile if and only if* $\lambda_i = \lambda_j$. *Then* $\mathcal{F}(GT_\lambda) \cong T_\lambda$.

Proof. The isomorphism between the two is given by taking a valid point in GT_λ, drawn in the diagram depicted in Definition 1.1, and placing the diagram to fit inside $\mathcal{I}_n$. Then, adjacent boxes with equal entries are placed into the same tile. Figure 2 shows an example of such transformation. For more details, see [McA06].

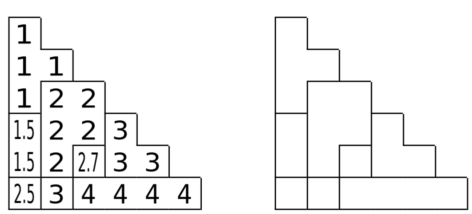

Figure 2. Left: example of mapping a GT pattern to a GT tiling. Right: the GT tiling.

□

Remark 2.3. Given a point $x \in \mathrm{GT}_\lambda$, this theorem gives us a way to determine the dimension of the minimal face containing x. Namely, map x to it's tiling, and the dimension of this minimal face will be the number of tiles - m. One way to see this is that T_λ is graded by number of tiles, and the minimal number of tiles is m.

2.2. **Ladder Diagrams.** For every λ, we define a graph Γ_λ such that faces of GT_λ correspond to subgraphs of Γ_λ with certain restrictions. These subgraphs were introduced in [ACK16] as ladder diagrams and will be helpful combinatorial models for the faces of GT_λ.

Let Q be the infinite graph corresponding to first quadrant of the Cartesian plane, i.e. let Q have vertices (i,j) for all $i,j \geq 0$ and edges $((i,j),(i+1,j))$ and $((i,j),(i,j+1))$. For convenience, define $a_0 = 0$ and $s_j = \sum_{i=0}^{j} a_i$ for $0 \leq j \leq m$.

Definition 2.4 (Ladder Diagrams)**.** For $\lambda = (1^{a_1}, \ldots, m^{a_m})$, the grid Γ_λ is an induced subgraph of Q constructed as follows. Let the *origin* be the vertex $(0,0)$. Define *terminal vertices* $t_j = (s_j, n - s_j)$ for $0 \leq j \leq m$. Γ_λ consists of all vertices and edges appearing on any North-East path between the origin and a terminal vertex.

A *ladder diagram* is a subgraph of Γ_λ such that

(1) the origin is connected to every terminal vertex by some North-East path.
(2) every edge in the graph is on a North-East path from the origin to some terminal vertex.

An example of the grid Γ_λ and some of its ladder diagrams are shown in Figure 3. All fixed entries are shaded. The terminal vertices lie along the *main diagonal* of Γ_λ.

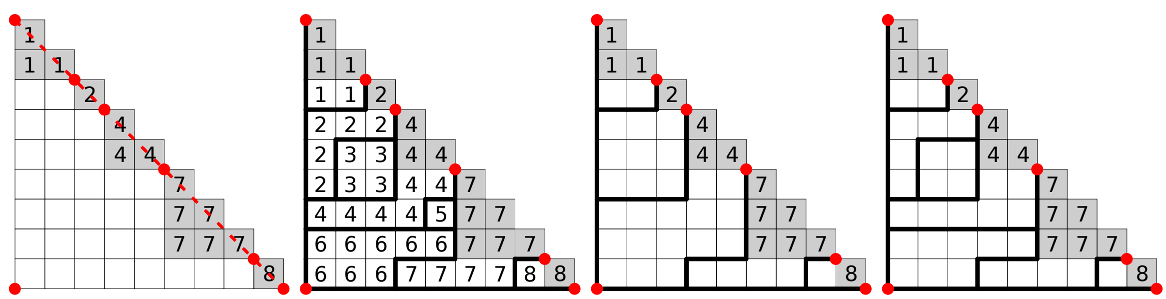

FIGURE 3. Let $\lambda = (1^2, 2^1, 4^2, 7^3, 8^1)$. From left to right: Γ_λ with origin and terminal vertices in red and a dashed line indicating the main diagonal, ladder diagram for a point in GT_λ, ladder diagram for a 0-dimensional face (vertex), and ladder diagram for a 2-dimensional face.

Definition 2.5 (Face Lattice of Ladder Diagrams)**.** The face lattice of Γ_λ, denoted as $\mathcal{F}(\Gamma_\lambda)$, is the set of all ladder diagrams ordered by inclusion.

Note that $\mathcal{F}(\Gamma_\lambda)$ is graded by number of bounded regions.

Theorem 2.6 (Theorem 1.9 in [ACK16])**.** $\mathcal{F}(GT_\lambda) \cong \mathcal{F}(\Gamma_\lambda)$.

Proof. An isomorphism is given by taking a point in GT_λ and drawing lines around adjacent groups of $x_{i,j}$ with equal value will produce a face of Γ_a. For more details, see [ACK16]. ☐

As with GT-tilings, given a point $x \in \mathrm{GT}_\lambda$, we can determine the dimension of the minimal face containing x. Map x to it's corresponding ladder diagram, and the number of bounded regions will be the dimension of this minimal face. Again, to see this note that the minimal elements of $\mathcal{F}(\Gamma_\lambda)$ are trees with 0 bounded regions.

Remark 2.7. Note that the posets T_λ and $\mathcal{F}(\Gamma_\lambda)$ only depend on the multiplicities a_i and not on the values of λ_i themselves. So when examining the purely combinatorial properties of GT_λ, it suffices to consider $\lambda = (1^{a_1}, 2^{a_2}, \dots, m^{a_m})$

3. Combinatorial Diameters

In this section, we give an exact formula for the diameter of the 1-skeleton of GT_λ, denoted $\mathrm{diam}(\mathrm{GT}_\lambda)$. We use diameter to refer to the smallest number of edges required to connect any two vertices. As explained in Remark 2.7, it suffices to consider $\lambda = (1^{a_1}, \dots, m^{a_m})$ where $a_1, \dots, a_m \in \mathbb{Z}_{>0}$.

In order to study the diameters of the 1-skeleton of GT_λ, we need to first understand what a vertex is in terms of ladder diagrams and under conditions two vertices are connected.

Definition 3.1. We say that two paths in a ladder diagram from the origin to terminal vertices are *noncrossing* if they do not meet again after their first separation.

In particular, vertices of GT_λ have lattice diagrams consisting of $m+1$ noncrossing paths. However, two of these paths go directly up and directly right from the origin, and are present in every ladder diagram so we usually ignore them. Two vertices are connected by an edge if the union of their ladder diagrams has exactly one bounded region.

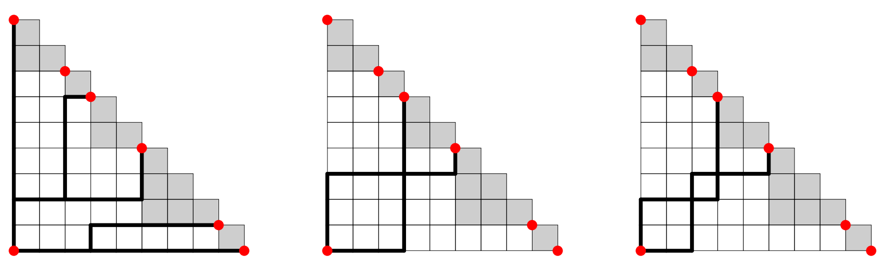

FIGURE 4. Left: 3 non-crossing paths. Middle and Right: 2 Crossing paths

Lemma 3.2. *Any two vertices v and w of GT_λ are separated by at most $2m-2-\delta_{1,a_1}-\delta_{1,a_m}$ edges.*

Proof. We give an algorithm to find a path between v and w of length at most $2m-\delta_{1,a_1}-\delta_{1,a_m}$. Assume that in the ladder diagram representation, vertex v corresponds to noncrossing paths

$v_1, \ldots, v_{m-1}$ where v_j connects the origin $(0,0)$ to terminal vertex $t_j = (s_j, n - s_j)$ where $s_j = \sum_{i=0}^{j} a_i$. Similarly denote the noncrossing paths corresponding to w as $w_1, \ldots, w_{m-1}$.

Essentially, we want to change v_1 to w_1, v_2 to w_2, ..., v_{m-1} to w_{m_1} and making sure that the $m-1$ paths we have are always noncrossing, and the common refinement before and after changing some paths has exactly one bounded region.

Phase 1: If $a_1 = 1$, then v_1, w_1 are paths that go from $(0,0)$ to $(1, n-1)$. Therefore, there exists a unique index r_v such that path v_1 passes through both $(0, r_v)$ and $(1, r_v)$. In other words, r_v is the vertical index for v_1 to go from column 0 to column 1. Similarly we can define r_w. WLOG, assume that $r_v \geq r_w$. Because of this inequality, we know that path w_1 is contained inside of v_1 and therefore, the ladder diagram consisted of $w_1, v_1, v_2, \ldots, v_{m-1}$ have exactly 1 bounded region and it is thus an edge e of the GT polytope containing v. Let $v' = (v'_1, \ldots, v'_{m-1})$ be the other side of this edge (one side is v). Notice that the inner edge $\big((0, r_w), (1, r_w)\big)$ is in the ladder diagram of e but not in the ladder diagram of v'. Therefore, $\big((0, r_w), (1, r_w)\big)$ must be in the ladder diagram of v' and it means $v'_1 = w_1$ since $a_1 = 1$. Similarly, we can use one move to make v_{m-1} and w_{m-1} equal if $a_m = 1$.

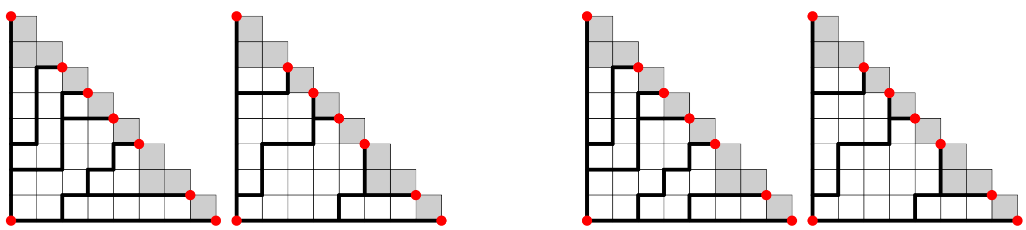

FIGURE 5. Phase 1 of the algorithm for Lemma 3.2. From left to right: the ladder diagram for v, w, v', $w'(= w)$.

Phase 2: Now we describe an algorithm that take v' to some vertex u in at most $m-1-\delta_{1,a_1}-\delta_{1,a_m}$ steps. The algorithm works as follows: for each $i = 1 + \delta_{1,a_1}, \ldots, m-1-\delta_{1,a_m}$, change path v_i so that it starts at terminal vertex t_i, goes horizontally to the left until it meets and merges with path v_{i-1}. First, the ladder diagram after this change is clearly a vertex. Also, if we take the common refinement of the two ladder diagrams before and after the change, or equivalently, start with the old ladder diagram and add a new path v'_i described above, then this new path simply cuts the tile bounded by v_{i-1} and v_i into two parts and thus there exists an edge between these two vertices. Figure 6 shows an example of this algorithm.

Similarly, we can apply the same algorithm to w' to get to the same vertex u in the same number of steps. And finally, the total number of steps is at most

$$(\delta_{1,a_1} + \delta_{1,a_m}) + 2(m - 1 - \delta_{1,a_1} - \delta_{1,a_m}) = 2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}.$$

□

Lemma 3.3. *There exist two vertices separated by at least $2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}$ edges.*

In order to prove this lemma, we need to first construct vertices in terms of ladder diagrams.

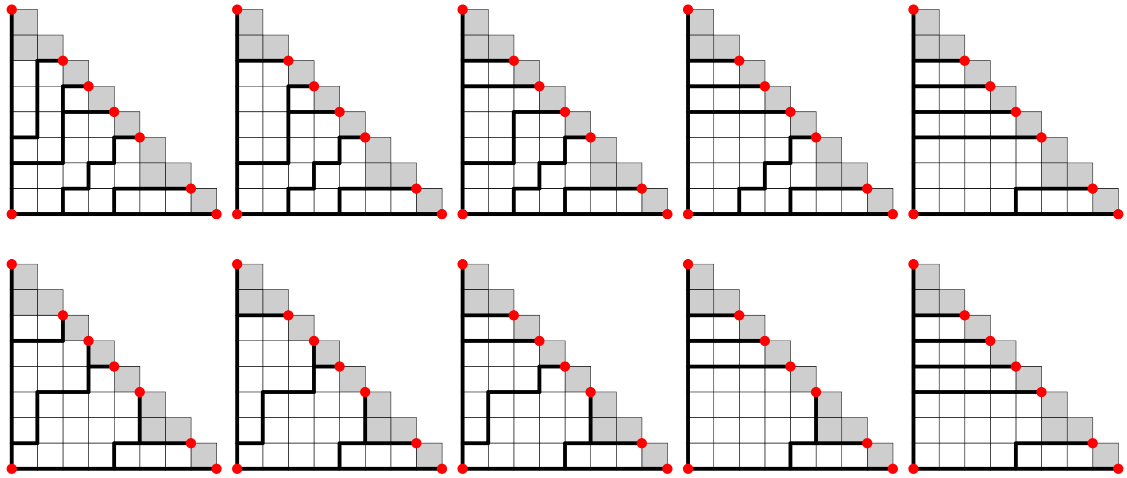

FIGURE 6. Phase 2 of the algorithm for Lemma 3.2. The first line shows steps that move from v' to u; the second line shows steps that move from w' to u.

Definition 3.4 (Zigzag lattice path)**.** We construct two vertices z_h and z_v that will be used in the proof for Lemma 3.3. Let $\lambda = (1^{a_1}, \dots, m^{a_m})$. If $a_1 > 1$, meaning that $(1, n-1)$ is not a terminal vertex, we call $(1, n-1)$ a *virtual terminal vertex*. Similarly, if $(n-1, 1)$ is not an actual terminal vertex, meaning that $a_m > 1$, we call it a virtual terminal vertex.

We will consider the ladder diagram for a vertex of GT_λ as $m-1$ southwest lattice paths from terminal vertices to the origin. For $j = 1, \dots, m-1$, define a horizontal zigzag path, h_j, to be the path that starts at terminal vertex t_j, goes horizontally left until reaching a column where there exists a terminal vertex or a virtual terminal vertex on it, then goes vertically down until reaching a row where there exists a terminal vertex or a virtual terminal vertex, and so on and so forth until the paths reaches column 0 or row 0 so that it will then go to the origin in a unique way. Similarly define a vertical zigzag path, v_j, with the only difference that it will start vertically instead of horizontally. Finally, let z_h be the vertex of GT_λ represented by the ladder diagram $(h_1, \dots, h_{m-1})$ and let z_v be the vertex of GT_λ represented by $(v_1, \dots, v_{m-1})$. Figure 7 shows the construction of an example.

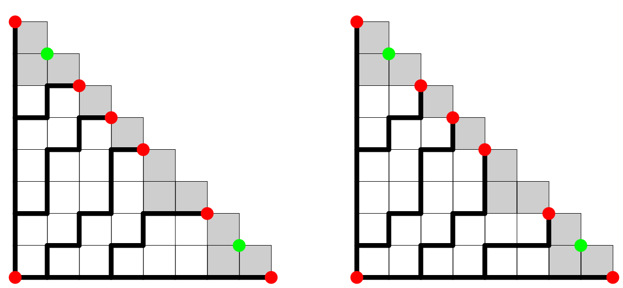

FIGURE 7. Vertices z_h (left) and z_v (right) of GT_λ with $\lambda = (1^2, 2^1, 3^1, 4^2, 5^2)$. Virtual terminal vertices are labeled as green dots.

Proof of Lemma 3.3. We will first consider the case where $a_1, a_m \geq 2$ so that the idea of the proof can be shown clearly. Afterward, we will deal with the details coming from either of them being 1.

Consider the vertices given in Definition 3.4. Assume that there is a sequence of vertices $z_h = y_0, y_1, \ldots, y_\ell = z_v$ in GT_λ such that y_k and y_{k+1} are connected by an edge. Since we can uniquely represent each vertex in this sequence as a union of $m-1$ lattice paths from the original to each terminal vertex, each step $y_k \to y_{k+1}$ can be thought of as simply changing these paths, such that the union of y_k and y_{k+1} has at most 1 bounded region. For convenience of notation, for each $i = 1, \ldots, m-1$, let p_i be the path that goes from the origin to the terminal vertex i generically for all y_k's. For $k = 1, \ldots, \ell$, define $X_k = \{i : p_i \text{ changes as we go from } y_{k-1} \text{ to } y_k\}$. Since in each step we can have at most 1 bounded region between y_{k-1} and y_k, it is clear that X_k is of the form $\{i, i+1, \ldots, j\}$ for some $i \le j$.

We will now show that $X_1, \ldots, X_\ell$ must have the following conditions:

(1) For each $s \in [m-1]$, s appears in at least two of the sets X_i's. In other words, for each $s \in [m-1]$, there exists $1 \le i < j \le \ell$ such that $s \in X_i$, $s \in X_j$.
(2) For each $s \neq s' \in [m-1]$, the last time s appears in any of the sets (which is the unique index b_s such that $s \in X_{b_s}$ and $k \notin X_i$ for all $i > b_k$) is different from the last time that k' appears.
(3) If $X_k = \{i, i+1, \ldots, j\}$, then at least $j-i$ of $i, i+1, \ldots, j$ must appear in some $X_{k'}$ for $k' < k$
(4) If $X_k = \{i, i+1, \ldots, j\}$ and it is the last time that s appears, then each one of $\{i, i+1, \ldots, j\} \backslash \{s\}$ must appear in (possibly different) $X_{k'}$ for $k' < k$.

Notice that if the last time s appears is in some X_k, it means that p_s has already become v_s (Definition 3.4) in y_k and stays the same for the rest of the steps. We will then explain these conditions in details one by one.

Condition (1). If for some $s \in [m-1]$, it appears in only one set X_k, then it means when we go from vertex y_{k-1} to y_k, the path p_s is changed from v_s to h_s in exactly one step. But by construction, superimposing h_s and v_s will create at least 2 bounded regions, instead of one. Therefore, y_{k-1} and y_k are not connected by an edge, a contradiction. Therefore, each $s \in [m-1]$ appears in at least two sets.

Condition (2). If X_k is the last time that both s and s' appears, then it means that from y_{k-1} to y_k, paths $p_{s'}$ and p_s are changed to $v_{s'}$ and v_s simultaneously. However, $v_{s'}$ and v_s do not have any intersection in the interior of $\mathcal{I}_n$. So if we want to go back from y_k to y_{k-1}, we have to change v_s' and v_s simultaneously, which will create two bounded regions when we superimpose y_{k-1} and y_k, which is a contradiction.

Condition (3). Since initially, h_i and h_j do not have any intersection in the interior of $\mathcal{I}_n$, in order to change multiple paths at the same time, we need to merge these paths first. Specifically, if we want to change paths $i, i+1, \ldots, j$ simultaneously, we need to modify at least all but one of these paths to join them together.

Condition (4). This condition is crucial and it justifies our choices for z_h and z_v. Assume that $X_k = \{i, i+1, \ldots, j\}$ is the last time that s appears, meaning that when we go from vertex y_{k-1} to y_k, we change p_s to v_s. If there exists $s' \neq s \in X_k$, such that path s' hasn't changed before, meaning that $p_{s'} = h_{s'}$. Notice that since we change paths $i, i+1, \ldots, j$ simultaneously, all of these paths before and after the change will have an interior vertex in $\mathcal{I}_n$ in common. Therefore, v_s and $h_{s'}$ must have a common interior vertex. Since $s \neq s'$, we must have $s' = s+1$. As we have assumed $a_1, a_m \geq 2$, superimposing the ladder diagram of y_{k-1} and y_k will create at least 2 bounded regions, because of the definition of v_s and h_{s+1}. Therefore, we have a contradiction and thus all $s' \neq s$ must have already appeared at least once.

As we have proved all these conditions, we will jump out of the general setting of GT polytopes and look at any sequence of sets $X_1, \ldots, X_\ell$ that satisfies all these four conditions. We will show that any such sequence will have length $\ell \geq 2m-2$.

To do this, for $i = \ell, \ldots, 1$, we look at the first set X_i that is not a singleton. Say that it is X_k. If X_k is the last time that some $s \in [m-1]$ appears, then we claim that changing X_k to $\{s\}$ will still satisfy all four conditions. According to Condition (4), for each $s' \neq s$ that is in X_k, s' must have appeared before. According to Condition (2), for each $s' \neq s$, this is not the last time that s' appears so s' will appear sometime later. Therefore, condition (1) still holds after changing X_k to $\{s\}$. Condition (2) also holds because this change does not modify the indices of the sets where each $s' \in [m-1]$ appears last. Condition (3) and (4) hold trivially because we have less non-singleton sets to worry about. Another case is that X_k is not the last time that any of the path appears last. According to condition (3), there exists $s \in X_k$ such that each one of $X_k \backslash \{s\}$ has appeared before. Similarly, we claim that all these four conditions will hold after changing X_k to $\{s\}$. For each one of $s' \in X_k \backslash \{s\}$, as it appears before and X_k is not the last time that it appears, we know that s' will appear at least twice even after this change. The number of appearance of s does not change. Therefore, condition (1) is satisfied. Condition (2) holds because each one of $s' \in X_k$ will appear sometime later. Condition (3) and (4) hold trivially because similarly we have less non-singleton sets to worry about.

Continuing this procedure inductively, we will eventually end up with a sequence of sets $Y_1, \ldots, Y_\ell$ where each one is a singleton. As Condition (1) still holds, we conclude that $\ell \geq 2m-2$ as desired.

Now we consider the case where a_1 and a_m may be 1. We claim that if $a_1 = 1$, then $\{1\}$ must appear as a singleton as one of the term in $X_1, \ldots, X_\ell$ and if $a_m = 1$, then $\{m-1\}$ must appear as a singleton as one of the term in the sequence. Notice the reasons are slightly different.

If $a_1 = 1$, path h_1 has only 1 interior edge, namely $(0, n-1) - (1, n-1)$. In order for it to merge with other paths, it must change on its own first, meaning that $\{1\}$ will appear, because h_1 and $p_2, \ldots, p_{m-1}$ will never have any interior intersection. If it never merges with other paths, then it must appear as $\{1\}$ at some point so that p_1 can be actually changed from h_1 to v_1. If $a_m = 1$, at some point, path p_{m-1} must be changed to v_{m-1}, which has only 1 interior edge. This change is recorded as $\{m-1\}$ because $p_1, \ldots, p_{m-2}$ cannot have any interior intersection with v_{m-1}.

Therefore, if $a_1 = 1$ (or/and $a_m = 1$), in the sequence $X_1, \ldots, X_\ell$, we can take out the singleton $\{1\}$ (or/and $\{m-1\}$) and delete all other 1's (or/and $m-1$'s) in the sets. The remaining sequence will satisfy the four conditions mentioned above. We have taken out at least $\delta_{1,a_1} + \delta_{1,a_m}$ singletons of $\{1\}$ and $\{m-1\}$ and there are $m - 1 - \delta_{1,a_1} - \delta_{1,a_m}$ paths remaining. So by the same argument, the total length of the sequence

$$\ell \geq (\delta_{1,a_1} + \delta_{1,a_m}) + 2(m - 1 - \delta_{1,a_1} - \delta_{1,a_m}) = 2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}$$

as desired. □

Theorem 3.5. *diam(GT_λ)* $= 2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}$.

Proof. By combining the upper and lower bounds in Lemma 3.2 and Lemma 3.3, we are done. □

4. Combinatorial Automorphisms

In this section, we completely describe the combinatorial automorphisms of GT_λ, denoted Aut(GT_λ). A combinatorial automorphism of GT_λ is a permutation of its vertices that preserves $\mathcal{F}(\mathrm{GT}_\lambda)$, and we will refer to combinatorial automorphisms simply as automorphisms. In Section 4.1 we exhibit generators for the automorphism group and use relations between them to describe the group structure. In Section 4.2 we develop some necessary background and use it in Section 4.3 to show that these generators in fact generate the entire automorphism group. As explained in Remark 2.7, it generally suffices to consider $\lambda = (1^{a_1}, \ldots, m^{a_m})$ where $a_1, \ldots, a_m \in \mathbb{Z}_{>0}$, but we make note of special properties arising from particular values of λ_i.

4.1. Symmetries.

We begin by identifying the generators of Aut(GT_λ). By Theorem 2.6, it suffices to show that these maps are automorphisms of $\mathcal{F}(\Gamma_\lambda)$ to show that they are automorphisms of GT_λ. We also examine these maps beyond their purely combinatorial properties.

Proposition 4.1 (The Corner Symmetry). *For any λ, there is a $\mathbb{Z}_2$ automorphism μ on $\mathcal{F}(\Gamma_\lambda)$ given by swapping two pairs of edges in any positive path leaving $(0,0)$: $((0,0),(1,0))$ is swapped with $((0,0),(0,1))$ and $((1,0),(1,1))$ with $((0,1),(1,1))$.*

Proof. Consider the linear map $\mu : \mathbb{R}^{\frac{n(n+1)}{2}} \to \mathbb{R}^{\frac{n(n+1)}{2}}$ that acts as the identity on all the $x_{i,j}$, except maps $x_{n,1} \mapsto x_{n,2} + x_{n-1,1} - x_{n,1}$. Note that $x_{n-1,1} \leq x_{n-1,1} + x_{n,2} - x_{n-1} \leq x_{n,2}$, and $\mu^2 = Id$, and so $\mu(GT_\lambda) = GT_\lambda$. This linear map thus induces a combinatorial $\mathbb{Z}_2$ automorphism, which we will abuse notation and call μ. Checking the cases when the above inequalities are equalities shows that μ acts on ladder diagrams in the described fashion. □

As demonstrated in the proof of the previous proposition, this combinatorial automorphism μ is always affine (in fact, linear) regardless of λ.

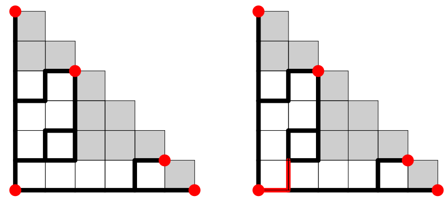

Figure 8. Action of μ on a ladder diagram

Proposition 4.2 (The k-Corner Symmetry)**.** *Label the k^{th} terminal vertex $t_k = (n-i,i)$, and suppose that $a_k, a_{k+1} \geq 2$. There is a $\mathbb{Z}_2$ automorphism μ_k on $\mathcal{F}(\Gamma_\lambda)$ given by swapping two pairs of edges, $((n-i,i)(n-i,i-1))$ with $((n-i,i)(i-1,i)$ and $((n-i,i-1),(n-i-1,i-1))$ with $((n-i-1,i),(n-i-1,i-1))$ in any positive path going to t_k.*

Proof. Let $x_{i',j'}$ be the coordinate in 1 which is immediately adjacent to the values λ_k and λ_{k+1}. Consider the linear map $\mu_k : \mathbb{R}^{\frac{n(n+1)}{2}} \to \mathbb{R}^{\frac{n(n+1)}{2}}$ which acts as the identity on the $x_{i,j}$, except $x_{i',j'} \mapsto x_{i',j'-1} + x_{i'+1,j'} - x_{i',j'}$. Since $a_k, a_{k+1} \geq 2$, we have $\lambda_k \leq x_{i',j'-1} \leq x_{i',j'-1} + x_{i'+1,j'} - x_{i',j'} \leq x_{i'+1,j'} \leq \lambda_{k+1}$, so all 4 constraints on the coordinate $x_{i',j'}$ are satisfied. Noting that $\mu_k^2 = Id$, we conclude that $\mu_k(GT_\lambda) = GT_\lambda$. This linear transformation induces a combinatorial automorphism on GT_λ, which we again refer to as μ_k. Examining when the above inequalities are equalities shows μ_k has the required description in terms of ladder diagrams. □

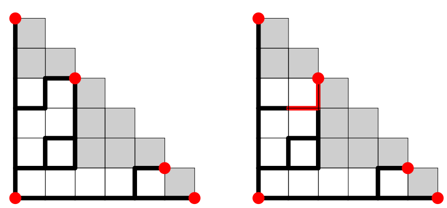

Figure 9. Action of μ_k

This map is very similiar to the Corner Symmetry, only occuring at the k^{th} terminal vertex instead of at the origin. It also shares the property that it can always be thought of as a linear map, regardless of the values of λ.

Proposition 4.3 (Symmetric Group Symmetry)**.** *Suppose that $a_1 = 1$. Then $S_{a_2} \subset Aut(GT_\lambda)$. Similiarly if $a_m = 1$, $S_{a_{m-1}} \subset Aut(GT_\lambda)$.*

Proof. It suffices to describe the action of S_{a_2} on ladder diagrams. Take the first column of possible horizontal edges, and label the top a_2 edges 1 though a_2. S_{a_2} then acts by if $\sigma(i) = j$, any path with an edge corresponding to i is mapped to a path with an edge corresponding to j. This map preserves inclusions, and since only the first a_2 edges are permuted, the number of bounded regions is also preserved. □

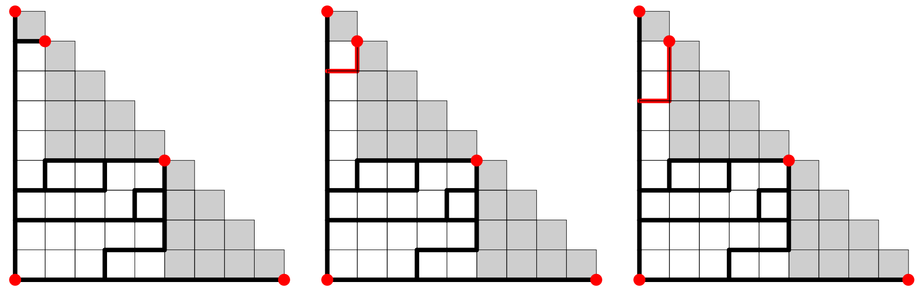

FIGURE 10. Action of $(123) \in S_4$ on a ladder diagram. Note that if (15) were applied, there would be an extra bounded region

Proposition 4.4 (The Flip Symmetry)**.** *Suppose that* $\lambda = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m}) = (1^{a_m}, 2^{a_{m-1}}, \ldots, m^{a_1})$. *Then there is a* $\mathbb{Z}_2$ *automorphism* ρ *of* GT_λ *which reflects a ladder diagram over the line* $y = x$.

Proof. Since $a_i = a_{m-j+1}$, the terminal vertices of any Γ_λ are symmetric about the $y = x$ axis, so reflection about this line gives another ladder diagram with the same number of bounded regions. □

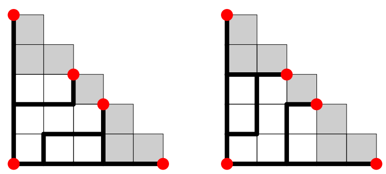

FIGURE 11. Action of ρ on ladder diagram.

Whenever λ satisfies the required condidition in Prop. 4.4 we say that λ is *reverse symmetric*. When $\lambda = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m})$, this symmetry is actually affine– it can be realized as the map $f(x) = -Px + m \cdot \mathbf{1}$, where P is a permutation matrix, and $\mathbf{1}$ is the all 1's vector. However, for λ of a different form, this is no longer true, even in small cases. For example when $\lambda = (1, 2, 4)$, a straightforward computations shows that this symmetry is no longer affine.

Proposition 4.5 (The $m = 2$ Rotation Symmetry)**.** *Suppose that* $m = 2$. *Note that any ladder diagram only has 3 terminal vertices, two on the the* x *or* y *axis and one not on the axes, call it* v. *There is a* $\mathbb{Z}_2$ *automorphism* τ *on* $\mathcal{F}(GT_\lambda)$ *taking paths from* $(0,0)$ *to* v *and rotating them* $180°$ *so that they are paths from* v *to* $(0,0)$.

Proof. Clear from the description of τ. □

Proposition 4.6 (The $m = 2$ Vertex Symmetry)**.** *When* $m = 2$, *there are two paths* p_1 *and* p_2 *to the terminal vertex* t_1 *that turn exactly once. The map* α *sending these two paths to each other is a* $\mathbb{Z}_2$ *automorphism of* GT_λ.

Proof. Suppose without loss of generality that a ladder diagram contains p_1 but not p_2. Note removing p_1 will unbound exactly 1 region, the south-eastmost bounded region. Adding in

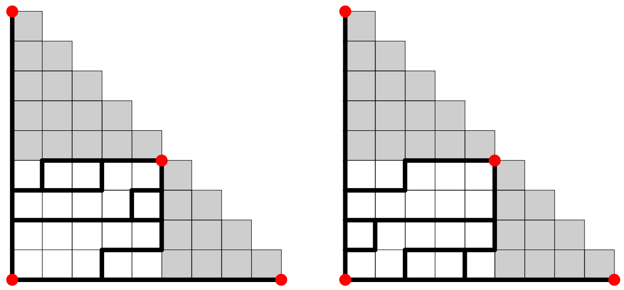

FIGURE 12. Action of τ on a ladder diagram

p_2 will create exactly one bounded region in the north-west corner of the diagram, so α preserves number of bounded regions. ☐

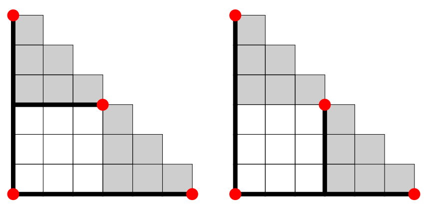

FIGURE 13. Paths p_1 and p_2

Note that these two paths themselves correspond to vertices in our polytope that are connected to every other vertex by an edge. This symmetry corresponds to exchanging these two vertices.

It is clear that the group formed by these possible generators is contained in $\mathrm{Aut}(\mathrm{GT}_\lambda)$. Note that if $m = 1$, the polytope is a single point with only the trivial automorphism. If $m = 2$ and either $\lambda = (1, 2^{a_2})$ or $\lambda = (1^{a_1}, 2)$, then the polytope is a simplex and its automorphism group is the symmetric group. In all other cases, either $m = 2$ with $a_1, a_2 \geq 2$, or $m \geq 3$. For both cases, we describe the subgroup generated by these symmetries, and postpone the proof that this subgroup is in fact $\mathrm{Aut}(GT_\lambda)$ to Section 4.3.

Theorem 4.7 (m=2 Automorphisms)**.** *Suppose $\lambda = (1^{a_1}, 2^{a_2})$ and $a_1, a_2 \geq 2$.*

If $a_1 = a_2 = 2$, then

$$D_4 \times \mathbb{Z}_2 \subseteq Aut(GT_\lambda).$$

Otherwise,

$$D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}} \times \subseteq Aut(GT_\lambda).$$

Proof. Suppose $a_1 \neq a_2$. Collecting the generators applicable in this case from previous propositions, we have the subgroup of automorphisms $\langle \mu, \mu_1, \tau, \alpha \rangle$, and note these generators satisfy the following relations $\mu^2 = \mu_1^2 = \tau^2 = \alpha^2 = 1, \mu\tau = \tau\mu_1$ and all other variables commute. It is not hard to recognize the subgroup $\langle \tau\mu, \mu \rangle$ as D_4. Adding in the generator α, will generate the entire group, and α commutes with all generators, so the resulting group is

isomorphic to $D_4 \times \mathbb{Z}_2$.
Assuming $a_1 = a_2 \geq 4$, we have the subgroup described in the previous case $a_1 \neq a_2$, but with the additional generator ρ. Note ρ commutes with all of these generators from the previous case, so the resulting group is $D_4 \times \mathbb{Z}_2^2$. When $a_1 = a_2 = 2$, note that $\rho = \tau$, resulting in a smaller subgroup. □

Theorem 4.8 ($m \geq 3$ Automorphisms)**.** *Suppose $\lambda = 1^{a_1} \dots m^{a_m}$ and $m \geq 3$. Let $t = 1$ if λ is reverse-symmetric and let $t = 0$ otherwise. Let j be the number of pairs $a_k, a_{k+1} \geq 2$. Then*

$$\mathbb{Z}_2^t \ltimes_\varphi (S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{j+1}) \subseteq Aut(GT_\lambda).$$

Where $\varphi(1)$ acts on $S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{j+1}$ via
$\varphi(1)(\sigma_1, \sigma_2, z_1, \dots, z_j, z_{j+1}) = (\sigma_2, \sigma_1, z_j, z_{j-1} \dots, z_2, z_1, z_{j+1})$

Proof. From the previous propositions, we have as possible generators $\mu, \mu_1, \dots, \mu_1, S_{a_2}, S_{a_{m-1}}$, with $S_{a_2}, S_{a_{m-1}}$ possibly omitted. However, whichever symmetries are present commute, as they permute disjoint sets of edges in the ladder diagram.

If λ is reverse symmetric, we also have the generator ρ. Note that since ρ flips every ladder diagram about $y = x$, ρ satisfies the following commutation relations: $\rho\mu = \mu\rho$, $\rho\mu_i = \mu_{l-i}\rho$, and for $\sigma_1 \in S_{a_2}$, $\sigma_2 \in S_{a_{m-1}}$, where σ_1 and σ_2 have the same cycle notation, $\rho\sigma_1 = \sigma_2\rho$. These relations are enough to give a subgroup of the stated form. □

In the following sections, we will show that the groups described in Theorems 4.7 and 4.8 are in fact all of $\mathrm{Aut}(\mathrm{GT}_\lambda)$. We will do this by showing the size of $\mathrm{Aut}(GT_\lambda)$ agrees with the size of these subgroups. Specifically, we will bound the size of $\mathrm{Aut}(GT_\lambda)$ by closely examining the action of this group on the facet structure of GT_λ and applying the Orbit-Stabilizer theorem. In order to do this, we need a few ways to classify and partition facets in GT_λ.

4.2. **Chains of Facets.** In this section, we study the facet structure of GT_λ in preparation for Section 4.3. The facet's of GT_λ will be in bijection with certain edges of ladder diagrams, which will allow us to analyze the action of $\mathrm{Aut}(GT_\lambda)$ on these facets, motivated by the following lemma.

Lemma 4.9. *An automorphism of GT_λ is completely determined by where it sends the facets of GT_λ.*

Proof. This follows from the fact that for a general polytope P, every face of P can be written as an intersection of the facets of P □

In particular, if an automorphism of GT_λ fixes every facet, it must be the identity.

Definition 4.10. We define the *interior edges* of Γ_λ to be the edges lying inside Γ_λ and any edges of the form $((s_j, n - s_{j+1}), (s_j, n - s_{j+1} + 1))$ and $((s_j, n - s_{j+1}), (s_j + 1, n - s_{j+1}))$. All other edges of Γ_λ are considered *boundary edges*.

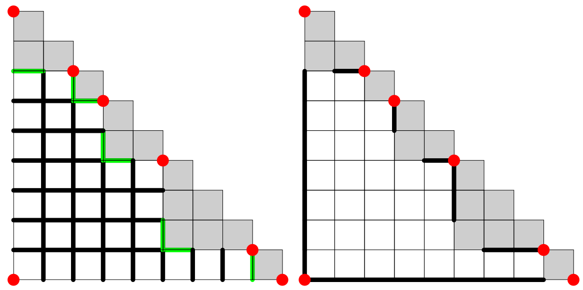

FIGURE 14. Left: interior edges of Γ_λ. Interior edges of the form $((s_j, n - s_{j+1}), (s_j, n - s_{j+1} + 1))$ and $((s_j, n - s_{j+1}), (s_j + 1, n - s_{j+1}))$ are in green. Right: boundary edges of Γ_λ.

Observe that interior edges are exactly the edges that can be erased from Γ_λ and still have a valid ladder diagram.

Proposition 4.11. *The facets of GT_λ are in bijection with the interior edges of Γ_λ.*

Proof. Since ladder diagrams in $\mathcal{F}(\Gamma_\lambda)$ are graded by number of bounded regions, any facet will correspond to a ladder diagram with all possible edges except one, which will be an interior edge as defined in Definition 4.10. The bijection is given by mapping a facet F to the edge not contained in $\Gamma_\lambda(F)$. □

This proposition gives us an easy way to think about the facets of GT_λ, so for most arguments we will simply refer to a facet by it's corresponding interior edge.

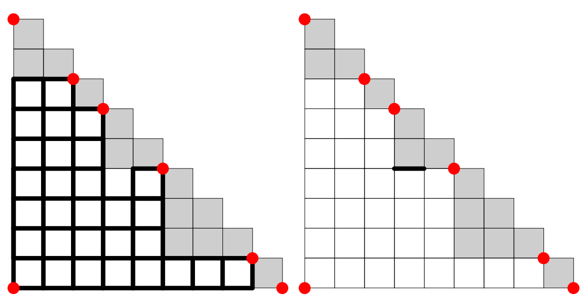

FIGURE 15. Left: ladder diagram of a facet. Right: complement of ladder diagram.

Definition 4.12. Two facets are called *dependent* if their intersection is a $d-3$ dimensional face.

There is a way to visualize the intersection of two facets. Remove both of their corresponding edges from Γ_λ, and then remove any necessary edges to create a ladder diagram. In most cases, no edges will require removal, meaning the intersection of the two facets is $d-2$ dimensional. However, if the removed edges are arranged fully as in Figure 16 and fully in the interior of Γ_λ, extra edges will require removal, implying the intersection is $d-3$ dimensional.

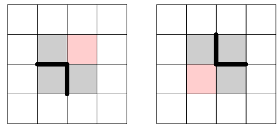

FIGURE 16. The gray boxes indicate entries $x_{i,j}$ that are equal on each facet. The red box indicates the entry forced to be equal to the other three.

Facets being arranged as in Figure 16 is a necessary condition for being dependent, from which we can see that any facet will be dependent on at most 2 other facets. However, towards the boundary of Γ_λ there are edges which are arranged as in Figure 16 which are not dependant.

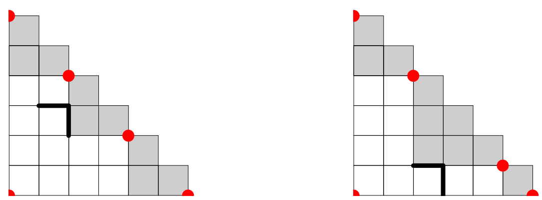

FIGURE 17. Examples of facets that are not dependant.

Facet dependency naturally partitions facets of GT_λ into visually simple sets called chains that will allow us to determine the orbits of facets under $\mathrm{Aut}(GT_\lambda)$ and various stabilizers.

Definition 4.13 (Facet Chains)**.** We can form maximal *facet chains* $C_i = \{F_{i,1}, ..., F_{i,l(i)}\}$ consisting of facets $F_{i,1}, ..., F_{i,l(i)}$ such that $F_{i,j+1}$ is dependent on $F_{i,j}$. Visually, a chain is the set of edges $e(F_{i,j})$ of Γ_λ and these edges form a zig-zag pattern. Let $\mathcal{C} = \{C_i\}$ denote the set of these facet chains. These chains partition the interior edges of Γ_λ.

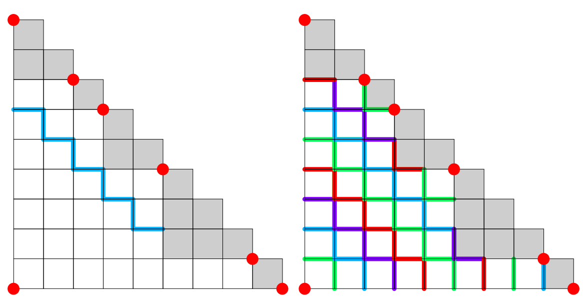

FIGURE 18. Left: a facet chain. Right: the partition of interior edges of Γ_λ.

Note that the condition of being a chain relies only on the dimension of an intersection of facets. In particular, if $\phi \in \mathrm{Aut}(GT_\lambda)$, and C_i is a chain, then $\phi(C_i)$ is a chain. One should also note that any chain C_i has exactly two facets, usually denoted $F_{i,1}, F_{i,l(i)}$ that are dependent on exactly one other facet in the chain.

Definition 4.14 (Adjacent Facets)**.** Two pairs of facets $\{F_1, F_2\}$ and $\{J_1, J_2\}$ are adjacent if $F_1 \cap F_2 = J_1 \cap J_2$.

Definition 4.15 (Adjacent Chains)**.** Two chains $C_i, C_j \in \mathcal{C}$ are *adjacent* if there exists $\{F_1, F_2\} \subset C_i$, $\{J_1, J_2\} \subset C_j$ such that $\{F_1, F_2\}$ and $\{J_1, J_2\}$ are adjacent.

Again, the condition of chain adjacency is given by intersection, so if chains C_i, C_j are adjacent then $\phi(C_i), \phi(C_j)$ must be adjacent. Visually, adjacency occurs iff one chain sits directly to the North-East of the other chain.

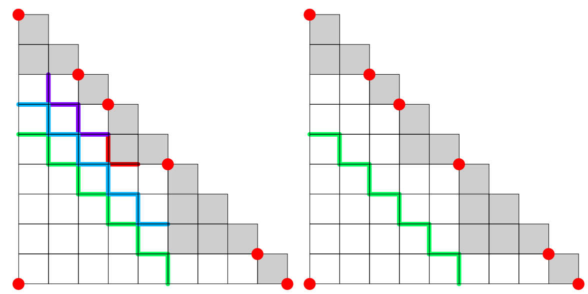

FIGURE 19. Left: example of adjacent chains. The purple, red, and green chains are each adjacent to the blue chain and are not adjacent to each other.
.

We now state two useful lemmas that will determine how automorphisms act on chains and pairs of adjacent chains.

Lemma 4.16. *Let $C = \{F_1, F_2, \ldots, F_k\}$ be a chain of facets, with F_1, F_k dependent on exactly one other facet, and suppose $\phi(C) = C$. Then either $\phi(F_i) = F_i$ for all i, or $\phi(F_i) = F_{k+1-i}$ for all i.*

Proof. If $\phi(F_1) = F_1$, the chain of dependencies of the F_i will determine the image of each F_i, so $\phi(F_i) = F_i$. If $\phi(F_1) \neq F_1$, the assumption of F_1 being dependent on exactly on facet means we must have $\phi(F_1) = F_k$, and again the chain of dependencies means that $\phi(F_2)$ must be dependent on $\phi(F_1)$, so $\phi(F_2) = F_{k-1}$, etc. □

Whenever $\phi(C) = C$ with $\phi(F_1) = F_k$, we will say that ϕ *flips* C, and otherwise we say that C is *not flipped* by ϕ.

However, chains cannot always be flipped independently of other chains, as the following lemma demonstrates.

Lemma 4.17. *Suppose C_i and C_j are adjacent chains of facets each with two distinct pairs of facets $\{F_1, F_2\}, \{F_1', F_2'\} \subset C_i, \{J_1, J_2\}, \{J_1', J_2'\} \subset C_j$ such that $\{F_1, F_2\}, \{J_1, J_2\}$ are adjacent and $\{F_1', F_2'\}, \{J_1', J_2'\}$ are adjacent. Then $\phi \in Aut(GT_\lambda)$ flips C_i iff ϕ flips C_j.*

Proof. Suppose ϕ does not flip C_i, so $\phi(F_1) = F_1$ and $\phi(F_2) = F_2$. Since $F_1 \cap F_2 = J_1 \cap J_2$, and the intersection of any pair of facets in a chain is unique, we must have $\phi(J_1), \phi(J_2) \in \{J_1, J_2\}$.

Suppose for contradiction that $\phi(J_1) = J_2$ implying C_j is flipped. Note flipping can preserve at most one adjacent pair of facets in a chain, we must have $\phi(\{J_1', J_2'\}) \neq \{J_1', J_2'\}$. Again the intersection of two facets in a chain is unique, implying $\phi(J_1') \cap \phi(J_2') \neq J_1' \cap J_2'$, but $\phi(J_1') \cap \phi(J_2') = \phi(F_1') \cap \phi(F_2') = F_1' \cap F_2' = J_1' \cap J_2'$, a contradiction. The reverse direction is similar. □

4.3. **The Automorphism Group.** We now show that the groups described in Theorems 4.7 and 4.8 form the entire automorphism group. We begin with the case when $m = 2$, and then prove the general case when $m \geq 3$.

Theorem 4.18 ($m = 2$ Automorphism Group)**.** *Suppose* $\lambda = (1^{a_1}, 2^{a_2})$ *and* $a_1, a_2 \geq 2$. *If* $a_1 = a_2 = 2$, *then*

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2.$$

Otherwise,

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}}$$

Proof. For convenience, define $G := \mathrm{Aut}(GT_\lambda)$, $G_a :=$ the stabilizer of a in G, $G_{a,b} :=$ the stabilizer of a, b, etc. and finally $Orb_H(a) :=$ the orbit of a with respect to a (sub)group H.

We will use the Orbit-Stabilizer theorem to show that $|\mathrm{Aut}(GT_\lambda)| = 16 \cdot 2^{\delta_{a_1,a_2}}$ when $a_1, a_2 \geq 3$, which will suffice to show equality of the groups in Thm. 4.7. There are two unique facets that are in maximal dependent chains of length 1, corresponding to the edges $((a_1, 0), (a_1, 1))$ and $((0, a_2), (1, a_2))$. Label these facets F_1 and F_2

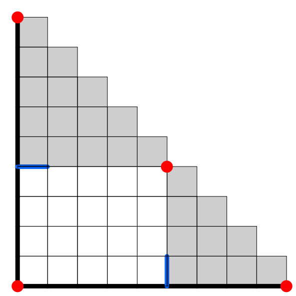

FIGURE 20. F_1 and F_2 drawn in blue.

These two facets are mapped to each other under α, and they cannot map to any other facets. Thus $|Orb_G(F_1)| = 2$, and if F_1 is fixed, F_2 must also be fixed.
Now consider the four facets contained in maximal chains of length two. Denote them J_1, J_2, J_3, J_4, where $\{J_1, J_2\}$ is the south-western chain and $\{J_3, J_4\}$ is the north-eastern chain.

These facets can only be mapped to each other. Note that $\mu, \mu_1, \alpha\tau \in G_{F_1}$, the compositions of which demonstrate that $|Orb_{G_{F_1}}(J_1)| = 4$. Now, $\mu_1 \in G_{F_1,J_1}$, which demonstrates that $|Orb_{G_{F_1,J_1}}(J_3)| = 2$. The results of Section 4.2 show that fixing J_1 fixes J_2, and similarly

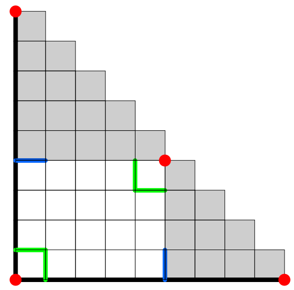

FIGURE 21. J_1, J_2, J_3, J_4 drawn in green.

fixing J_3 fixes J_4
Consider the chain K_1 immediately north-east of $\{J_1, J_2\}$, and label it's north-westernmost facet K_1'. Note that when $a_1 = a_2 = 2$, this chain is the chain $\{J_3, J_4\}$, so the following discussion will not apply.

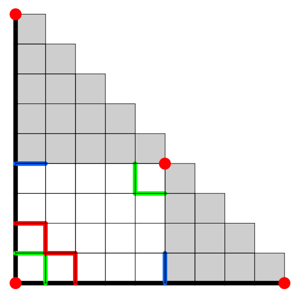

FIGURE 22. Chain K_1 drawn in red

Assuming J_1, J_2 are fixed, the discussion in Section 4.2 shows that K_1 can only be mapped to itself or its flip. If $a_1 = a_2$, the map $\rho\alpha\mu\mu_1 \in G_{F_1J_1J_3}$ flips K_1, and in particular shows that $|Orb_{G_{F_1J_1J_3}}(K_1')| = 2$. So now we must show that if $a_1 \neq a_2$, then K_1 cannot be flipped. Note we have a sequence of chains $K_1, K_2, \ldots, K_n$, where the chain K_n is defined as follows. Usually the length of K_{i+1} is the length of $K_i + 2$. K_n will be the first chain such that the length of K_n = length of $K_{n-1} + 1$. See Fig. **??**.

We will show that K_n cannot be flipped. There is a facet in K_n has exactly one dependent facet, and the intersection of these two facets is not equal to the intersection of any two pairs of facets in K_{n-1}, a property not other facet in K_n has (This corresponds to the south-easternmost facet of K_n in Fig. 23). Thus K_n is fixed. By lemma 4.17, this implies that K_1 is fixed, and we have shown that $|Orb_{G_{F_1J_1J_3}}(K_1')| = 2^{\delta_{a_1,a_2}}$ when $a_1, a_2 \geq 3$.
Combining the results so far, we have that

$$|G| = |Orb_G(F_1)||Orb_{G_{F_1}}(J_1)||Orb_{G_{F_1J_1}}(J_3)||Orb_{G_{F_1J_1J_3}}(K')||G_{F_1J_1J_3K'}| = 16 \cdot 2^{\delta_{a_1,a_2}}|G_{F_1J_1J_3K'}|$$

so it remains to show that $G_{F_1J_1J_3K'}$ is the trivial group. So suppose $\phi \in G$ fixes F_1, J_1, J_3 and K'. As remarked above, ϕ then fixes every facet in every chain of length ≤ 2. Since ϕ fixes the facet K', ϕ must also fix K_1. Every chain of length ≥ 3 is connected to K_1 by a

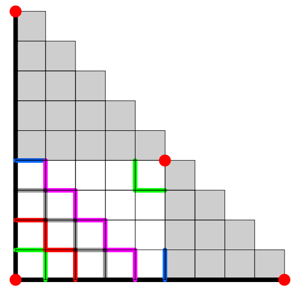

FIGURE 23. Chain K_n drawn in purple, which is K_3 in this example. Note this chain only exists if $a_1 \neq a_2$. Set of adjacent chains (in this case, just K_2) drawn in grey.

sequence of adjacent chains all of length ≥ 3, so ϕ maps every chain to itself and by Lemma 4.17, no chain can be flipped by ϕ, so ϕ in fact fixes every facet. □

Remark. Considering $\lambda = (a_1, a_2)$, denote G_λ to be the graph corresponding to the 1-skeleton of GT_λ, we can observe that $\mathrm{Aut}(GT_\lambda) \subseteq G_\lambda$, and it is natural to ask when we have equality. Numerical computations suggest that we have equality whenever $a_1, a_2 \neq 3$. However, the approach of computing orbit sizes for $\mathrm{Aut}(G_\lambda)$ is exceedingly difficult in general.

Theorem 4.19 ($m \geq 3$ Automorphism Group)**.** *Suppose $\lambda = 1^{a_1} \dots m^{a_m}$ and $m \geq 3$. Let $t = 1$ if $\lambda = \lambda'$ and let $t = 0$ otherwise. Let j be the number of pairs $a_k, a_{k+1} \geq 2$. Then*

$$Aut(GT_\lambda) \cong \mathbb{Z}_2^t \ltimes_\varphi (S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{j+1}).$$

Proof. For convenience, let $H_\lambda := \mathbb{Z}_2^t \ltimes_\varphi (S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{j+1})$. Theorem 4.8 shows that H_λ is a subgroup of $\mathrm{Aut}(\mathrm{GT}_\lambda)$. Now we show that the order of $\mathrm{Aut}(\mathrm{GT}_\lambda)$ is at most the order of H_λ. By Lemma 4.9, it suffices to consider the possible ways an automorphism $\phi \in \mathrm{Aut}(\mathrm{GT}_\lambda)$ can permute the facets of GT_λ.

The action of ϕ on the facets of GT_λ can be extended naturally to the facet chains in $\mathcal{C}$ as defined Definition 4.13. In particular, ϕ must send every chain to another chain of the same length.

Assume $a_1 \neq 1$ and $a_m \neq 1$. Then there are two chains consisting of a single facet. Denote these as $C_1 = (F_{1,1})$ and $C_{2m-1} = (F_{2m-1,1})$. These facets are the uppermost horizontal edge and the rightmost vertical edge respectively.

Now consider the chains of length 2. Let C_0 denote the length 2 chain near the origin. Let $C_2, \dots, C_{2m-2}$ denote the length 2 chains that occur along the border of Γ_λ where C_{2k} occurs near terminal vertex t_k and C_{2k-1} occurs at the corner of the kth shaded triangular subgrid corresponding to i^{a_i} in the the partition λ. Finally, we call the C_{2k-1} *type A* chains and the C_{2k} *type B* chains. Note that some type B chains may not exist if some $a_i = 1$. See Figure 25 for an example of this. When referring to all type B chains, we only refer to those that exist.

A type A chain $C_{2j-1} = (F_{2j-1,1}, F_{2j-1,2})$ and a type B chain $C_{2k} = (F_{2k,1}, F_{2k,2})$ are called *incompatible* if there does not exist a vertex containing both $e(F_{2j-1,1}), e(F_{2j-1,2})$ and at least

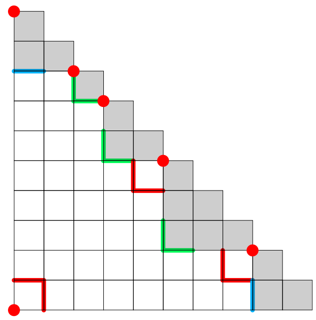

FIGURE 24. Length 1 chains shown in blue. Type A chains shown in green. Type B chains shown in red.

one of $e(F_{2k,1}), e(F_{2k,2})$. C_1 (resp. C_{2m-1}) is *incompatible* with type B chain $(F_{2k,1}, F_{2k,2})$ if there does not exist a vertex containing both $e(F_{2k,1}), e(F_{2k,2})$ and $e(F_{1,1})$ (resp. $e(F_{2m-1,1})$).

C_0 is compatible with all type A chains. For $1 \leq k \leq m-1$, the type B chain C_{2k} is incompatible with its neighboring type A chains C_{2k-1} and C_{2k+1}.

C_0 must be mapped to itself since all other type B chains are incompatible with some type A chain. (This is not true if $m = 2$ which is handled below.) The map ϕ must send C_1 to itself or to C_{2m-1}. If the former occurs, then $\phi(C_2) = C_2$ and so on, such that $\phi(C_i) = C_i$ for all $1 \leq i \leq 2m-1$. If the latter occurs, then $\phi(C_i) = C_{2m-i}$ for all $1 \leq i \leq 2m-1$.

Assume $\phi(C_i) = C_i$ for all $1 \leq i \leq 2m-1$, i.e. the order of the sequence of type A and type B chains is preserved. Facets in chain C_i must be mapped to other facets in C_i so $\phi(F_{1,1}) = F_{1,1}$ and $\phi(F_{2m-1,1}) = F_{2m-1,1}$. For each $0 \leq k \leq m-1$, either C_{2k} is flipped or not flipped. These two possibilities account for whether ϕ contains any of $\mu, \mu_1, \ldots, \mu_{m-1}$. On the other hand, ϕ cannot interchange $F_{2k-1,1}$ and $F_{2k-1,2}$ because there exist vertices containing $e(F_{2k-1,2})$ and one of $e(F_{2k,1}), e(F_{2k,2})$, but there do not exist vertices containing $e(F_{2k-1,1})$ and one of $e(F_{2k,1}), e(F_{2k,2})$. If $\phi(C_i) = C_{2m-i}$ for all $1 \leq i \leq 2m-1$, the analysis is similar with each C_{2k} being sent to C_{2m-2k} of the flip of C_{2m-2k}.

Say C_{long} has endpoints at $(0, a)$ and $(a, 0)$. For $1 \leq i \leq a$, define $C'_i \in \mathcal{C}$ to be the chain with endpoints at $(0, i)$ and $(i, 0)$. Then $C'_1 = C_0$ and $C'_a = C_{long}$. For $i \leq a-1$, C'_i is only adjacent to C'_{i-1} and C'_{i+1}. The map ϕ must preserve adjacencies between chains. Since $\phi(C'_1) = C'_1$, $\phi(C'_2) = C'_2$ and either C'_2 is flipped or not flipped. In either case, C'_k has the same orientation as C'_{k-1} for all $3 \leq k \leq a$.

We have four cases:

(1) $\phi(C_i) = C_i$ for all $1 \leq i \leq 2m-1$. All of C'_k for $2 \leq k \leq a$ are flipped.
(2) $\phi(C_i) = C_i$ for all $1 \leq i \leq 2m-1$. All of C'_k for $2 \leq k \leq a$ are not flipped.
(3) $\phi(C_i) = C_{2m-i}$ for all $1 \leq i \leq 2m-1$. All of C'_k for $2 \leq k \leq a$ are flipped.
(4) $\phi(C_i) = C_{2m-i}$ for all $1 \leq i \leq 2m-1$. All of C'_k for $2 \leq k \leq a$ are not flipped.

We show that cases (1) and (4) lead to contradictions and that case (3) occurs only if $\lambda = \lambda'$, accounting for a factor of 2 due to ρ.

Consider all the chains above C_{long}. Starting with the chains adjacent to C_{long}, we note that all chains are of length ≤ 2 which have been accounted for. All chains of greater length share at least 2 vertices with C_{long} so their position is already determined by C_{long}'s orientation. If C_{long} is not flipped, then all of its adjacent chains are not flipped. We continue this argument with each of C_{long}'s adjacent chains, moving towards the main diagonal. After all chains have been dealt with, note that the adjacencies of type A and type B chains with chains below them are only preserved if $\phi(C_i) = C_i$ for all $1 \leq i \leq 2m-1$. This shows case (4) cannot occur.

If C_{long} is flipped, then for every chain C'' above C_{long} that is adjacent to C_{long}, there must a chain of the same length that is a reflection of C'' across $x = y$. Again, we continue this argument upwards moving towards the main diagonal. By looking at adjacencies of type A and type B chains with chains below them, we conclude that case (1) cannot occur. In order for reflections of chains to exist as each step, Γ_λ must be symmetric about $x = y$ which occurs iff $\lambda = \lambda'$. This accounts for case (3).

If WLOG $a_m = 1$, then we have additional type 1 chains as shown in Figure 25. If any length 1 chain in the bottom right is sent to a length 1 chain in the upper left, then all of them are. The set of length 1 chains in the bottom right can only be sent to the set of length 1 chains in the upper left if $a_1 = a_m = 1$ and $a_2 = a_{m-1}$. Following the argument above shows that this can only occur if $\lambda = \lambda'$.

If the set of length 1 chains in the bottom right is sent to itself, then accounting for any permutation of the a_{m-1} length 1 chains gives us a factor of $a_{m-1}!$ in the order of $\text{Aut}(\text{GT}_\lambda)$. Similarly for chains in the upper left. Then we can fix any permutations in $S_{a_2}, S_{a_{m-1}}$ for the length 1 chains and apply the same argument above.

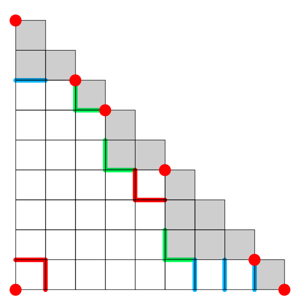

FIGURE 25. Length 1 chains shown in blue. Type A chains shown in green. Type B chains shown in red.

Note that for each factor of 2 or $a_2!$ or $a_{m-1}!$, the conditions under which we include the factor in the order of $\text{Aut}(\text{GT}_\lambda)$ exactly match the conditions under which we include Z_2's or S_{a_2} or $S_{a_{m-1}}$ in H_λ. Thus we have bounded the order of $\text{Aut}(\text{GT}_\lambda)$ by the order of H_λ. Since H_λ is contained in $\text{Aut}(\text{GT}_\lambda)$, we have $\text{Aut}(\text{GT}_\lambda) \cong H_\lambda$, as desired. ☐

Computing the number of vertices in GT_λ is a difficult question, and one that is addressed in both [GKT13] and [ACK16]. When $m = 3$, this question actually has a complete answer. For the purposes of this discussion, let $V(\mathrm{GT}_\lambda)$ denote the number of vertices in GT_λ, and recall that $n = a_1 + a_2 + a_3$. Gusev-Kiritchenko-Timorin originally derive the following expression analytically from a generating function. However, there is a combinatorial proof of this formula, which we give here.

Theorem 5.1 (Theorem 1.4 in [GKT13])**.** *If $\lambda = (1^{a_1}, 2^{a_2}, 3^{a_3})$, then we have the following equality*

$$V(GT_\lambda) = \binom{n}{a_1}\binom{n}{a_2} + 2\sum_{i=1}^{a_1}(-1)^i\binom{n}{a_1 - i}\binom{n}{a_2 - i}$$

Proof. The ladder diagram of a vertex will be a tree with root $(0,0)$ and 4 terminal vertices. The paths to $(n,0)$ and $(0,n)$ are always straight lines, but the middle paths form a tree, branching at a single node and having these two branches never intersect again. We can interpret this as the number of pairs of positive paths that both start at $(0,0)$, never intersect once they diverge, and have endpoints (a_1, a_2+a_3) and (a_1+a_2, a_3) respectively. The number of all pairs of paths from $(0,0)$ to (a_1, a_2+a_3) and (a_1+a_2, a_3) is exactly $\binom{n}{a_1}\binom{n}{a_2}$, so now we must subtract off the paths that diverge and then intersect again.
♠ To do: Add in rest of proof □

Using this combinatorial interpretation, we hoped find a similar formula when $m = 4$. In this case however, there is one common base point for 3 paths which branch at 2 distinct points. It's not clear how to pull apart one pair without accidentally intersecting with the other path. There is another approach that this problem to a question about non-intersecting lattice paths, which is well-studied and may eventually give an answer.

Proposition 5.2. *When $m = 4$, let $NI(p_1, p_1', p_2, p_2', p_3, p_3', p_4, p_4')$ denote the number of quadruples of non-intersecting positive lattice paths starting at p_i and ending at p_i'. Let x, y and z denote the terminal vertices $(a_1, n - a_1)$, $(n - a_4, a_4)$ and $(n - a_1 - a_2, a_3 + a_4)$. We have the following equality*

$$V(GT_\lambda) = \sum_{\substack{0 \le i \le a_1 \\ 0 \le j \le a_4}} \binom{i+j}{i} \Big(\sum_{\substack{i+1 \le k \le a_1+a_2 \\ j \le l \le a_4}} NI((i+1,j),(k,l),(i,j+1),x,(k+1,l),y,(k,l+1),z)$$
$$+ \sum_{\substack{i \le k \le a_1 \\ j+1 \le l \le a_3+a_4}} NI((i,j+1),(k,l),i+1,j),y,(k+1,l),z,(k,l+1),x)\Big)$$

Proof. The ladder diagram of a vertex in GT_λ is the union of positive paths that diverge into two paths at (i,j) and diverge into two more paths at (k,l). Note at the divergence at (i,j), there are two types of vertices. The path that diverges to the right from (i,j) can either diverge once more and go to y and z or not diverge and just go to y. The first inner sum corresponds to the former, and sums over all possible places this right path can diverge.

The second sum computes the latter case, so the path going up from (i,j) is the one that must diverge. □

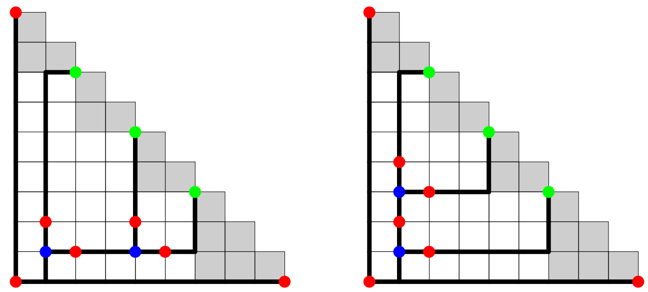

FIGURE 26. The sums count the possible placements of the blue dots where the paths diverge. The left figure corresponds to the first inner sum.

Remark 5.3. For some configurations of points, it is possible to count the number of non-intersecting lattice paths between them. However, this is only doable when the paths meet certain conditions. Label the start points of the paths $1,2,\dots,n$ and the endpoints $1,2,\dots,n$. We can think of a valid set of paths as giving a permutation on $[n]$ in 1-line notation, where the i^{th} position is given by the endpoint of the path that begins at i. The usual determinant formula for non-intersecting lattice paths will give the correct answer provided that all arrangements of paths give a permutation having the same sign. Unfortunately this condition is not met in the previous proposition, which is why we have a formula in terms of this NI function. Note that this proposition can theoretically be generalized to any m by summing over all cases of points where paths can diverge and those possible sets of endpoints. However, the final answer will always rely on some version of this NI function, which we cannot in general compute.

6. A Special Case: $\lambda = 12^{n-2}3$

As we can see in the previous section, the enumeration of vertices and faces in general is difficult. However, in the special case when $\lambda = (1, 2^{n-2}, 3)$, the face poset of GT_λ admits a fairly simple combinatorial description, from which we can write down a nice expression counting the number of faces of each dimension.

Proposition 6.1. *For $\lambda = (1, 2^{n-2}, 3)$, $\mathcal{F}(GT_\lambda)$ is isomorphic to the set of pairs $(A,B) \subseteq [n]\times[n]$ ordered by inclusion with the following conditions*

(1) $A = \varnothing$ iff $B = \varnothing$
(2) $n-1 \in A \cap B$ iff $n \in A \cap B$
(3) The rank of an element $\rho((A,B))$ is given by

$$\rho((A,B)) = \begin{cases} |A| + |B| - 2 & \{n-1,n\} \notin A \cap B \\ |A| + |B| - 3 & \{n-1,n\} \in A \cap B \end{cases}$$

Proof. Note that any tiling of GT_λ has exactly one column and one row of unforced squares that intersect at exactly one tile. Label the borders between squares of the column $1,\dots,n$

going down and the borders between the squares of the rows $n, n-1, \ldots, 1$ going left to right. So note the bottom-left most square has all of its edges labeled either $n-1$ or n. Given a tiling of GT_λ, read off the labels in the column that are borders between tiles. Stop if you reach a tile of width > 1. These labels will be the set A. Note if the bottom-most square in the column is reached, the label n is included. Repeat this same algorithm starting from the right in the last row, checking to make sure that the tiles are now of height 1. These labels will be the set B. Note that if n has been included in A but $n-1$ has not been, the bottom-leftmost square is in a tile of size > 1 extending into the column, and so cannot be in B, so B cannot have the label n. Now if $n-1, n \in A$, the bottom-leftmost square must be in it's own tile. If this tile borders a tile of height > 1, B contains neither n or $n-1$. If this tile borders a tile of height ≤ 1, B will contain both $n-1$ and n. Noting that the number of tiles, and thus the dimension, increases for every additional label added to either A or B except when n is added to B completes the proof. See Figure 27. □

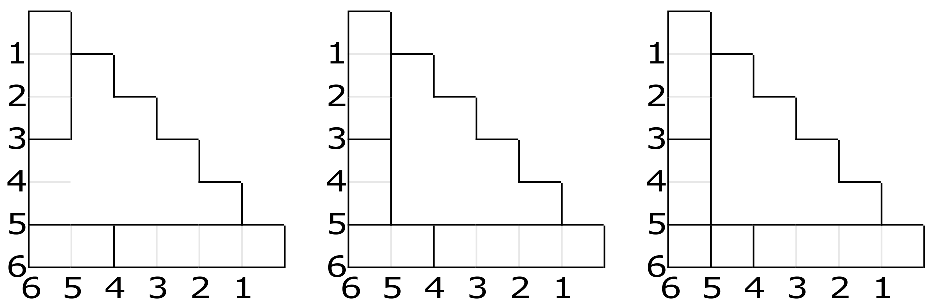

FIGURE 27. The sets (A, B) for these tilings are $(\{3\}, \{4, 6\})$, $(\{3, 5\}, \{4, 6\})$ and $(\{3, 5, 6\}, \{4, 5, 6\})$

Remark. Note that the face poset for $\Delta_n \times \Delta_n$ (the cross product of two simplices on n vertices) is exactly the poset described in Prop. 6.1, except without condition (2) and with condition (3) being changed to $\rho((A, B)) = |A| + |B| - 2$. This is not a coincidence, as we will see in the next section.

Using this simple combinatorial description of the face poset of GT_λ, we can derive a simple expression counting the number of faces of each dimension in GT_λ.

Definition 6.2. Given a polytope P, the f-polynomial $f(P)$ is defined as

$$f(P) = \sum_{e \in F(P)} t^{\dim(e)} = \sum_{i=0}^{\dim(P)} |\text{faces of dimension } i| t^i$$

Proposition 6.3. *If $\lambda = (1, 2^{n-2}, 3)$, the following holds*

$$\begin{aligned} f(GT_\lambda) &= \sum_{\substack{\emptyset \neq A \subset [n] \\ \emptyset \neq B \subset [n]}} t^{|A|+|B|-2} - \sum_{\substack{A \subset [n-2] \\ B \subset [n-2]}} t^{|A|+|B|-2}(2t^2 + 4t^3) + \sum_{\substack{A \subset [n-2] \\ B \subset [n-2]}} t^{|A|+|B|-2}(t^3 - t^4) \\ &= \left(\frac{(1+t)^n - 1}{t}\right)^2 - (t^2 + 3t + 2)(1+t)^{2n-4} \end{aligned}$$

Proof. Note the term $\sum_{\substack{\emptyset\neq A\subset[n]\\ \emptyset\neq B\subset[n]}} t^{|A|+|B|-2}$ corresponds to the poset in Proposition 6.1 without conditions (2) and (3). To exclude all pairs (A,B) not satisfying (2), there are a few cases that we must account for:

(1) $n-1 \in A \bigcap B$ but $n \notin A \bigcup B$
(2) $n \in A \bigcap B$ but $n-1 \notin A \bigcup B$
(3) $n-1 \in A \bigcap B$ and $n \in A$ but $n \notin B$
(4) $n \in A \bigcap B$ and $n-1 \in A$ but $n-1 \notin B$
(5) $n-1 \in A \bigcap B$ and $n \in B$ but $n \notin A$
(6) $n-1 \in A \bigcap B$ and $n \in B$ but $n \notin A$

For every pair $(A', B') \subset [n-2] \times [n-2]$, we can extend it to a subset of $[n] \times [n]$ falling in cases 1 or 2 by adding 2 elements, and extend it to a subset in cases 3, 4, 5 or 6 by adding 3 elements. Thus the sets (A,B) not satisfying (2) can be accounted for by the term

$$\sum_{\substack{A\subset[n-2]\\ B\subset[n-2]}} t^{|A|+|B|-2}(2t^2 + 4t^3)$$

Now we must take all subsets (A,B) affected by condition (3) and change their rank. For every pair (A', B'), we can extend it to a set with $n-1, n \in A \bigcap B$ by adding 4 elements. However, this should count the rank as $|A| + |B| - 3$, instead of $|A| + |B| - 2$. So we subtract off every such pair, and add it back with the correct rank. This gives the term

$$\sum_{\substack{A\subset[n-2]\\ B\subset[n-2]}} t^{|A|+|B|-2}(-t^4 + t^3)$$

From here, the binomial theorem gives that

$$\sum_{\substack{\emptyset\neq A\subset[n]\\ \emptyset\neq B\subset[n]}} t^{|A|+|B|} = (\sum_{\emptyset\neq A\subset[n]} t^{|A|})^2 = (1+t)^n - 1$$

$$\sum_{\substack{A\subset[n-2]\\ B\subset[n-2]}} t^{|A|+|B|} = (\sum_{A\subset[n-2]} t^{|A|})^2 = (1+t)^{2n-4}$$

completing the proof. □

Remark. From the poset in Prop. 6.1, one could write down expressions for the flag-f vector as well, as sums of products of multinomials. Essentially, a face corresponds to a pair of sets (A,B). A chain of faces corresponds to a sequence of pairs of sets $(A_i, B_i)_{1\leq i\leq k}$ such that $A_1 \subseteq A_2 \subseteq \ldots \subseteq A_k$ and $B_1 \subseteq B_2 \subseteq \ldots \subseteq B_k$. By doing casework on when n first appears in the sequences $\{A_i\}$ and $\{B_i\}$, we can explicitly count the number of such chains where each set has a specified size. However, this expression is fairly complicated and un-enlightening, and it is unknown if there is a concise way to describe it in a nicer form.

7. Observations, Rederived Results, and some Failed Ideas

As remarked in the previous section, it is not a coincidence that $\lambda = (1, 2^{n-2}, 3)$ the face poset of GT_λ closely resembles that of $\Delta_n \times \Delta_n$. This GT polytope can be realized as a Minkowski sum of two simplices with a certain embedding. In fact more is true– any GT-polytope can be realized as a sum of smaller GT_λ, a notion we make precise here.

Definition 7.1. Given polytopes $P, Q \subset \mathbb{R}^d$, the Minkowski sum of P and Q is defined as

$$P + Q = \{p + q | p \in P, q \in Q\}$$

Proposition 7.2. *Let $\lambda = (0^{a_0}, 1^{a_1}, \ldots, m^{a_m})$. Let GT_{λ_k} be the GT polytope corresponding to $\lambda_k = (0^{a_1+a_2+\ldots+a_k}, 1^{a_{k+1}+\ldots+a_m})$. Then $GT_\lambda = GT_{\lambda_1} + GT_{\lambda_2} + \ldots + GT_{\lambda_{m-1}}$*

Proof. Note the set on the right hand side has fixed $x_{i,i}$ coordinates, and that $x_{i,i} = \lambda_i$. Furthermore, any point on the right hand side is a sum of points which satisfy the inequalities in Figure 1, and so satisfies the inequalities in Figure 1. Since the Minkowski sum preserves convexity, it suffices to show that this sum surjects onto the vertices of GT_λ. Consider the tiling T of a vertex v of GT_λ. It must have exactly $m + 1$ tiles, each correspoding to the values 0 to m. Consider the tiling in each GT_{λ_k} that has every tile in T corresponding to a value $> k$ in one tile (corresponding to 1), and every tile in T corresponding to a value $\leq k$ in another tile (corresponding to 0). These vertices in each GT_{λ_k} corresponding to these tiles sum to v, so this sum is indeed surjective. □

To be quite honest, we are currently unsure if this is useful or not– still wading through the literature on that. It is interesting though.

Proposition 7.3. *A (known) formula for the number of SSYT of shape λ in terms of non-intersecting lattice paths.*

Proof. Add in counting integer points in terms of non-intersecting lattice paths □

Proposition 7.4. *When $\lambda = (0^{a_1}, 1^{a_2})$, GT_λ is isomorphic to the order polytope for the product of a chain of length a_1 and a chain of length a_2.*

Proof. Talk a bit about order polytopes and how GT-polytopes are related. □

Proposition 7.5. *The following idea to compute the f-vector for GT_λ when $\lambda = (0^{a_1}, 1^{a_2})$ fails.*

Proof. Talk about stripping away columns, idea to keep track of 'heights' of paths. □

Acknowledgements

This research was carried out as part of the 2016 REU program at the School of Mathematics at University of Minnesota, Twin Cities, and was supported by NSF RTG grant DMS-1148634 and by NSF grant DMS-1351590. The authors would like to thank Elise delMas and Craig Corsi for their valuable advice and comments. The authors are especially grateful to Victor Reiner for his mentorship and support, and for many fruitful conversations.

References

[ACK16] Byunghee An, Yunhyung Cho, and Jang Soo Kim. On the f-vectors of gelfand-cetlin polytopes. *arXiv preprint arXiv:1606.05957*, 2016.

[GKT13] Pavel Gusev, Valentina Kiritchenko, and Vladlen Timorin. Counting vertices in gelfand-zetlin polytopes. *J. Comb. Theory, Ser. A*, 120(4):960–969, 2013.

[McA06] Tyrrell B McAllister. *Applications of Polyhedral Geometry to Computational Representation Theory*. PhD thesis, University of California Davis, 2006.

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139

E-mail address: gaoyibo@mit.edu

Department of Mathematics, Yale University, New Haven, CT 06520

E-mail address: benjamin.krakoff@yale.edu

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139

E-mail address: lisayang@mit.edu