DECIDING POSITIVITY OF LITTLEWOOD–RICHARDSON COEFFICIENTS

PETER BÜRGISSER* AND CHRISTIAN IKENMEYER†

Abstract. Starting with Knutson and Tao's hive model [KT99], we characterize the Littlewood–Richardson coefficient $c^{\nu}_{\lambda,\mu}$ of given partitions $\lambda,\mu,\nu \in \mathbb{N}^n$ as the number of capacity achieving hive flows on the honeycomb graph. Based on this, we design a polynomial time algorithm for deciding $c^{\nu}_{\lambda,\mu} > 0$. This algorithm is easy to state and takes $\mathcal{O}\big(n^3 \log \nu_1\big)$ arithmetic operations and comparisons. We further show that the capacity achieving hive flows can be seen as the vertices of a connected graph, which leads to new structural insights into Littlewood–Richardson coefficients.

Key words. Littlewood–Richardson coefficients, hive model, polynomial time algorithm, flows in networks

AMS subject classifications. 05E05, 22E46, 90C27

1. Introduction. Let $\lambda,\mu,\nu \in \mathbb{Z}^n$ be nonincreasing n-tuples of integers. The *Littlewood–Richardson coefficient* $c^{\nu}_{\lambda,\mu}$ is defined as the multiplicity of the irreducible $\mathsf{GL}_n(\mathbb{C})$-representation V_ν with dominant weight ν in the tensor product $V_\lambda \otimes V_\mu$. These coefficients appear not only in representation theory and algebraic combinatorics, but also in topology and enumerative geometry (Schubert calculus): for instance, they determine the multiplication in the cohomology ring of the Grassmann varieties. Littlewood–Richardson coefficients gained further prominence due to their role in the proof of Horn's conjecture [HR95, Kly98, KT99, KTW04] on the relation of the eigenvalues of a triple A, B, C of Hermitian matrices satisfying $C = A + B$. The latter problem is of relevance in perturbation and quantum information theory. We refer to Fulton [Ful00] for an excellent account of these more recent developments.

Different combinatorial characterizations of the Littlewood–Richardson coefficients are known. The classic Littlewood–Richardson rule (cf. [Ful97]) counts certain skew tableaux, while in Berenstein and Zelevinsky [BZ92], the number of integer points of certain polytopes are counted. A beautiful characterization was given by Knutson and Tao [KT99], who characterized Littlewood–Richardson coefficients either as the number of honeycombs or hives with prescribed boundary conditions.

The focus of this paper is on the complexity of computing the Littlewood–Richardson coefficient $c^{\nu}_{\lambda,\mu}$ on input λ,μ,ν. Without loss of generality we assume that the components of λ,μ,ν are nonnegative integers and put $|\lambda| := \sum_i \lambda_i$. Moreover we write $\ell(\lambda)$ for the number of nonzero components of λ. Then $|\nu| = |\lambda| + |\mu|$ and $\nu_1 \geq \max\{\lambda_1, \mu_1\}$ are necessary conditions for $c^{\nu}_{\lambda,\mu} > 0$. We think of λ,μ,ν as encoded in binary and interpret $\sum_{i=1}^n (\log \lambda_i + \log \mu_i + \log \nu_i) \leq 3n \log \nu_1$ as a measure of the input size. All the algorithms derived from the above mentioned characterizations of Littlewood–Richardson coefficients take exponential time in the worst case. Narayanan [Nar06] proved that this is unavoidable: the computation of $c^{\nu}_{\lambda,\mu}$ is a $\#\mathbf{P}$-complete problem. Hence there does not exist a polynomial time algorithm for computing $c^{\nu}_{\lambda,\mu}$ under the widely believed hypothesis $\mathbf{P} \neq \mathbf{NP}$.

Main results. We first characterize $c^{\nu}_{\lambda,\mu}$ as the number of *capacity achieving hive flows* on the honeycomb graph G, cf. Figures 2.1–2.2. Besides capacity constraints given by λ,μ,ν, these flows have to satisfy rhombus inequalities corresponding to the ones considered in [KT99, Buc00]. We then develop a polynomial time algorithm (Algorithm 2) for deciding $c^{\nu}_{\lambda,\mu} > 0$ with $\mathcal{O}\big(n^3 \log \nu_1\big)$ arithmetic operations and comparisons. This is basically a capacity scaling Ford-Fulkerson algorithm [FF62] on well-chosen residual networks. The

*Institute of Mathematics, University of Paderborn, D-33098 Paderborn, Germany, pbuerg@math.upb.de, partially supported by DFG grant BU 1371/3-2.

†Institute of Mathematics, University of Paderborn, D-33098 Paderborn, Germany, ciken@math.upb.de, partially supported by DFG grant BU 1371/3-2.

algorithm is easy to state and implement: we encourage the reader to try out our Java applet at `http://www-math.upb.de/agpb/flowapplet/flowapplet.html`. We also show that the set of capacity achieving hive flows is the vertex set of a natural connected graph, which is relevant for efficiently enumerating these flows. In fact, our work is the basis of a follow-up paper [Ike12], in which more algorithmic insights are obtained, notably an algorithm for deciding $c^{\nu}_{\lambda,\mu} \geq t$ in time polynomial in $n \log \nu_1$ and t. This implies that "small" Littlewood–Richardson coefficients can be efficiently computed. In [Ike12] we also prove a conjecture on stretched Littlewood–Richardson coefficients posed by King, Tollu, and Toumazet in [KTT04].

Motivation and previous work. Our investigations are motivated by Geometric Complexity Theory, an approach towards proving fundamental complexity lower bounds by means of algebraic geometry and representation theory, that was initiated by Mulmuley and Sohoni [MS01, MS08] (see [Mul11] for recent pointers to the literature). To the best of our knowledge, the existence of a polynomial time algorithm for deciding $c^{\nu}_{\lambda,\mu} > 0$ was first pointed out in [DLM06], and one day later in [MS05]. Indeed, by the saturation property, $c^{\nu}_{\lambda,\mu} > 0$ is equivalent to $\exists N\, c^{N\nu}_{N\lambda,N\mu} > 0$, which can be rephrased as the feasibility problem of a certain rational polyhedron, whose elements are called *hives*. Feasibility of rational polyhedra is a basic problem in linear programming, well-known to be solvable in polynomial time, cf. [GLS93]. More specifically, the existence of a hive with prescribed boundary conditions can be expressed as the feasibility of a system $Ax \leq b$ of linear inequalities, where the entries of the matrix A are in $\{-1, 0, 1\}$ and the components of the vector b are either zero or among the components of λ, μ, ν. For the format $M \times N$ of the matrix A we have $M, N = \mathcal{O}(n^2)$. The basic ellipsoid method (cf. [GLS93, p. 80]) for solving this feasibility problem takes $\mathcal{O}(MN^3\ell)$ arithmetic steps and comparisons, where $\ell = \mathcal{O}(n^2 \log \nu_1)$ is the encoding length of the linear program. This gives a bound of $\mathcal{O}(n^{10} \log \nu_1)$, which is considerably worse than the bound $\mathcal{O}(n^3 \log \nu_1)$ proven for Algorithm 2 in this paper. We also note that standard interior point methods at least require $n^9 \log(n\nu_1)$ arithmetic operations, cf. [BC13, Chap. 10].

The starting point for the present work was a question in [MS05] asking for a combinatorial algorithm for deciding $c^{\nu}_{\lambda,\mu} > 0$ in polynomial time, using ideas similar to the max-flow or weighted matching problems in combinatorial optimization.

The algorithm in this paper is a considerable improvement over the one presented by the authors at FPSAC 2009 [Ike08, BI09], both with regard to simplicity and running time. The reason is that there, before each augmentation step, the flow had to be substituted by a nondegenerate flow using a costly routine. (Nondegenerate meaning that small triangles and small rhombi are the only flatspaces, cf. [Buc00] and Remark 4.21.) The present algorithm does not suffer from this deficiency anymore.

Outline of paper. Section 2 describes the setting and introduces the main terminology. We define the notion of hive flows on honeycomb graphs and associate with a triple λ, μ, ν of partitions the polytope $B(\lambda, \mu, \nu)$ of bounded hive flows, along with a linear function δ measuring the overall throughput of a flow. A flow $f \in B(\lambda, \mu, \nu)$ is called *capacity achieving* if $\delta(f) = |\nu|$. We denote by $P(\lambda, \mu, \nu)$ the polytope consisting of these flows and by $P(\lambda, \mu, \nu)_{\mathbb{Z}}$ the set of its integral points. It turns out that the Littlewood–Richardson coefficient $c^{\nu}_{\lambda,\mu}$ counts the elements of $P(\lambda, \mu, \nu)_{\mathbb{Z}}$ (Proposition 2.7).

In Section 3 we obtain some structural insights into the set of hive flows. We show that $P(\lambda, \mu, \nu)_{\mathbb{Z}}$ is the vertex set of a natural connected graph (*Connectedness Theorem* 3.12). The connectedness immediately implies the property $c^{\nu}_{\lambda,\mu} = 1 \Rightarrow \forall N\, c^{N\nu}_{N\lambda,N\mu} = 1$, that was conjectured by Fulton and proved in [KTW04] (Corollary 3.14). The connectedness of $P(\lambda, \mu, \nu)_{\mathbb{Z}}$ is also relevant for efficiently enumerating the points of $P(\lambda, \mu, \nu)_{\mathbb{Z}}$ and for proving the implication $c^{\nu}_{\lambda,\mu} = 2 \Rightarrow \forall N\, c^{N\nu}_{N\lambda,N\mu} = N + 1$, which was conjectured by King, Tollu, and Toumazet [KTT04], cf. [Ike12].

Proposition 2.7 suggests to decide $c^{\nu}_{\lambda,\mu} > 0$ by optimizing the overall throughput function δ on the polytope $B(\lambda, \mu, \nu)$ of bounded hive flows. We imitate the basic Ford-Fulkerson idea and construct, for a given integral hive flow f, a "residual digraph" R_f, such that f optimizes δ on $B(\lambda, \mu, \nu)$ iff R_f does not contain an s-t-path. In Section 4 we define the residual digraph R_f and study the partition of the triangular graph into f-flatspaces. We present and analyze a first max-flow algorithm for deciding $c^{\nu}_{\lambda,\mu} > 0$ (Algorithm 1). The proof of correctness of this algorithm requires an in-depth understanding of the properties of hives and it has two main ingredients. The *Shortest Path Theorem* 4.8 states that the rhombus inequalities are not violated after augmenting the current flow f by a shortest path in the residual network R_f. This is remarkable since, unlike in the usual max-flow situation, the polytopes of hive flows are not integral, cf. [Buc00]. The other ingredient, needed for the optimality criterion, is the *Rerouting Theorem* 4.19, which tells us how to replace an augmenting flow direction d by a flow in the residual network without changing the overall throughput. This amounts to a rerouting of d along the borders of the flatspaces of the current flow f (cf. Figure 4.1). Here the main difficulty is the analysis of the degenerate situation of large flatspaces, a topic not pursued in detail in the previous papers [KT99, Buc00].

In Section 5 we state and analyze our polynomial time Algorithm 2 for deciding the positivity of Littlewood–Richardson coefficients.

The remainder of the paper is devoted to the combinatorially quite intricate proofs of the Rerouting Theorem, the Shortest Path Theorem, and the Connectedness Theorem.

2. Flow description of LR coefficients.

2.1. Flows on digraphs. We fix some terminology regarding flows on directed graphs, compare [AMO93]. Let D be a digraph with vertex set $V(D)$ and edge set $E(D)$. We assume that $s, t \in V(D)$ are two different distinguished vertices, called source and target, respectively. Let $e_- := u$ denote the vertex where the edge e starts and $e_+ := v$ the vertex where e ends. The *inflow* and *outflow* of a map $f\colon E(D) \to \mathbb{R}$ at a vertex $v \in V(D)$ are defined as

$$\mathsf{inflow}(v, f) := \sum_{e_+ = v} f(e), \quad \mathsf{outflow}(v, f) := \sum_{e_- = v} f(e),$$

respectively. A *flow on* D is defined as a map $f\colon E(D) \to \mathbb{R}$ that satisfies Kirchhoff's conservation laws: $\mathsf{inflow}(v, f) = \mathsf{outflow}(v, f)$ for all $v \in V(D) \setminus \{s, t\}$.

The set of flows on D is a vector space that we denote by $F(D)$. A flow is called *integral* if it takes only integer values and we denote by $F(D)_{\mathbb{Z}}$ the group of integral flows on D. The quantity $\delta(f) := \sum_{e_-=s} f(e) - \sum_{e_+=s} f(e)$ is called the *overall throughput* of the flow f.

By a *walk* p in D we understand a sequence $x_0, \dots, x_\ell$ of vertices of D such that $(x_{i-1}, x_i) \in E$ for all $1 \le i \le \ell$. A *path* p in D is defined as a walk such that the vertices $x_0, \dots, x_\ell$ are pairwise distinct. We will say that $x_0, \dots, x_\ell$ are the *vertices used by* p. The path p is called an *s-t-path* if $x_0 = s$ and $x_\ell = t$; p is called a *t-s-path* if $x_0 = t$ and $x_\ell = s$. A sequence $x_0, \dots, x_\ell$ of vertices of D is called a *cycle* c if $x_0, \dots, x_{\ell-1}$ are pairwise distinct, $x_\ell = x_0$, and $(x_{i-1}, x_i) \in E$ for all $1 \le i \le \ell$. Again we say that $x_0, \dots, x_\ell$ are the vertices used by c. We call c a *proper cycle* if c does not use s or t. It will be sometimes useful to identify a path or a cycle with the set of its edges $\{(x_0, x_1), \dots, (x_{\ell-1}, x_\ell)\}$. Since the starting vertex x_0 of a cycle is not relevant, this does not harm. By a *complete path* p *in* D we understand an s-t-path, t-s-path, or a cycle in D. (It is not excluded that the cycle passes through s or t.)

A complete path p in D defines a flow f on D by setting $f(e) := 1$ if $e \in p$ and $f(e) := 0$ otherwise. It will be convenient to denote this flow with p as well. We note that $\delta(p) = 1$ for an s-t-path p, $\delta(p) = -1$ for a t-s-path p, and $\delta(c) = 0$ for a cycle c.

A flow is called *nonnegative* if $f(e) \ge 0$ for all edges $e \in E$. We call $\operatorname{supp}(f) := \{e \in E(D) \mid f(e) \neq 0\}$ the *support* of f.

An important method for analyzing flows is the fact that they can be decomposed into paths and cycles [AMO93].

LEMMA 2.1. *For any nonnegative flow $f \in F(D)$ there exists a family $p_1, \dots, p_m$ of complete paths in D contained in* $\mathrm{supp}(f)$*, and positive real numbers $\alpha_1, \dots, \alpha_m$ such that $f = \sum_{i=1}^m \alpha_i p_i$ Moreover, if the flow f is integral, then the α_i may be assumed to be integers.* □

We will study flows in two rather different situations. The residual digraph R introduced in Section 4 has the property that it never contains an edge (u, v) and its reverse edge (v, u). Only nonnegative flows on R will be of interest.

On the other hand, we also need to look at flows on digraphs resulting from a undirected graph G by replacing each of its undirected edges $\{u, v\}$ by the directed edge $e = (u, v)$ and its reverse $-e := (v, u)$. We shall denote the resulting digraph also by G. To a flow f on G we assign its *reduced representative* $\tilde{f}$ defined by $\tilde{f}(e) := f(e) - \min\{f(e), f(-e)\}$. Hence $\tilde{f}(e) = f(e) - f(-e) \geq 0$ and $\tilde{f}(-e) = 0$ if $f(e) \geq f(-e)$. It will be convenient to interpret f and $\tilde{f}$ as manifestations of the same flow. Formally, we consider the linear subspace $N(G) := \{f \in \mathbb{R}^{E(G)} \mid \forall e \in E(G) : f(e) = f(-e)\}$ of "null flows" and the factor space

$$\overline{F}(G) := F(G)/N(G). \tag{2.1}$$

We call the elements of $\overline{F}(G)$ *flow classes on* G (or simply flows) and denote them by the same symbols as for flows. No confusion should arise from this abuse of notation in the context at hand. We usually identify flow classes with their reduced representative. We note that the overall throughput function factors to a linear function $\delta \colon \overline{F}(G) \to \mathbb{R}$. A flow class is called *integral* if its reduced representative is integral and we denote by $\overline{F}(G)_{\mathbb{Z}}$ the group of integral flow classes on G.

We remark that in the literature on flows, the subtle distinction between flows and their classes is not relevant, as the goal usually is to optimize the throughput of a flow subject to certain capacity constraints. But in the context of LR coefficients, we are interested in *counting* the number of capacity achieving flow classes, so that this distinction is necessary.

2.2. Flows on the honeycomb graph G. We start with a triangular array of vertices, $n+1$ on each side, as seen in Figure 2.1(a). The resulting planar graph Δ shall be called the *triangular graph* with parameter n, we denote its vertex set with $V(\Delta)$ and its edge set with $E(\Delta)$. A triangle consisting of three edges in Δ is called a *hive triangle*. Note that there are two types of hive triangles: upright and downright oriented ones. A *rhombus* is defined to be the union of an upright and a downright hive triangle which share a common side. In contrast to the usual geometric definition of the term *rhombus* we use this term here in this very restricted sense only. Note that the angles at the corners of a rhombus are either acute of 60° or obtuse of 120°. Two distinct rhombi are called *overlapping* if they share a hive triangle.

To realize the dual graph of Δ, as in [Buc00], we introduce a black vertex in the middle of each hive triangle and a white vertex on each hive triangle side, see Figure 2.1(b). Moreover, in each hive triangle T, we introduce edges connecting the three white vertices of T with the black vertex. Additionally (not depicted in Figure 2.1(b)), we introduce a source vertex s and a target vertex t. The source s is connected by an edge with each white vertex v on the right or on the bottom border of Δ, and the target t is connected by an edge with each white vertex v on the left border of Δ. Clearly, the resulting (undirected) graph G is bipartite and planar. We shall call G the *honeycomb graph* with parameter n.

We study now the vector space $\overline{F}(G)$ of flow classes on G introduced in Section 2.1. Recall that for this, we have to replace each edge of G by the corresponding two directed edges. Correspondingly, we will consider G as a directed graph. Any complete path p in the digraph G defines a flow and thus a flow class on G, that we denote by p as well. According to Lemma 2.1 we can write each flow class $f \in \overline{F}(G)$ as a nonnegative linear combination of complete paths. (Note that the reduced representative of any flow on G is nonnegative.)

In order to characterize the flow class $f \in \overline{F}(G)$ in a concise way, we introduce the notion of the throughput of f through edges of Δ. For each edge $k \in E(\Delta)$, there is exactly one upright hive triangle having k as a side: let $e_k \in E(G)$ denote the directed edge in

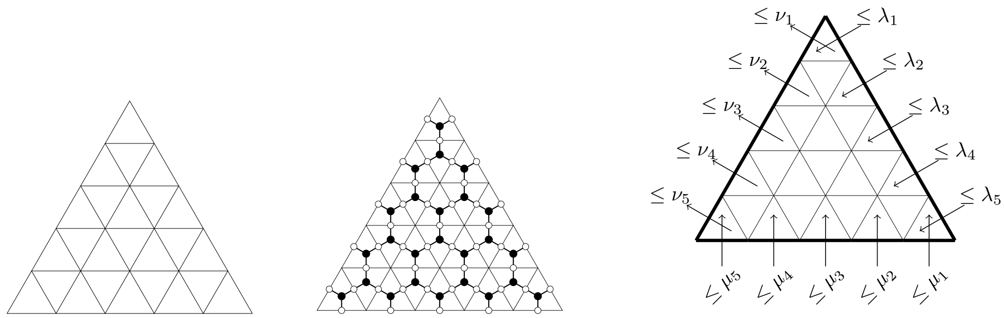

(a) The triangular graph Δ for $n = 5$.

(b) The honeycomb graph G. Source, target and its incident edges are omitted.

(c) The throughput capacities.

Fig. 2.1: Graph constructions.

this triangle pointing from the white vertex on k towards the black vertex in this upright triangle. Then we call $\delta(k, f) := f(e_k) - f(-e_k)$ the *throughput of* f *through* k, which is clearly independent of the choice of the representative. As for the choice of sign: this should be interpreted as the total flow of f *going into the upright hive triangle* through k. Note that $\overline{F}(G) \to \mathbb{R}, f \mapsto \delta(k, f)$ is a linear form.

It is obvious that a flow class f on G is completely determined by the throughput function $\delta\colon E(\Delta) \to \mathbb{R}, k \mapsto \delta(k, f)$. Furthermore, Kirchhoff's conservation laws translate to the closedness condition

$$\delta(k_1, f) + \delta(k_2, f) + \delta(k_3, f) = 0 \tag{2.2}$$

holding for each hive triangle (upright or downright) with sides denoted by k_1, k_2, k_3. So we see that the vector space $\overline{F}(G)$ of flow classes on G can be identified with the subvector space $Z \subseteq \mathbb{R}^{E(\Delta)}$ consisting of the functions δ satisfying (2.2) for all hive triangles. Moreover, under this identification, integral flow classes f correspond to functions in the subgroup $Z_{\mathbb{Z}}$ consisting of functions δ taking integer values.

By adding up (2.2) for all upright hive triangles and subtracting (2.2) for all downright hive triangles, taking into account the cancelling of throughputs on all inner sides k, we see that the sum of $\delta(k, f)$ over all border edges k of Δ vanishes. Therefore, we can express the overall throughput $\delta(f)$ as

$$\delta(f) \; = \sum_{k \in E_r \cup E_b} \delta(k, f) \; = \; - \sum_{k' \in E_\ell} \delta(k', f), \tag{2.3}$$

where E_ℓ, E_r, and E_b denotes the set of edges of Δ on the left side, right side, and bottom side, respectively.

The flow classes on G can be characterized in yet another way. Let x_0 be the top vertex of Δ and define the vector space H of functions $h\colon V(\Delta) \to \mathbb{R}$ satisfying $h(x_0) = 0$. We denote by $H_{\mathbb{Z}}$ the subgroup of functions $h \in H$ taking integer values.

For a moment, think of the edges k of Δ as oriented such that all upright hive triangles get clockwise oriented. Consider the linear map $\partial\colon H \to \mathbb{R}^{E(\Delta)}, h \mapsto \delta$ defined by $\delta(k) = h(k_+) - h(k_-)$, where k points from k_- to k_+. Then it is obvious that ∂ is injective, and it is straightforward to check that $\mathrm{im}\partial \subseteq Z$. In order to show equality, suppose $\delta \in Z$. For a vertex $x \in V(\Delta)$, choose a directed path p (in the sense of the above chosen orientations) from the top vertex x_0 to x. The closedness condition (2.2) easily implies that the sum $h(x) := \sum_{k \subseteq p} \delta(k)$ is independent of the choice of p. It follows that $\partial(h) = \delta$.

So we have a linear isomorphism $\partial\colon H \to Z$, which induces an isomorphism $H_{\mathbb{Z}} \to Z_{\mathbb{Z}}$.

Remark 2.2. *The reader familiar with basic algebraic topology will recognize ∂ as a coboundary map of the simplicial complex provided by Δ, and hence* $\mathrm{im}\partial = Z$ *as a consequence of the fact that the triangle underlying Δ is simply connected.*

2.3. Hives and hive flows. Following [KT99, Buc00] we define a *hive* on Δ as a function $h \in H$ such that for all rhombi ϱ, the sum of the values of h at the two obtuse vertices of ϱ is greater than or equal to the sum of the values of h at the two acute vertices of ϱ. In pictorial notation,

$$\sigma(\Diamond, h) := h(\bullet\!\Diamond) + h(\Diamond\!\bullet) - h(\overset{\bullet}{\Diamond}) - h(\underset{\bullet}{\Diamond}) \;\geq\; 0 \tag{2.4}$$

where $\bullet\!\Diamond, \Diamond\!\bullet, \overset{\bullet}{\Diamond}, \underset{\bullet}{\Diamond} \in V(\Delta)$ denote the corner vertices of $\Diamond$. We call the $\sigma(\varrho, h)$ the *slack* of the rhombus ϱ with respect to the hive h.

If one interprets $h(v)$ as the height of a point over $v \in V(\Delta)$ and interpolates these points linearly over each hive triangle of Δ one gets a continuous function $h\colon \Delta \to \mathbb{R}$. (Here the triangle Δ is to be interpreted as a convex subset of $\mathbb{R}^2$.) Then the conditions (2.4) mean that h is a concave function. The function h is linear over a rhombus ϱ iff $\sigma(\varrho, h) = 0$, in which case we call the rhombus ϱ *h-flat*.

Lemma 2.3. *For a hive $h \in H$ and $x \in V(\Delta)$ we have $\min_{\partial\Delta} h \;\leq\; h(x) \;\leq\; n \max_{\partial\Delta} h$, where $\partial\Delta$ denotes the boundary of the convex set $\Delta \subseteq \mathbb{R}^2$.*

Proof. Let $x(m, i)$ denote the vertex of Δ in the mth line parallel to the ground side (counting from the top) and on the ith side parallel to the left side (counting from the left), for $0 \leq i \leq m \leq n$. So $x(0, 0)$ is the top vertex and $h(x(0, 0)) = 0$ for $h \in H$. Put $a := h(x(1, 0))$ and $b := h(x(1, 1))$.

Since h is a concave function, its subgraph $S := \{(x, y) \in \Delta \times \mathbb{R} \mid y \leq h(x)\}$ is convex. Hence S is bounded from above by the plane spanned by $((0, 0), 0)$, $((1, 0), a)$, and $((1, 1), b)$. This plane's height at $x(m, i)$ is $am + (b - a)i$. This implies that $h(x(m, i)) \leq am + (b - a)i$ for all $h \in H$. Therefore, $h(x(m, i)) \leq m \max\{a, b\}$, proving the upper bound.

The lower bound follows easily from the convexity of S. $\square$

In this paper, it will be extremely helpful to have some graphical way of describing rhombi and throughputs. We shall denote a rhombus ϱ of Δ by the pictogram $\Diamond$, even though ϱ may lie in any of the three positions "$\Diamond$", "$\angle\!\!\!\!\!\!\diagup$" or "$\diagdown\!\!\!\!\!\!\angle$" obtained by rotating with a multiple of 60°. Let $\ominus\!\!\!\!\Diamond$ denote the edge k of Δ given by the *diagonal* of ϱ connecting its two obtuse angles. Then we denote by $\overset{\uparrow}{\ominus}(f) := \delta(k, f)$ the throughput of f through k (going into the upright hive triangle). Similarly, we define the throughput $\overset{\downarrow}{\ominus}(f) := -\delta(k, f)$. The advantage of this notation is that if the throughput is positive, then the flow goes in the direction of the arrow. For instance, using the symbolic notation, we note the following consequence of the flow conservation laws:

$$\swarrow\!\!\Diamond(f) + \Diamond\!\!\swarrow(f) \;=\; \nwarrow\!\!\Diamond(f) + \Diamond\!\!\searrow(f). \tag{2.5}$$

If f is the flow corresponding to the hive $h \in H$ under the isomorphisms $H \simeq Z \simeq \overline{F}(G)$, then (2.4) and the definition of the coboundary map ∂ imply that

$$\sigma(\Diamond, h) \;=\; \Big(h(\bullet\!\Diamond) - h(\overset{\bullet}{\Diamond})\Big) + \Big(h(\Diamond\!\bullet) - h(\underset{\bullet}{\Diamond})\Big) \;=\; \nwarrow\!\!\Diamond(f) + \Diamond\!\!\swarrow(f).$$

We define now the slack of a rhombus with respect to a flow f as the slack with respect to the corresponding hive h.

Definition 2.4. *The* slack *of the rhombus $\Diamond$ with respect to $f \in \overline{F}(G)$ is defined as*

$$\sigma(\Diamond, f) \;:=\; \nwarrow\!\!\Diamond(f) + \Diamond\!\!\swarrow(f).$$

*The rhombus $\Diamond$ is called f-*flat *if $\sigma(\Diamond, f) = 0$.*

It is clear that $\overline{F}(G) \to \mathbb{R}, f \mapsto \sigma(\varrho, f)$ is a linear form. Note also that by (2.5), the slack can be written in varous different ways:

$$\sigma(\Diamond, f) \;=\; \nwarrow\!\!\Diamond(f) + \Diamond\!\!\swarrow(f) \;=\; \nwarrow\!\!\Diamond(f) - \Diamond\!\!\nearrow(f) \;=\; \swarrow\!\!\Diamond(f) - \Diamond\!\!\nwarrow(f) \;=\; \swarrow\!\!\Diamond(f) + \Diamond\!\!\searrow(f).$$

Definition 2.5. *A flow $f \in \overline{F}(G)$ is called a* hive flow *iff $\sigma(\varrho, f) \geq 0$ for all rhombi ϱ in Δ.*

By definition, the hives correspond to the hive flows under the isomorphism $H \simeq \overline{F}(G)$. Note that the set of hive flows is a cone in $\overline{F}(G)$. Figure 2.2 provides an example of a hive flow. We encourage the reader to verify the slack inequalities there to get some idea of the nature of these constraints.

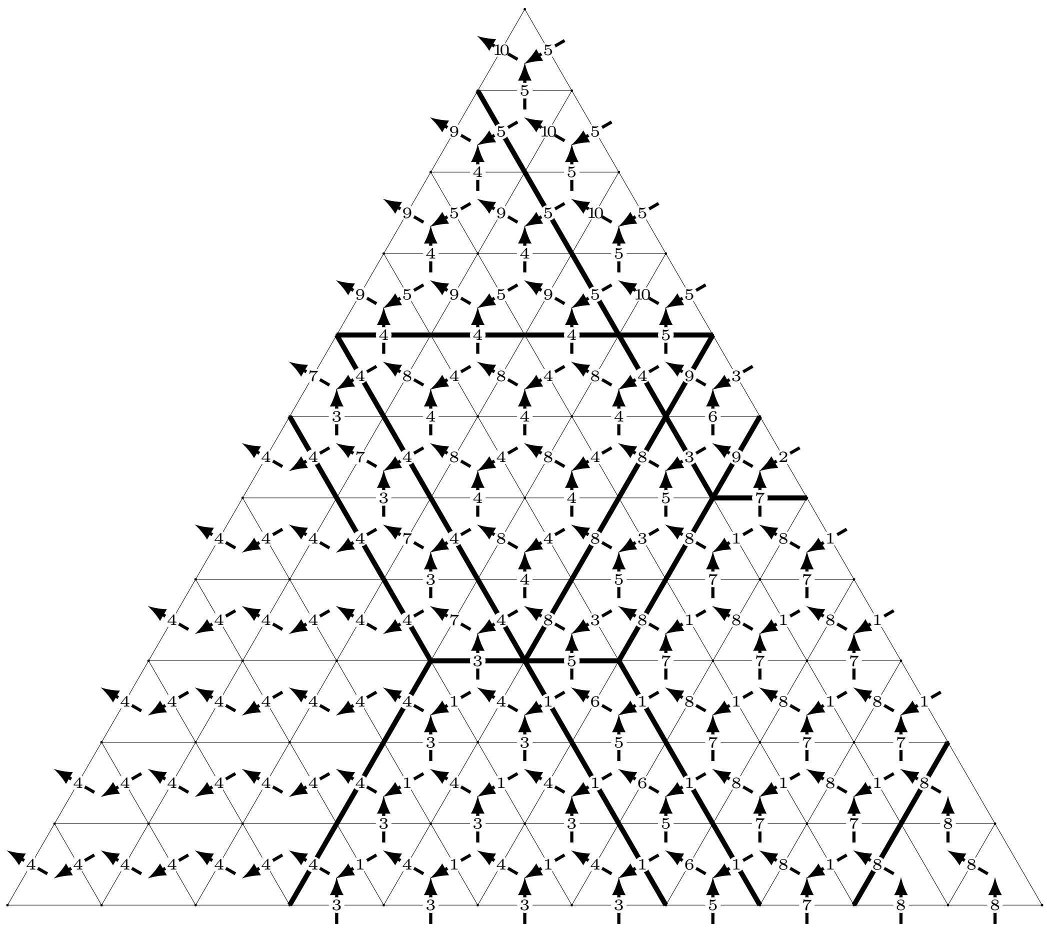

Fig. 2.2: A hive flow for $n = 11$, $\lambda = (5,5,5,5,3,2,1,1,1,0,0)$, $\mu = (8,8,7,5,3,3,3,3,0,0,0)$, $\nu = (10,9,9,9,7,4,4,4,4,4,4)$, and its corresponding partition of Δ into flatspaces. The numbers give the throughputs through edges of Δ in the directions of the arrows. The properties of Lemma 4.11 and Lemma 4.13 are readily verified.

We formulate now capacity constraints for the throughputs of hive flows at the border of Δ, depending on a chosen triple $\lambda, \mu, \nu \in \mathbb{N}^n$ of partitions satisfying $|\nu| = |\lambda| + |\mu|$. Hereby, we treat the left border of Δ differently from the right and bottom border of Δ with regard to orientations. To the ith border edge k of Δ on the right border of Δ, counted from top to bottom, we assign the *throughput capacity* $b(k) := \lambda_i$, see Figure 2.1(c). Further, we set $b(k) := \mu_i$ for the ith edge k on the bottom border of Δ, counted from right to left. Finally, we set $b(k') := \nu_i$ for the ith edge k' on the left border of Δ, counted from top to bottom. Recall that $\delta(k, f)$ denotes the throughput of a flow f into Δ, while $-\delta(k', f)$ denotes the throughput of f out of Δ.

Definition 2.6. *Let $\lambda, \mu, \nu \in \mathbb{N}^n$ be a triple of partitions satisfying $|\nu| = |\lambda| + |\mu|$. The* polytope of bounded hive flows $B := B(\lambda, \mu, \nu) \subseteq \overline{F}(G)$ *is defined to be the set of hive flows $f \in \overline{F}(G)$ satisfying*

$$0 \leq \delta(k, f) \leq b(k) \quad \textit{and} \quad 0 \leq -\delta(k', f) \leq b(k')$$

for all border edges k on the right or bottom border of Δ, and for all border edges k' on the left border of Δ. The polytope of capacity achieving hive flows $P := P(\lambda, \mu, \nu)$ *consists of those $f \in B(\lambda, \mu, \nu)$ for which $\delta(k, f) = b(k)$ and $-\delta(k', f) = b(k')$ for all k and k' as above. We also set $B_{\mathbb{Z}} := B \cap \overline{F}(G)_{\mathbb{Z}}$ and $P_{\mathbb{Z}} := P \cap \overline{F}(G)_{\mathbb{Z}}$.*

Lemma 2.3 and the isomorphism $\overline{F}(G)_{\mathbb{Z}} \simeq H_{\mathbb{Z}}$ imply that B is bounded and thus B and P are indeed polytopes.

We note that by (2.3), we have $\delta(f) \leq |\nu|$ for any $f \in B(\lambda, \mu, \nu)$. Moreover, $f \in B(\lambda, \mu, \nu)$ is capacity achieving iff $\delta(f) = |\nu|$.

Knutson and Tao [KT99] (see also [Buc00]) characterized the Littlewood–Richardson coefficient $c^{\nu}_{\lambda,\mu}$ as the number of integral hives taking fixed values on the border vertices of Δ, prescribed by the partitions λ, μ, ν. Their description via the isomorphism $\overline{F}(G)_{\mathbb{Z}} \simeq H_{\mathbb{Z}}$ immediately translates to the following fundamental result.

Proposition 2.7. *The Littlewood–Richardson coefficient $c^{\nu}_{\lambda,\mu}$ equals the number of capacity achieving integral hive flows, i.e., $c^{\nu}_{\lambda,\mu} = |P(\lambda, \mu, \nu)_{\mathbb{Z}}|$.* □

To advocate the advantage of the flow interpretation of Littlewood–Richardson coefficients, we show in the next section that $P_{\mathbb{Z}} := P(\lambda, \mu, \nu)_{\mathbb{Z}}$ can be interpreted as the set of vertices of a graph in a natural way. This will be important for searching and enumerating $P_{\mathbb{Z}}$ in an efficient way. Our investigations will be purely structural though. We leave the (more complicated) algorithmic aspects of searching to the forthcoming paper [Ike12].

3. Properties of hive flows. We recall that any complete path p defines a flow on G, denoted by the same symbol. In order to describe the slack of a rhombus with respect to p, we introduce some further terminology.

Definition 3.1. *A* turn *is defined to be a path in G of length 2 that lies inside Δ, starts at a white vertex and ends with a different white vertex, see Figure* 2.1(b).

Note that there are six turns in each hive triangle. We shall denote turns pictorially by $\Diamond$, $\Diamond$ etc. with the obvious interpretation. Similarly, $\Diamond$ and $\Diamond$ stand for a path consisting of four edges.

In order to describe the different ways a complete path p may pass a rhombus ϱ, we consider the following sets of paths in ϱ.

Definition 3.2. *The sets of paths, interpreted as subsets of $E(G)$,*

$$\Psi_+(\Diamond) := \{\Diamond, \Diamond, \Diamond, \Diamond\}, \quad \Psi_-(\Diamond) := \{\Diamond, \Diamond, \Diamond, \Diamond\}, \ \textit{and} \ \Psi_0(\Diamond) := \{\Diamond, \Diamond, \Diamond, \Diamond\}$$

are called the sets of of positive*,* negative*, and* neutral slack contributions *of the rhombus $\Diamond$, respectively.*

For later use the reader should remember that the turns in $\Psi_+(\Diamond)$ at the acute angles are clockwise, while the concatenations of two turns at the obtuse angles are counterclockwise.

The verification of the following is immediate using Definition 2.4 of the slack.

Observation 3.3. *Let p be a complete path in G and E_{ϱ} be the set of edges of G contained in a rhombus ϱ. Then $p \cap E_{\varrho}$ is either empty, or it is a union of one or two slack contributions q. The slack $\sigma(\varrho, p)$ is obtained by adding* 1*,* 0*, or* -1 *over the contributions q contained in p, according to whether q is is positive, negative, or neutral.*

We remark that the only situations, in which $p \cap E_{\varrho}$ is a union of two slack contributions, is when p uses both counterclockwise turns $\Diamond$ and $\Diamond$ at acute angles, or both clockwise turns $\Diamond$ and $\Diamond$ at acute angles, in which case $\sigma(\varrho, p) = -2$ or $\sigma(\varrho, p) = 2$, respectively. It is not possible that c uses both $\Diamond$ and $\Diamond$ since otherwise, due to the planarity of Δ, c would have to intersect itself.

3.1. The support of flows on G. Recall the definition of the support $\mathrm{supp}(d)$ of a flow class $d \in \overline{F}(G)$. By the definition, $\mathrm{supp}(d)$ cannot contain an edge and its reverse. We note the following:

$$\big(◇ \subseteq \mathrm{supp}(d) \ \text{ or } \ ◇ \subseteq \mathrm{supp}(d)\big) \quad \Longleftrightarrow \quad ◇(d) > 0.$$

Recall from Definition 3.2 the sets $\Psi_+(\varrho)$, $\Psi_-(\varrho)$, and $\Psi_0(\varrho)$ of positive, negative, and neutral slack contributions of a rhombus ϱ, respectively, interpreted as sets of directed edges of G. We assign to any slack contribution $p \in \Psi_+(\varrho) \cup \Psi_-(\varrho)$ of a rhombus ϱ its *antipodal contribution* $p' \in \Psi_+(\varrho) \cup \Psi_-(\varrho)$, which is defined by reversing p and then applying a rotation of 180°. For instance, ◇ is the antipodal contribution of ◇ and ◇ is the antipodal contribution of ◇. Clearly, $p \mapsto p'$ is an involution.

The following lemma on antipodal contributions will be of great use.

Lemma 3.4. *Let $d \in \overline{F}(G)$ such that $\sigma(\varrho, d) \geq 0$ for a rhombus ϱ. If $p \subseteq \mathrm{supp}(d)$ for a negative slack contribution p of ϱ, then $p' \subseteq \mathrm{supp}(d)$ for its antipodal contribution p'.*

Proof. 1. Suppose that ◇ $\subseteq \mathrm{supp}(d)$, which means $\delta_1 :=$ ◇$(d) > 0$ and $\delta_2 :=$ ◇$(d) > 0$. Since $\delta_3 :=$ ◇$(d) -$ ◇$(d) = \sigma($◇$, d) \geq 0$ we get ◇$(d) = \delta_1 + \delta_3 > 0$. Moreover, ◇$(d) =$ ◇$(d) -$ ◇$(d) = (\delta_1 + \delta_3) - (\delta_1 - \delta_2) = \delta_3 + \delta_2 > 0$. Altogether, ◇ $\subseteq \mathrm{supp}(d)$.

2. Suppose that ◇ $\subseteq \mathrm{supp}(d)$, which means $\delta_1 :=$ ◇$(d) > 0$, $\delta_2 :=$ ◇$(d) > 0$, and $\delta_3 :=$ ◇$(d) > 0$. Hence ◇$(d) = \delta_3 - \delta_1$. We have $\delta_4 :=$ ◇$(d) -$ ◇$(d) = \sigma($◇$, d) \geq 0$ and thus ◇$(d) = \delta_1 + \delta_4 > 0$. Therefore $\delta_3 = \delta_2 + (\delta_1 + \delta_4)$ and thus ◇$(d) = \delta_3 - \delta_1 = (\delta_2 + \delta_1 + \delta_4) - \delta_1 = \delta_2 + \delta_4 > 0$. Altogether, ◇ $\subseteq \mathrm{supp}(d)$. ∎

Applying Lemma 3.4 successively can provide important information about the support of a flow class d. This is stated in the following lemma on "flow propagation".

It will be convenient to use symbols like ◇, ◇ etc., which stand for the rhombi in the positions relative to ◇ as indicated by the shaded regions.

Lemma 3.5. *Given $d \in \overline{F}(G)$ such that $\sigma($◇$, d) \geq 0$ and ◇ $\subseteq \mathrm{supp}(d)$. Then ◇ $\subseteq \mathrm{supp}(d)$. If additionally $\sigma($◇$, d) \geq 0$, then ◇ $\subseteq \mathrm{supp}(d)$. Similarly, if additionally $\sigma($◇$, d) \geq 0$, then ◇ $\subseteq \mathrm{supp}(d)$.*

For an example on how Lemma 3.5 can be used, see Figure 3.1.

Fig. 3.1: The rhombi of the pentagon have nonnegative slack with respect to the flow d. If the turns in the left picture are in $\mathrm{supp}(d)$, then, by applying Lemma 3.5 several times, we see that all the turns in the right picture are in $\mathrm{supp}(d)$.

Proof of Lemma 3.5. The first assertion is a direct application of Lemma 3.4. Suppose that $\sigma($◇$, d) \geq 0$. Since ◇ $\subseteq \mathrm{supp}(d)$, flow conservation implies that ◇ $\subseteq \mathrm{supp}(d)$ or ◇ $\subseteq \mathrm{supp}(d)$. We want to show ◇ $\subseteq \mathrm{supp}(d)$. If ◇ $\subseteq \mathrm{supp}(d)$, then ◇ $\subseteq \mathrm{supp}(d)$ and ◇ is a negative contribution in ◇. Hence by Lemma 3.4, we have ◇ $\subseteq \mathrm{supp}(d)$. The other assertion is proved analogously. ∎

3.2. The graph of capacity achieving integral hive flows. Fix λ, μ, ν and recall the polytopes B and P from Definition 2.6. We show now that $P_{\mathbb{Z}}$ can be naturally seen as the vertex set of a graph.

Definition 3.6. *We say that $f, g \in P_{\mathbb{Z}}$ are* neighbours *iff $g - f$ is a cycle in G. The resulting graph with the set of vertices $P_{\mathbb{Z}}$ is also denoted by $P_{\mathbb{Z}}$.*

The neighbour relation is clearly symmetric. We also remark that a cycle of the form $g - f$ must be proper, i.e., it neither uses the source or target. The reason is that the flow $g - f$ vanishes on the edges touching the border of Δ, as f and g are both capacity achieving.

For an explicit characterization of the neighbour relation we need the following concepts.

Definition 3.7. *Let $f \in B$ and c be a proper cycle in G.*

1. *We call a rhombus* ϱ nearly f-flat *iff* $\sigma(\varrho, f) = 1$.

2. c *is called* f-hive preserving *iff* c *does not use negative contributions in* f*-flat rhombi.*

3. c *is called* f-secure *iff* c *is* f*-hive preserving and* c *does not use both counterclockwise turns at acute angles in nearly* f*-flat rhombi ( ).*

We remark that c is f-hive preserving iff $f + \varepsilon c \in B$ for sufficiently small $\varepsilon > 0$.

Proposition 3.8. *Assume* $f \in P_{\mathbb{Z}}$. *If* $g \in P_{\mathbb{Z}}$ *is a neighbour of* f, *then* $g - f$ *is an* f*-secure cycle. Conversely, if* c *is an* f*-secure cycle, then* $f + c \in P_{\mathbb{Z}}$ *is a neighbour of* f.

Proof. Assume that f and g are neighbours in $P_{\mathbb{Z}}$, so $c := g - f$ is a proper cycle in G. Hence $\sigma(\varrho, f + c) \geq 0$ for each rhombus ϱ. This implies that $\sigma(\varrho, c) \geq 0$ for each f-flat rhombus, that is, c is f-hive preserving. Moreover, if $\sigma(\varrho, f) = 1$, then $\sigma(\varrho, c) \geq -1$. Hence c is f-secure. The argument can be reversed. □

Apparently, the symmetry of the neighbour relation in $P_{\mathbb{Z}}$ does not seem to be obvious from the characterization in Proposition 3.8.

Before continuing, we state a useful observation. The union of two overlapping rhombi ϱ_1 and ϱ_2 forms a trapezoid. Glueing together two such trapezoids (ϱ_1, ϱ_2) and (ϱ_1', ϱ_2') at their longer side, we get a hexagon. The verification of the following *hexagon equality* is straightforward and left to the reader: $\sigma(\varrho_1, f) + \sigma(\varrho_2, f) = \sigma(\varrho_1', f) + \sigma(\varrho_2', f)$ for any flow $f \in \overline{F}(G)$. In pictorial notation, the hexagon equality can be succinctly expressed as

$$\sigma(\;\,, f) + \sigma(\;\,, f) = \sigma(\;\,, f) + \sigma(\;\,, f), \tag{3.1}$$

As an immediate consequence we obtain the following.

Corollary 3.9. *For all hive flows* f, *if* and *are* f*-flat, then also* and *are* f*-flat.* □

Remark 3.10. *The slacks of rhombi are exactly the numbers in Berenstein-Zelevinsky triangles [PV05] and the hexagon equality (*3.1*) is just the condition for their validity.*

The next result tells us how f-secure cycles may arise.

Theorem 3.11. *Let* $f \in B_{\mathbb{Z}}$ *and* c *be an* f*-hive preserving cycle in* G *of minimal length. Then* c *is* f*-secure.*

Proof. We argue by contradiction. Suppose that c is an f-hive preserving cycle in G of minimal length, but not f-secure. So there is a nearly f-flat rhombus in which c uses both turns and . Let us call such rhombi *bad*.

Since c has minimal length, it cannot be rerouted via . Hence we cannot be in the following case, in which c can easily be rerouted via :

(is not f-flat or c uses) and (is not f-flat or c uses).

In the remaining two cases

(A) is f-flat and c uses or (B) is f-flat and c uses (3.2)

the f-hive preserving cycle c cannot be rerouted via by Definition 3.7(2).

Let us assume that we are in the situation (A). So we have the bad, nearly f-flat rhombus and the shaded f-flat rhombus. The hexagon equality (3.1) implies that either is f-flat and is nearly f-flat, or is nearly f-flat and is f-flat. These two possibilities are indicated on the left and right side of the following picture, respectively, where the shaded rhombi are f-flat and the diagonals of nearly f-flat rhombi are drawn thick. Further, parts of c which run in f-flat rhombi, are drawn with straight arrows:

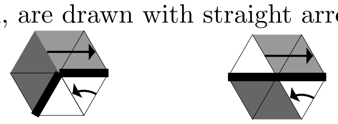

The fact that c uses no negative contributions in f-flat rhombi (and no vertex of G twice) forces c to run exactly as depicted in the following picture:

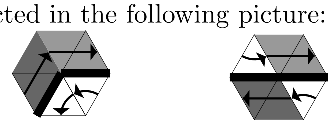

Hence the second nearly f-flat rhombus in the left and right picture, respectively, is bad as well.

So we see that the diagonal $\Diamond$ of the bad rhombus $\Diamond$ shares a vertex $\bullet\!\Diamond$ with the diagonal of another bad rhombus and that both diagonals either lie on the same line or include an angle of $120°$. By symmetry, the same conclusion can be drawn in the case (B).

By induction, this implies that there is a region bounded by diagonals of bad rhombi. This is impossible, because c would have to run both inside and outside of this region. □

The following is an important insight into the structure of $P_{\mathbb{Z}}$. We postpone the proof to Section 8.

Theorem 3.12 (Connectedness Theorem). *The graph $P_{\mathbb{Z}}$ is connected.*

As an application of our insights, we obtain the following characterization of multiplicity freeness. Recall that $c^{\nu}_{\lambda,\mu} = |P(\lambda,\mu,\nu)_{\mathbb{Z}}|$.

Proposition 3.13. *Suppose that $f \in P(\lambda,\mu,\nu)_{\mathbb{Z}}$. Then we have $c^{\nu}_{\lambda,\mu} > 1$ iff there exists an f-hive preserving cycle in G.*

Proof. If there exists an f-hive preserving cycle in G, then there is also one of minimal length, call it c. Theorem 3.11 implies that c is f-secure. Proposition 3.8 tells us that $f + c \in P_{\mathbb{Z}}$. It follows that $|P_{\mathbb{Z}}| \geq 2$.

Conversely, assume that $|P_{\mathbb{Z}}| \geq 2$. Since $f \in P_{\mathbb{Z}}$ and $P_{\mathbb{Z}}$ is connected by Theorem 3.12, there exists a neighbour $g \in P_{\mathbb{Z}}$ of f. Proposition 3.8 tells us that $g - f$ is an f-secure cycle. □

A proof of Fulton's conjecture, first shown in [KTW04] by different methods, is obtained as an easy consequence.

Corollary 3.14. *If $c^{\nu}_{\lambda,\mu} = 1$, then $c^{N\nu}_{N\lambda,N\mu} = 1$ for all $N \geq 1$.*

Proof. By definition, c is an f-hive preserving cycle in G iff c is an Nf-hive preserving cycle in G. Now apply Proposition 3.13. □

The characterization in Proposition 3.13 points to a way of algorithmically deciding whether $c^{\nu}_{\lambda,\mu} > 1$. However, it is not obvious how to efficiently search for f-hive preserving cycles in the graph G. For this, and even for the simpler task of deciding $c^{\nu}_{\lambda,\mu} > 0$, we have to construct suitable "residual digraphs", which brings us to the topic of the next section. More details on the complexity of testing $c^{\nu}_{\lambda,\mu} > 1$ can be found in [Ike12].

4. The residual digraph R_f. Proposition 2.7 suggests to decide $c^{\nu}_{\lambda,\mu} > 0$ by solving the problem of optimizing the linear (overall throughput) function δ on the polytope $B = B(\lambda,\mu,\nu)$ of bounded hive flows. In fact, we will show later that the optimum is always obtained at an integral flow. Maximizing a certain linear functional on the hive polytope and showing that the maximum is attained at an integer point is the basic idea in [KT99, Buc00]. However, they do not present any algorithmic result.

We follow a Ford-Fulkerson approach and try to construct for a given integral hive flow $f \in B_{\mathbb{Z}}$ a digraph R_f, such that adding an s-t-path p in R_f to f leads to a bounded hive flow. We have to guarantee that $f + p$ does not lead to negative slacks of rhombi so that $f + p$ is a hive flow. On the other hand, we want to make sure that f is optimal, when there is no s-t-path in R_f.

4.1. Turnpaths and turncycles. The intuition is to consider paths in G in which each node remembers its predecessor. This can be formally achieved by studying paths in an auxiliary digraph that we define next.

Recall that a turn is a path in G of length 2 that lies inside Δ, starts at a white vertex and ends with a different white vertex.

Definition 4.1. *A* turnedge *is an ordered pair of turns that can be concatenated to a path in G.*

Note that a turnedge defines a path in G of length 4. We write turnedges pictorially like $\Diamond := (\Diamond, \Diamond)$ etc. We construct now the auxiliary digraph R.

Definition 4.2. *The* digraph *R has as vertices the turns, henceforth called* turnvertices, *and the source and target of G. The edges of R are the turnedges and the following additional edges: the digraph R contains an edge (s,ϑ) from the source s to any turnvertex ϑ starting* *at the right or bottom border of Δ. Vice versa, for any turnvertex ϑ' pointing at the right or bottom border of Δ, there is a turnedge (ϑ', s) in R. Similarly, for the target t, there are edges (ϑ, t) for each turnvertex ϑ pointing at the left border of Δ and vice versa, there are edges (t, ϑ') for each turnvertex ϑ' starting at the left border of Δ.*

The reader should check that R never contains an edge and its reverse. In fact, the digraph R is rather complicated, for instance one can show that it is not planar for $n \geq 2$.

We have a well-defined notion of flows on R as R is a digraph with two distinguished vertices s and t. We assign now to a flow f on R a flow class $\tilde{f}$ on G by defining the corresponding throughput map $E(\Delta) \to \mathbb{R}, k \mapsto \delta(k, \tilde{f})$ as follows. An edge k of Δ lies in exactly one upright hive triangle $\triangle$. Let $\triangle$ and $\triangle$ denote the two turns in $\triangle$ pointing towards the white vertex on k. Further, let $\triangle$ and $\triangle$ denote the turns obtained when reversing $\triangle$ and $\triangle$. We define

$$\delta(k, \tilde{f}) \; = \; \diamondsuit(f) \; := \; \mathsf{inflow}(\triangle, f) + \mathsf{inflow}(\triangle, f) - \mathsf{outflow}(\triangle, f) - \mathsf{outflow}(\triangle, f). \qquad (4.1)$$

More explicitly,

$$\diamondsuit(f) = f(\lozenge) + f(\lozenge) + f(\lozenge) + f(\lozenge) - f(\lozenge) - f(\lozenge) - f(\lozenge) - f(\lozenge).$$

From (4.1) it is straightforward to check that the closedness condition (2.2) is satisfied in the hive triangle $\triangle$.

Therefore, the flow class $\tilde{f} \in \overline{F}(G)$ is well defined by (4.1). So we have defined the linear map

$$\pi \colon F(R) \to \overline{F}(G), f \mapsto \tilde{f} \qquad (4.2)$$

that moreover maps integral flows to integral flows.

Let us stress that we are only interested in the *cone $K(R)$ of nonnegative flows* on R. We define the *slack* $\sigma(\varrho, f)$ of rhombus ϱ with respect to a flow $f \in K(R)$ by $\sigma(\varrho, f) := \sigma(\varrho, \pi(f))$. Similarly, we define the *throughput* $\diamondsuit(f) := \diamondsuit(\pi(f))$ of f through an edge k, and we call $\delta(f) := \delta(\pi(f))$ the *overall throughput* $\delta(f)$ of $f \in K(R)$.

For the sake of clarity, paths and cycles in R shall be called *turnpaths* and *turncycles*. Correspondingly, we have the notions of s-t-turnpaths, t-s-turnpaths. By a *complete turnpaths* we understand an s-t-turnpaths, a t-s-turnpaths, or a turncycle (which may pass through s or t). A complete turnpath p defines a flow on R, again denoted by p, by putting the flow value of 1 on each turnedge used.

Example 4.3. *The flow $\pi(p)$ induced by a complete turnpath p in R is not necessarily given by a complete path on G. E.g., it is possible that p uses both turnedges $\lozenge$ and $\lozenge$ (which do not share a turnvertex). Then $\diamondsuit(\pi(p)) = 2$, while for a complete path q of G we always have $\diamondsuit(q) \in \{-1, 0, 1\}$.*

If p is a complete turnpath in R and x is a turnvertex or turnedge, then it will be convenient to write $\mathbf{1}_p(x) = 1$ if x occurs p, and $\mathbf{1}_p(x) = 0$ otherwise.

We reconsider now Definition 3.2 and interpret the sets $\Psi_+(\lozenge)$, $\Psi_-(\lozenge)$, and $\Psi_0(\lozenge)$ of slack contributions of a rhombus $\lozenge$ —instead of as subsets of $E(G)$— as subsets of $V(R) \cup E(R)$. Note that these sets are pairwise disjoint.

Lemma 4.4. *Let p be a complete turnpath in R. Then we have for any rhombus ϱ,*

$$\sigma(\varrho, p) = \sum_{x \in \Psi_+(\varrho)} \mathbf{1}_p(x) - \sum_{x \in \Psi_-(\varrho)} \mathbf{1}_p(x).$$

Moreover, $\sigma(\varrho, p) \in \{-4, -3, \ldots, 3, 4\}$. □

Proof. For each rhombus $\Diamond$ we have in pictorial notation $\mathbf{1}_p(\Diamond)+\mathbf{1}_p(\Diamond)=\mathbf{1}_p(\Diamond)+\mathbf{1}_p(\Diamond)$. Moreover,

$$\begin{aligned}
\sigma(\Diamond,p) &= \Diamond(p)+\Diamond(p)\\
&= \big(\mathbf{1}_p(\Diamond)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)+\big(\mathbf{1}_p(\Diamond)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)\\
&= \big(\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)+\big(\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\\
&= \big(\mathbf{1}_p(\Diamond)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\\
&\qquad+\big(\mathbf{1}_p(\Diamond)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\\
&= \big(\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)+\big(\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\big)+\mathbf{1}_p(\Diamond)-\mathbf{1}_p(\Diamond)\\
&= \sum_{x\in\Psi_+(\Diamond)}\mathbf{1}_p(x)-\sum_{x\in\Psi_-(\Diamond)}\mathbf{1}_p(x).
\end{aligned}$$

The assertion on the possible values of $\sigma(\varrho,p)$ follows immediately. ∎

Example 4.5. *There are complete turnpaths p and q, that use all the turnvertices in $\Psi_+(\varrho)$ and $\Psi_-(\varrho)$, respectively, resulting in the slacks $\sigma(\varrho,p)=4$ and $\sigma(\varrho,q)=-4$.*

We construct now the digraph R_f from R by deleting the negative slack contributions in f-flat rhombi, and removing all edges of R crossing capacity achieving edges of Δ at the border of Δ. Recall from Definition 2.6 the definition of the throughput capacities $b(k)$ of the border edges of Δ, given by λ,μ,ν.

Definition 4.6. *Let $f\in B(\lambda,\mu,\nu)$. The* residual digraph $R_f:=R_f(\lambda,\mu,\nu)$ *is obtained from R by deleting the turnvertices and turnedges in $\Psi_-(\varrho)$ in f-flat rhombi ϱ. Moreover, for all edges k on the right and bottom border of Δ satisfying $\delta(k,f)=b(k)$, we delete all four edges of R crossing k. Similarly, for all edges k' on the left border of Δ satisfying $-\delta(k',f)=b(k')$, we delete all four edges of R crossing k'. Let $\mathcal{P}_f$ denote the set of* complete turnpaths in R_f.

The following is an immediate consequence of the construction of R_f and Lemma 4.4.

Lemma 4.7. *We have $\sigma(\varrho,p)\geq 0$ for any $p\in\mathcal{P}_f$ and any f-flat rhombus ϱ.* ∎

We denote by $K(R_f)$ the cone of nonnegative flows on R_f. As R_f is a subgraph of R, a nonnegative flow $f\in K(R_f)$ can be interpreted as a flow $f\in K(R)$ with value zero on the turnedges not present in R_f.

Let $f\in B_{\mathbb{Z}}$ and p be an s-t-turnpath in R_f. Is $f+\pi(p)\in B_{\mathbb{Z}}$?

By construction of R_f, if p crosses the border edge k, then $\delta(k,f)<b(k)$ if k is on the right or bottom border of Δ. Similarly, $-\delta(k',f)<b(k')$ if k' is on the left border of Δ. Thus the flow $f+\pi(p)$ does not violate the border capacity constraints.

In order to see whether $f+\pi(p)$ is a hive flow, we note that if ϱ is an f-flat rhombus, then $\sigma(\varrho,f+\pi(p))=\sigma(\varrho,f)+\sigma(\varrho,\pi(p))=\sigma(\varrho,p)\geq 0$ by Lemma 4.7. However, for rhombi ϱ that are not f-flat, it may be that $\sigma(\varrho,f)+\sigma(\varrho,\pi(p))<0$. Fortunately, it turns out that if p is an s-t-turnpath of minimal length, then this cannot happen!

The proof of the following result is astonishingly delicate and postponed to Section 7.

Theorem 4.8 (Shortest Path Theorem). *Let $f\in B_{\mathbb{Z}}$ and let p be a shortest s-t-turnpath in R_f. Then $f+\pi(p)\in B_{\mathbb{Z}}$.*

To investigate in a more general context to what extent the hive conditions are preserved when adding a flow $d\in\overline{F}(G)$ to $f\in B$, we make the following definition, extending Definition 3.7.

Definition 4.9. *For a hive flow $f\in B$, a flow $d\in\overline{F}(G)$ is called f-*hive preserving *if $f+\varepsilon d\in B$ for sufficiently small $\varepsilon>0$.*

We note that the set of f-hive preserving flows forms a cone C_f, which was called “cone of feasible directions” in [BI09].

Lemma 4.10. *Let $f\in B$ and $d'\in K(R_f)$. Then $\pi(d')$ is f-hive preserving.*

Proof. According to Lemma 2.1, there are complete turnpaths $p_1,\dots,p_m\in\mathcal{P}_f$ and $\alpha_1,\dots,\alpha_m\geq 0$ such that $d'=\sum_{i=1}^m\alpha_i p_i$. Lemma 4.7 tells us that $\sigma(\varrho,p_i)\geq 0$ if ϱ is f-flat.

By construction of R_f, if p_i crosses a border edge k on the right or bottom side of Δ, then $\delta(k,f) < b(k)$. This implies that $\delta(k, f + \varepsilon d') = \delta(k,f) + \varepsilon \sum_i \alpha_i \delta(k, p_i) < b(k)$ for sufficiently small $\varepsilon > 0$. The argument is analogous for the left border edges.

We show now that $f + \varepsilon\pi(d')$ is a hive flow for sufficiently small $\varepsilon > 0$. By the linearity of the slack, this means to show that for all rhombi ϱ, we have $\sigma(\varrho, f) + \varepsilon \sum_i \alpha_i \sigma(\varrho, p_i) \geq 0$ for sufficiently small $\varepsilon > 0$. In the case $\sigma(\varrho, f) > 0$, this is obvious. On the other hand, if $\sigma(\varrho, f) = 0$, this follows from $\sigma(\varrho, p_i) \geq 0$. □

4.2. Flatspaces. Our goal here is to get a detailed understanding of how turnpaths in R_f behave. For this,we first have to recall the concept of flatspaces from [KT99, Buc00]. In the following we fix $f \in B$.

By a *convex set L in the triangular graph Δ* we shall understand a union of hive triangles, which is convex. It is obvious that the angles at the corners of a convex set L are either acute of 60° or obtuse of 120°.

We call two hive triangles *adjacent* if they share a side and form an f-flat rhombus. This defines a graph whose vertices are the hive triangles. An *f-flatspace* is defined to be a connected component of this graph, see Figure 2.2. We simply write *flatspace* if it is clear, which flow f is meant. Also, we will identify flatspaces with the union of their hive triangles.

The following was observed in [Buc00].

Lemma 4.11.

1. *f-flatspaces are convex sets.*
2. *A side of an f-flatspace is either on the border of Δ, or it is also a side of a neighbouring flatspace.*
3. *There are exactly five types of convex sets: triangles, parallelograms, trapezoids, pentagons and hexagons.*

Proof. The first and second claim follow from Corollary 3.9. The third claim is just the enumeration of convex shapes on the triangular grid. □

We show next that turnpaths in R_f can move in f-flatspaces only in a very limited way: namely along the border in counterclockwise direction, and they can enter and leave the flatspace only through a few distinguished edges that we define next (cf. Figure 4.1).

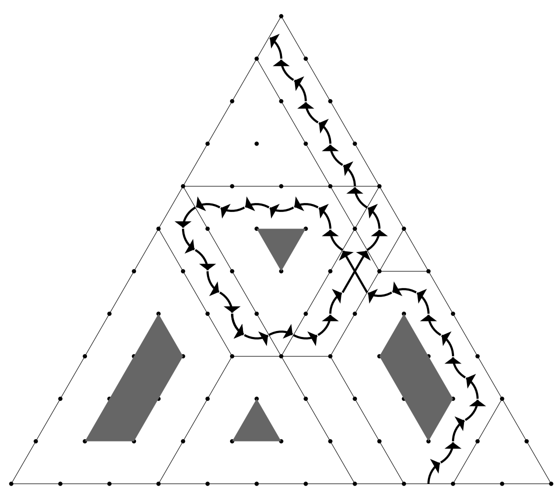

Fig. 4.1: The same situation as in Figure 2.2. The inner triangles are shaded while the border triangles are white. An exemplary turnpath in R_f is also depicted, where the two straight arrows represent crossing turnedges.

Let L be a convex set and a be one of its sides. For $k \in E(\Delta)$ we write $k \subseteq a$ to express that k is contained in a. Let $k_1, \ldots, k_r$ denote the edges contained in a in clockwise order. We call $a_{\to L} := k_1$ the *L-entrance edge* of a and $a_{L\to} := k_r$ the *L-exit edge* of a. We may have $r = 1$ in which case the entrance and exit edges coincide. Note that if M is a convex set adjacent to L, sharing with it the joint side a, then the M-entrance edge of a is at the same time the L-exit edge of L, that is, $a_{\to M} = a_{L\to}$.

The hive triangles in a convex set L either touch the border of L or lie inside L. Correspondingly, we will speak about *border triangles* and *inner triangles* of L. Recall from Definition 4.6 the set $\mathcal{P}_f$ of complete turnpaths in R_f.

PROPOSITION 4.12. *Let $p \in \mathcal{P}_f$ and L be an f-flatspace. Then:*

1. *p can enter L only by crossing entrance edges of L. Similarly, p can leave L only by crossing exit edges of L.*

2. *p uses only turnvertices in border triangles of L and traverses the border of L in counterclockwise direction.*

Proof. We call turnvertices, which lie in L and start at entrance edges of L, *entrance turnvertices*. Diagonals of non-f-flat rhombi shall be called *dividers*.

(1) If p enters L with a counterclockwise turn , then must be a divider. Hence is an entrance turnvertex.

(2) If p enters L with a clockwise and a counterclockwise turn , then is a divider and hence is an entrance turnvertex.

(3) If p enters L with two clockwise turns , then is a divider and hence is an entrance turnvertex.

Analogous arguments hold for exits with the situations , and . This proves the first assertion.

We now show the second assertion. Consider an inner triangle △. All rhombi in the shaded area are f-flat. By the definition of R_f, the counterclockwise turnvertices , and are not vertices of R_f. For the same reason, the clockwise turnvertices , and have no incident turnedge in R_f. This shows that turnpaths and turncycles in R_f can only use turnvertices in border triangles.

Finally, the fact that a counterclockwise turn in p implies that is a divider, shows that p traverses the border triangles of f-flatspaces in counterclockwise direction. ∎

Let $d \in \overline{F}(G)$ be a flow and $k \in E(G)$ be an edge lying at the border of a convex set L. If the hive triangle in L having the side k is upright, we define $\delta(k, \to L, d) := \delta(k, f)$, otherwise we set $\delta(k, \to L, d) := -\delta(k, f)$. We call $\delta(k, \to L, d)$ the *throughput of d into L through k*. It will be convenient to call $\delta(k, L \to, d) := -\delta(k, \to L, d)$ the *throughput of d out of L through k*.

Note that if the convex sets L and M are adjacent, sharing an edge k, then $\delta(k, \to M, d) = \delta(k, L \to, d)$.

For some of the following properties of throughputs compare Figure 2.2.

LEMMA 4.13. *Let L be a convex set contained in an f-flatspace and a be a side of L. Further, let $k_1, \dots, k_r \in E(\Delta)$ be the edges contained in a in clockwise order. Then $\delta(k_1, \to L, f) = \dots = \delta(k_r, \to L, f)$. Moreover, if $d \in \overline{F}(G)$ is f-hive preserving, then $\delta(k_1, \to L, d) \geq \dots \geq \delta(k_r, \to L, d)$.*

Proof. It is sufficient to show this for adjacent edges $k_1 =$ and $k_2 =$, where the rhombi ◊ and are f-flat. Since $0 = \sigma(\lozenge, f) = (f) + (f)$ and $0 = \sigma(, f) = (f) + (f)$, it follows $(f) = (f)$.

We have $\sigma(\lozenge, d) \geq 0$ and $\sigma(, d) \geq 0$ as d is f-hive preserving and ◊ and are f-flat. The second statement follows now similarly as before. ∎

OBSERVATION 4.14. *Let L be an f-flatspace with a side a lying on the left border of Δ. Then the maximum of the capacities $b(k)$ (cf. Definition 2.6) over all edges $k \subseteq a$ is attained at the exit edge of a. An analogous statement holds for the right and bottom border and entrance edges.*

Proof. This follows directly from the fact that $\nu_1 \geq \dots \geq \nu_n$ and the definition of the throughput capacities $b(k)$ of the border edges k of Δ, cf. Figure 2.1(c). Similarly for λ and μ. ∎

It will be important to decompose the throughput $\delta(k, \to L, d)$ into its positive and negative part. Recall that $\delta(k, \to L, d) = -\delta(k, L \to, d)$.

Definition 4.15. *Let $d \in \overline{F}(G)$, L be an f-flatspace, and $k \in E(\Delta)$ be an edge at the border of L. The* L-inflow of d through k *and* L-outflow of d through k *are defined as*

$$\omega(k, \to L, d) := \max\{\delta(k, \to L, d), 0\}, \quad \omega(k, L\to, d) := \max\{\delta(k, L\to, d), 0\}.$$

Further, for a side a of L, we define the L-inflow of d through a *and the* L-outflow of d through a *by*

$$\omega(a, \to L, d) \ := \ \sum_{k \subseteq a} \omega(k, \to L, d), \quad \omega(a, L\to, d) \ := \ \sum_{k \subseteq a} \omega(k, L\to, d).$$

We write $\omega(a, \to\Delta, d) := \omega(a, \to L, d)$ and $\omega(a, \Delta\to, d) := \omega(a, L\to, d)$ if the side a is on the border of Δ.

Note that $\delta(k, \to L, d) = \omega(k, \to L, d) - \omega(k, L\to, d)$. Further, if L and M are adjacent convex sets sharing a side a, then $\omega(k, \to L, d) = \omega(k, M\to, d)$ for $k \subseteq a$ and hence

$$\omega(a, \to L, d) = \omega(a, M\to, d). \tag{4.3}$$

The partition of Δ into f-flatspaces leads to a partition of the border of Δ. Let $\mathcal{S}_f$ denote the set of sides of f-flatspaces that lie on the right or bottom border of Δ.

Lemma 4.16. *For $f \in B$ and $d \in \overline{F}(G)$ we have*

$$\delta(d) = \sum_{a \in \mathcal{S}_f} \big(\omega(a, \to\Delta, d) - \omega(a, \Delta\to, d)\big).$$

Proof. By the definition of the overall throughput, and since s is connected in G only to the vertices on the right or bottom border of Δ, we have

$$\delta(d) \ = \ \sum_{e_- = s} d(e) - \sum_{e_+ = s} d(e) \ = \ \sum_k \delta(k, \to\Delta, d),$$

where the right-hand sum is over all edges $k \in E(\Delta)$ on the right or bottom border of Δ. Recall that $\delta(k, \to\Delta, d) = \omega(k, \to\Delta, d) - \omega(k, \Delta\to, d)$ By Definition 4.15,

$$\sum_k \omega(k, \to\Delta, d) \ = \ \sum_{a \in \mathcal{S}_f} \sum_{k \subseteq a} \omega(k, \to\Delta, d) \ = \ \sum_{a \in \mathcal{S}_f} \omega(a, \to\Delta, d).$$

Similarly, $\sum_k \omega(k, \Delta\to, d) = \sum_{a \in \mathcal{S}_f} \omega(a, \Delta\to, d)$ and the assertion follows. ∎

4.3. The Rerouting Theorem. We fix $f \in B$. Recall the set $\mathcal{P}_f$ of complete turnpaths in R_f from Definition 4.6. Let $\mathcal{P}_{st}$, $\mathcal{P}_{ts}$, and $\mathcal{P}_c$ denote the sets of s-t-turnpaths, t-s-turnpaths, and turncycles in R_f, respectively. Then we have the disjoint decomposition $\mathcal{P}_f = \mathcal{P}_{st} \cup \mathcal{P}_{ts} \cup \mathcal{P}_c$. Note that every $p \in \mathcal{P}_{st}$ enters Δ through exactly one edge on the right or bottom side of Δ, and leaves Δ through exactly one edge on the left side of Δ (otherwise s or t would be used more than once). Similarly, every $p \in \mathcal{P}_{ts}$ enters Δ through exactly one edge on the left side of Δ and leaves Δ through the right or bottom side of Δ. The reader should also note that turncycles $p \in \mathcal{P}_c$ may pass through s or t (or both of them).

Definition 4.17. A weighted family φ of complete turnpaths *in R_f is defined as a map $\varphi\colon \mathcal{P}_f \to \mathbb{R}_{\geq 0}$. If φ takes values in $\mathbb{N}$, we call φ a* multiset of complete turnpaths *in R_f. In this case, we interpret $\varphi(p)$ as the multiplicity with which p occurs in the multiset φ.*

A weighted family φ of complete turnpaths in R_f defines the nonnegative flow $\sum_{p \in \mathcal{P}_f} \varphi(p)p$ in R_f. On the other hand, by Lemma 2.1, any nonnegative flow $d' \in F_+(R_f)$ can be written in this form.

The flow $d' := \sum_p \varphi(p)p$ on R_f defined by the weighted family φ satisfies

$$\delta(d') \;=\; \sum_{p\in\mathcal{P}_{st}} \varphi(p) - \sum_{p\in\mathcal{P}_{ts}} \varphi(p). \tag{4.4}$$

To motivate the next definition, recall from Proposition 4.12 that a complete turnpath $p \in \mathcal{P}_f$ can enter an f-flatspace L only through an entrance edge $a_{\to L}$ of a side a of L, and leave only through an exit edge $a_{L\to}$.

Definition 4.18. *Let a be a side of an f-flatspace L. We denote by $\mathcal{P}_f(\to L, a)$ the set of $p \in \mathcal{P}_f$ that enter L through the edge $a_{\to L}$. The set $\mathcal{P}_f(L\to, a)$ denotes the set of $p \in \mathcal{P}_f$ that exit L through the edge $a_{L\to}$. For a weighted family φ of complete turnpaths and an f-flatspace L we define the* entrance weight *and the* exit weight *of a side a of L as follows:*

$$\omega(a, \to L, \varphi) \;:=\; \sum_{p\in\mathcal{P}_f(\to L,a)} \varphi(p), \quad \omega(a, L\to, \varphi) \;:=\; \sum_{p\in\mathcal{P}_f(L\to,a)} \varphi(p).$$

If the side a is on the border of Δ we write $\omega(a, \to\Delta, \varphi) := \omega(a, \to L, \varphi)$, $\omega(a, \Delta\to, \varphi) := \omega(a, L\to, \varphi)$, and $\mathcal{P}_f(\to\Delta, a) := \mathcal{P}_f(\to L, a)$.

The following remarkable result tells us that for any f-hive preserving flow $d \in \overline{F}(G)$, there is a weighted family φ of complete turnpaths such that the inflows and outflows of d through the sides a of the f-flatspaces are given by the entrance weight and exit weight of φ through a, respectively.

Theorem 4.19 (Rerouting Theorem). *Let $f \in B$ and $d \in \overline{F}(G)$ be f-hive preserving. Then there exists a weighed family φ of complete turnpaths in R_f such that $\omega(a, \to L, d) = \omega(a, \to L, \varphi)$ and $\omega(a, L\to, d) = \omega(a, L\to, \varphi)$ for all f-flatspaces L and all sides a of L. If d is integral, then we may assume that φ is a multiset.*

Let us draw an immediate consequence.

Corollary 4.20. *Under the assumptions of Theorem 4.19, the nonnegative flow $d' := \sum_{p\in\mathcal{P}_f} \varphi(p)p$ on R_f satisfies $\delta(d') = \delta(d)$.*

Proof. Recall that $\mathcal{S}_f$ denotes the set of sides of f-flatspaces that lie on the right or bottom border of Δ. We have

$$\sum_{p\in\mathcal{P}_{st}} \varphi(p) + \sum_{\substack{p\in\mathcal{P}_c\\ s\in p}} \varphi(p) \;=\; \sum_{a\in\mathcal{S}_f} \sum_{p\in\mathcal{P}_f(\to\Delta,a)} \varphi(p) \;=\; \sum_{a\in\mathcal{S}_f} \omega(a, \to\Delta, \varphi) \;=\; \sum_{a\in\mathcal{S}_f} \omega(a, \to\Delta, d),$$

where we have used Theorem 4.19 for the last equality. Similarly,

$$\sum_{p\in\mathcal{P}_{ts}} \varphi(p) + \sum_{\substack{p\in\mathcal{P}_c\\ s\in p}} \varphi(p) \;=\; \sum_{a\in\mathcal{S}_f} \sum_{p\in\mathcal{P}_f(\Delta\to,a)} \varphi(p) \;=\; \sum_{a\in\mathcal{S}_f} \omega(a, \Delta\to, d).$$

Subtracting and using (4.4) we get

$$\delta(d') \;=\; \sum_{p\in\mathcal{P}_{st}} \varphi(p) - \sum_{p\in\mathcal{P}_{ts}} \varphi(p) \;=\; \sum_{a\in\mathcal{S}_f} \omega(a, \to\Delta, d) - \sum_{a\in\mathcal{S}_f} \omega(a, \Delta\to, d) \;=\; \delta(d),$$

where we have used Lemma 4.16 for the last equality. □

The proof of the Rerouting Theorem is postponed to Section 6. The rough idea of the proof is to define a notion of canonical turnpaths within a convex set L, that specializes to the complete turnpaths in R_f restricted to L, in case L is an f-flatspace. We shall successively cut L into convex subsets by straight lines and recursively build up the required canonical turnpaths by operations of concatenation and straightening.

Remark 4.21. *For given $f \in B$, let C_f denote the cone of f-hive preserving flows on G. Lemma 4.10 states that $\pi(K(R_f)) \subseteq C_f$. Using Proposition 4.12, it is easy to see that this inclusion may be strict for some $f \in B$. On the other hand, one can show that $\pi(K(R_f)) = C_f$ if the hive triangles and rhombi are the only f-flatspaces. Hive flows f satisfying the latter property were called* shattered *in [Ike08, BI09]. We note that if f is shattered, then the Rerouting Theorem is not needed, and hence the optimality criterion stated in Proposition 4.22 below is much easier to prove.*

4.4. A first max-flow algorithm. We have already said that our goal is to maximize the overall throughput function δ on the polytope B of bounded hive flows. In order to implement this idea, we need a criterion that tells us when $f \in B$ is optimal.

PROPOSITION 4.22 (Optimality Criterion). *Let $f \in B$. Then $\delta(f) = \max_{g \in B} \delta(g)$ iff there exists no s-t-turnpath in R_f.*

Proof. We call $f \in B$ *optimal* iff $\delta(f) = \max_{g\in B} \delta(g)$.

If p is an s-t-turnpath in R_f, then by Lemma 4.10, we have $f + \varepsilon\pi(p) \in B$ for some $\varepsilon > 0$. Since $\delta(f + \varepsilon p) = \delta(f) + \varepsilon > \delta(f)$, the flow f is not optimal.

Now suppose that f is not optimal and let $g \in B$ such that $\delta(g) > \delta(f)$. Clearly, $d := g - f$ is f-hive preserving and satisfies $\delta(d) > 0$. Let φ be the weighted family of complete turnpaths corresponding to d as provided by the Rerouting Theorem 4.19, and put $d' := \sum_p \varphi(p)p$. Corollary 4.20 shows that $\delta(d') = \delta(d) > 0$ and (4.4) implies that there exists an s-t-turnpath in R_f. □

Consider the following Algorithm 1 for deciding positivity of LR coefficients.

Algorithm 1

Input: partitions λ, μ, ν with $|\nu| = |\lambda| + |\mu|$.
Output: **TRUE**, if $c^{\nu}_{\lambda,\mu} > 0$. **FALSE** otherwise.
1: $f \leftarrow 0$.
2: **while** there is a shortest s-t-turnpath p in R_f **do**
3: $\quad f \leftarrow f + \pi(p)$.
4: **end while**
5: **return** whether $\delta(f) = |\nu|$.

THEOREM 4.23. *Algorithm 1 returns whether $c^{\nu}_{\lambda,\mu} > 0$.*

Proof. Clearly f stays integral during the run of Algorithm 1. The Shortest Path Theorem 4.8 ensures that during the run of Algorithm 1 we always have $f \in B_{\mathbb{Z}}$. If Algorithm 1 returns **TRUE**, then we know that the final value of f is an integral and capacity achieving hive flow in B. Hence Proposition 2.7 implies $c^{\nu}_{\lambda,\mu} > 0$.

On the other hand, if Algorithm 1 returns **FALSE**, we have $\delta(f) < |\nu|$ and according to Proposition 4.22, the flow f has the maximum value of δ among all flows in B. Hence there is no capacity achieving flow in B and Proposition 2.7 implies that $c^{\nu}_{\lambda,\mu} = 0$. □

We note the following important integrality property.

COROLLARY 4.24. *For all λ, μ, ν, the overall throughput function δ attains the maximal value on $B(\lambda, \mu, \nu)$ at an integer flow.*

Proof. In the last line executed by Algorithm 1, there exists no s-t-turnpath in R_f. Hence, by Proposition 4.22, the integral flow f has the maximal value on B. □

As an application of the foregoing, we deduce here the saturation property of the Littlewood–Richardson coefficients, which was first shown in [KT99].

COROLLARY 4.25. *$c^{N\nu}_{N\lambda,N\mu} > 0$ for some $N \geq 1$ implies $c^{\nu}_{\lambda,\mu} > 0$.*

Proof. If $c^{N\nu}_{N\lambda,N\mu} > 0$, then there exists an integral capacity achieving hive flow $f \in B(N\lambda, N\mu, N\nu)$, by Proposition 2.7 . Hence $\frac{f}{N} \in B(\lambda, \mu, \nu)$ satisfies $\delta(\frac{f}{N}) = |\nu|$ and maximizes δ on $B(\lambda, \mu, \nu)$. Even though $\frac{f}{N}$ may not be integral, Corollary 4.24 implies that there exists an integral optimal hive flow $\tilde{f} \in B(\lambda, \mu, \nu)$ such that $\delta(\tilde{f}) = |\nu|$. Hence $c^{\nu}_{\lambda,\mu} > 0$ by Proposition 2.7. □

5. A polynomial time algorithm. In this section we use the capacity scaling approach (see, e.g., [AMO93, ch. 7.3]) to turn Algorithm 1 into a polynomial-time algorithm. During this method, $f \in B$ stays 2^{ℓ}*-integral*, for $\ell \in \mathbb{N}$, which means that all flow values are an integral multiple of 2^{ℓ}. The incrementation step in Algorithm 1, line 3, is replaced by adding $2^{\ell}\pi(p)$. Further, ℓ is decreased in the course of the algorithm. So our algorithm at first does big increments which over time decrease.

To implement this idea, we will search for a shortest s-t-turnpath in the subgraph R_f^ℓ of R_f defined next. By construction we will have $R_f^0 = R_f$. Recall that the polytope $B = B(\lambda, \mu, \nu)$ has the border capacity constraints as in Definition 2.6.

Definition 5.1. *Let $\ell \in \mathbb{N}$ and let $f \in B$ be 2^ℓ-integral. The digraph R_f^ℓ is obtained from R_f by deleting all turnedges crossing an edge k on the right or bottom border of Δ satisfying $\delta(k, f) + 2^\ell > b(k)$, and by deleting all turnedges crossing an edge k' on the left border of Δ satisfying $-\delta(k', f) + 2^\ell > b(k')$.*

Algorithm 2 stated below is now fully specified.

Algorithm 2

Input: partitions λ, μ, ν with $|\nu| = |\lambda| + |\mu|$ and $\nu_1 \geq \max \lambda_1, \mu_1$.
Output: **TRUE**, if $c^\nu_{\lambda,\mu} > 0$. **FALSE** otherwise.

```
f ← 0.
for ℓ from ⌈log ν_1⌉ down to 0 do
    while there is a shortest s-t-turnpath p in R_f^ℓ do
        f ← f + 2^ℓ π(p).
    end while
end for
return whether δ(f) = |ν|.
```

It is clear that f stays 2^ℓ-integral during the run of Algorithm 2.

Claim 5.2. *During the run of Algorithm 2, the flow f always is in B.*

Proof. Given a 2^ℓ-integral hive flow $f \in B = B(\lambda, \mu, \nu)$. First we note that the set of s-t-turnpaths on R_f^ℓ equals the set of s-t-turnpaths on $R_{\tilde f}(\tilde\lambda, \tilde\mu, \tilde\nu)$, where $\tilde f := f/2^\ell$, $\tilde\lambda := \lfloor \frac{\lambda}{2^\ell} \rfloor$, $\tilde\mu := \lfloor \frac{\mu}{2^\ell} \rfloor$, $\tilde\nu := \lfloor \frac{\nu}{2^\ell} \rfloor$, and division and rounding of partitions is defined componentwise. Let p be a shortest s-t-turnpath on R_f^ℓ and hence also a shortest s-t-turnpath on $R_{\tilde f}(\tilde\lambda, \tilde\mu, \tilde\nu)$. According to the Shortest Path Theorem 4.8 we have $\frac{f}{2^\ell} + \pi(p) \in B(\tilde\lambda, \tilde\mu, \tilde\nu)$. Therefore, we obtain $f + 2^\ell \pi(p) \in B(2^\ell\tilde\lambda, 2^\ell\tilde\mu, 2^\ell\tilde\nu) \subseteq B(\lambda, \mu, \nu) = B$. ∎

The last iteration of the for-loop of Algorithm 2 (where $\ell = 0$) operates like Algorithm 1 and hence Theorem 4.23 implies that Algorithm 2 works according to its specification.

For the analysis of the running time we need the following auxiliary result, relying on the Rerouting Theorem 4.19.

Lemma 5.3. *Let $f \in B_\mathbb{Z}$ be 2^ℓ-integral and $\ell \in \mathbb{N}$ be such that R_f^ℓ has no s-t-turnpath. Then $\delta_{max} - \delta(f) < 3n2^\ell$, where $\delta_{max} := \max_{g \in B} \delta(g)$.*

Proof. Let $g \in B$ with $\delta(g) = \delta_{\max}$ and put $d := g - f \in \overline{F}(G)$. Hence $\delta_{\max} - \delta(f) = \delta(d)$. Let φ be the family of complete weighted turnpaths corresponding to d as provided by the Rerouting Theorem 4.19. We decompose the set $\mathcal{P}_f = \mathcal{P}_{st} \cup \mathcal{P}_{ts} \cup \mathcal{P}_c$ of complete turnpaths in R_f into the sets $\mathcal{P}_{st}$, $\mathcal{P}_{ts}$, and $\mathcal{P}_c$ of s-t-turnpaths, t-s-turnpaths, and turncycles, respectively. Then the flow $d' := \sum_p \varphi(p)p$ on R_f defined by φ satisfies by (4.4) and Corollary 4.20,

$$\delta(d) \;=\; \delta(d') \;=\; \sum_{p \in \mathcal{P}_{st}} \varphi(p) - \sum_{p \in \mathcal{P}_{ts}} \varphi(p) \;\leq\; \sum_{p \in \mathcal{P}_{st}} \varphi(p). \tag{i}$$

A turnpath $p \in \mathcal{P}_{st}$ enters Δ exactly once (through the right or bottom border) and leaves Δ exactly once (through the left border). For an edge k on the right or bottom border of Δ, let $\mathcal{P}_f(k)$ denote the set of $p \in \mathcal{P}_f$ that enter Δ through k. Further, for an edge k' on the left border of Δ, let $\mathcal{P}_f(k')$ denote the set of $p \in \mathcal{P}_{st}$ that leave Δ through k'.

We call an edge k on the right or bottom border of Δ *small*, if $\delta(k, f) + 2^\ell > b(k)$. Let $\mathscr{E}$ denote the set of these edges. Note that for $k \in \mathscr{E}$ we have

$$\delta(k, d) \;=\; \delta(k, g) - \delta(k, f) \;\leq\; b(k) - \delta(k, f) \;<\; 2^\ell. \tag{ii}$$

Similarly, we call an edge $k' \in E(\Delta)$ on the left border of Δ *small*, if $-\delta(k', f) + 2^\ell > b(k')$ and denote the set of these edges by $\mathscr{E}'$. Border edges that are not small are called *big*.

The point is that an s-t-turnpath $p \in \mathcal{P}_{st}$ in R_f that crosses two big edges is also an s-t-turnpath in R_f^ℓ. Hence, by our assumption, there are no s-t-turnpaths in R_f that cross two big edges. We conclude that for all $p \in \mathcal{P}_{st}$, there exists $k \in \mathscr{E} \cup \mathscr{E}'$ such that $p \in \mathcal{P}_f(k)$. Therefore,

$$\sum_{p\in\mathcal{P}_{st}} \varphi(p) \;\le\; \sum_{k\in\mathscr{E}} \sum_{p\in\mathcal{P}_f(k)} \varphi(p) + \sum_{k'\in\mathscr{E}'} \sum_{p\in\mathcal{P}_f(k')} \varphi(p). \tag{iii}$$

To bound the right-hand sums, suppose first that $k \in \mathscr{E}$. By Proposition 4.12, $p \in \mathcal{P}_{st}(k)$ implies that k is the entrance edge of the side a of an f-flatspace L, in which case $k = a_{\to L}$. We have $\mathcal{P}_f(k) = \mathcal{P}_f(\to L, a)$ and hence, by Definition 4.18,

$$\sum_{p\in\mathcal{P}_f(k)} \varphi(p) \;=\; \omega(a, \to L, \varphi) = \omega(a, \to L, d),$$

where the last equality is guaranteed by the Rerouting Theorem 4.19.

Lemma 4.13 and Observation 4.14 imply that, since k is a small edge, all the other edges contained in a are small as well. Note that $\delta(\tilde{k}, \to L, d) < 2^\ell$ implies $\omega(\tilde{k}, \to L, d) < 2^\ell$ for all edges $\tilde{k} \subseteq a$ by Definition 4.15. Therefore we can use (ii) to deduce

$$\omega(a, \to L, d) \;=\; \sum_{\tilde{k}\subseteq a} \omega(\tilde{k}, \to L, d) \;<\; |a|\, 2^\ell,$$

where $|a|$ denotes the number of edges of Δ contained in a. Summarizing, we conclude for $k \in \mathscr{E}$,

$$\sum_{p\in\mathcal{P}_f(k)} \varphi(p) \;<\; |a|\, 2^\ell,$$

where a is the side of the f-flatspace, in which k lies.

The same bound holds for $k' \in \mathscr{E}'$ by an analogous argument. Combining these bounds with (i) and (iii) we obtain

$$\delta(d) = \delta(d') \;\le\; \sum_{p\in\mathcal{P}_{st}} \varphi(p) \;<\; 3n2^\ell$$

since there are $3n$ edges on the border of Δ. ∎

THEOREM 5.4. *Algorithm* 2 *decides the positivity of the Littlewood–Richardson coefficient* $c^\nu_{\lambda,\mu}$ *with* $\mathcal{O}\big(\ell(\nu)^3 \log \nu_1\big)$ *arithmetic operations and comparisons.*

Proof. Again let $\delta_{\max} := \max_{g\in B} \delta(g)$. After ending the while-loop for the value ℓ, there is no s-t-turnpath in R_f^ℓ and hence $\delta_{\max} - \delta(f) < 3n2^\ell = 6n\, 2^{\ell-1}$. Hence in the next iteration of the while-loop, for the value $\ell - 1$, at most $6n$ s-t-turnpaths can be found. Moreover, note that the initial value of ℓ is so large, that in the first iteration of the while-loop at most one s-t-turnpath can be found.

So Algorithm 2 finds at most $6n\lceil \log \nu_1 \rceil$ many s-t-turnpaths and searches at most $\log \nu_1$ many times for an s-t-turnpath without finding one. Note that searching for a shortest s-t-turnpath requires at most time $\mathcal{O}(n^2)$ using breadth-first-search, since there are $\mathcal{O}(n^2)$ turnvertices and turnedges. Hence we get a total running time of $\mathcal{O}\big(n^3 \log \nu_1\big)$. ∎

6. Proof of the Rerouting Theorem. We will easily derive the Rerouting Theorem 4.19 by a glueing argument from the Canonical Turnpath Theorem 6.6 below. The latter is a general result for hive flows on convex sets in the triangular graph. In the next subsection we introduce the necessary terminology to state the Canonical Turnpath Theorem, which is then proved by induction in the remainder of this section, considering separately the different possible shapes of convex sets.

6.1. Canonical turnpaths in convex sets. Let L be a convex set in the triangular graph Δ. We define the graph G_L as the induced subgraph of the honeycomb graph G obtained by restricting to the set of vertices lying in L (including the vertices on the border of L, but omitting s and t). A *flow on* G_L is defined as a map $E(G_L) \to \mathbb{R}$ satisfying the flow conservation laws at all vertices of G_L that do not lie at the border of L. The vector space $\overline{F}(G_L)$ of *flow classes on* G_L is defined as in (2.1) by factoring out the null flows. As in Definition 2.5, we define a *hive flow* f *on* L as a flow class in $\overline{F}(G_L)$ satisfying $\sigma(\varrho, f) \geq 0$ for all rhombi ϱ lying in L. Similarly, we define the notion of a *flow on* R *(or* R_f*) restricted to* L, by restricting to the subgraph induced by the turnvertices lying in L.

For a fixed convex set L we are going to define a set $\mathcal{P}_L$ of distinguished turnpaths p in L starting at some entrance edge $a_{\to L}$ and ending at some exit edge $b_{L\to}$. The goal is to achieve that p is a turnpath in R_f whenever L is an f-flatspace. Proposition 4.12 provides the guiding principle for making the right definition.

Let $r \geq 3$ and $a_1, \ldots, a_r$ be a sequence of successive sides of L in counterclockwise order, where $a_2, \ldots, a_{r-1}$ are different. Further, we assume that the angles between a_{i-1} and a_i are obtuse for $i = 3, \ldots, r-1$. We then form a unique turnpath p moving within the border triangles of L from $(a_1)_{\to L}$ to $(a_r)_{L\to}$ in counterclockwise direction, cf. Figure 6.1. The

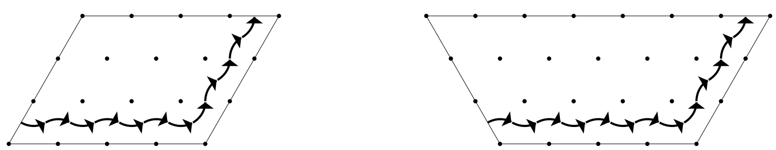

Fig. 6.1: Canonical turnpaths in a parallelogram and in a trapezoid. The left hand turnpath starts with a counterclockwise turn (acute angle) and the righthand turnpath starts with a clockwise turn (obtuse angle).

turnpath p alternatively takes clockwise and counterclockwise turns, except at the (obtuse) angles of L between a_{i-1} and a_i (for $i = 3, \ldots, i-1$), where p takes two consecutive counterclockwise turns to go around. If a_1, a_2 form an acute angle, then p starts with a counterclockwise turn, otherwise p starts with a clockwise turn. Moreover, if a_{r-1}, a_r form an acute angle, then p ends with a counterclockwise turn, otherwise p ends with a clockwise turn. We call the resulting turnpath a *canonical turnpath* of L. We shall also consider the turnpaths consisting of a single clockwise turnvertex at an acute angle as a canonical turnpath of L.

Definition 6.1. *The symbol* $\mathcal{P}_L$ *denotes the set of all canonical turnpaths of the convex set* L*. For* $p \in \mathcal{P}_L$*, we denote by* $\mathsf{start}(p) = (a_1)_{\to L}$ *the edge of* Δ *from which* p *starts and by* $\mathsf{end}(p) = (a_r)_{L\to}$ *the edge of* Δ *where* p *ends.*

Example 6.2. *A triangle has exactly six canonical turnpaths, cf. Figure* 6.2*. A parallelogram has exactly eight canonical turnpaths, cf. Figure* 6.3*. In particular, this holds true for rhombi. A trapezoid has exactly nine canonical turnpaths, cf. Figure* 6.4*. A pentagon has exactly 16 canonical turnpaths, cf. Figure* 6.5*. A hexagon has six canonical turnpaths up to rotations, which makes a total of 36 turnpaths, see Figure* 6.6*.*

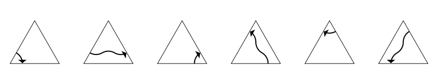

Fig. 6.2: The six canonical turnpaths in a triangle.

Lemma 6.3. *Let* f *be a hive flow and* L *be one of its* f*-flatspaces. Then any canonical turnpath* $p \in \mathcal{P}_L$ *of* L *is a turnpath in* R_f*.*

Proof. We need to ensure that p uses no negative contributions in f-flat rhombi. But p does not use any pair of successive clockwise turnvertices at all. And whenever p uses a counterclockwise turnvertex ◊, then ◊ is at the border of an f-flatspace. □

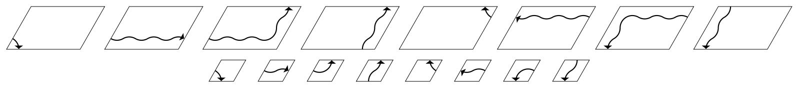

Fig. 6.3: Top row: The eight canonical turnpaths in a parallelogram. Bottom row: The same list in the special case of a rhombus.

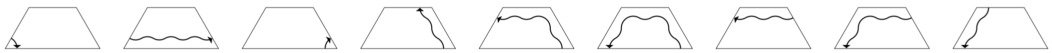

Fig. 6.4: The nine canonical turnpaths in a trapezoid.

We have to extend some of the notions introduced in Section 4.3.

DEFINITION 6.4. *A multiset φ of* canonical turnpaths *in a convex set L is defined as a map $\varphi\colon \mathcal{P}_L \to \mathbb{N}_{\geq 0}$. Let a be a side of L. The number of turnpaths of φ starting from $a_{\to L}$ and ending at $a_{L\to}$, respectively, is denoted by*

$$\omega(a, \to L, \varphi) := \sum_{\mathit{start}(p)=a_{\to L}} \varphi(p), \quad \omega(a, L\to, \varphi) := \sum_{\mathit{end}(p)=a_{L\to}} \varphi(p).$$

Note that the weighted sum $\sum_{p\in\mathcal{P}_L} \varphi(p)p$ defines a nonnegative flow d'_L on R restricted to L. Moreover, if L is an f-flatspace, then d'_L is a flow on R_f restricted to L, as a consequence of Lemma 6.3.

DEFINITION 6.5. *Let d be a hive flow on a convex set L. A multiset φ of canonical turnpaths on L is called* compatible with d, *if $\omega(a, \to L, \varphi) = \omega(a, \to L, d)$ and $\omega(a, L\to, \varphi) = \omega(a, L\to, d)$ for all edges a of L.*

The key result, amenable to an inductive proof along L, is the following.

THEOREM 6.6 (Canonical Turnpath Theorem). *Let L be a convex set and d be an integral hive flow on L. Then there exists a multiset φ of canonical turnpaths on L which is* compatible with d.

LEMMA 6.7. *The Canonical Turnpath Theorem* 6.6 *implies the Rerouting Theorem* 4.19.

Proof. We first note that it suffices to prove the Rerouting Theorem 4.19 for an integral flow $d \in \overline{F}(G)$. Indeed, then it trivially must hold for a rational d. A standard continuity argument then shows the assertion for a real d.

So let $f \in B$ and $d \in \overline{F}(G)$ be integral and f-hive preserving. Theorem 6.6 applied to every f-flatspace L and the hive flow d restricted to L yields a multiset φ_L of canonical turnpaths on L compatible with d restricted to L. By Lemma 6.3, canonical turnpaths of L are in R_f.

Suppose that L and M are adjacent f-flatspaces sharing the side a. Then we have, using Theorem 6.6 and (4.3),

$$\omega(a, \to L, \varphi_L) = \omega(a, \to L, d) = \omega(a, M\to, d) = \omega(a, M\to, \varphi_M). \tag{6.1}$$

We set up an arbitrary bijection between the turnpaths p_M in φ_M ending at $a_{M\to}$ and the turnpaths p_L in φ_L starting from $a_{\to L} = a_{M\to}$ and concatenate these turnpaths correspondingly. It is essential to note that the additional turnedges used for joining p_M and p_L lie in R_f, since the rhombus with diagonal $a_{\to L}$ is not f-flat!

Similarly, we have $\omega(a, \to M, \varphi_M) = \omega(a, L\to, \varphi_L)$ and we concatenate the turnpaths in φ_L ending at $a_{L\to}$ with the turnpaths in φ_M starting from $a_{\to M} = a_{L\to}$ correspondingly.

Doing so for all sides a shared by different f-flatspaces, we obtain a multiset of turncycles in R_f and a multiset of turnpaths in R_f going from a side of Δ to a side of Δ. These turnpaths can be extended to complete turnpaths. Altogether, we obtain a multiset φ of complete turnpaths in R_f.

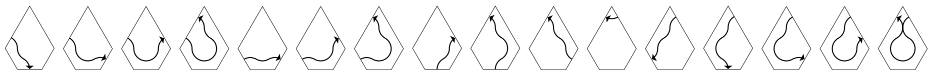

Fig. 6.5: The 16 canonical turnpaths in a pentagon.

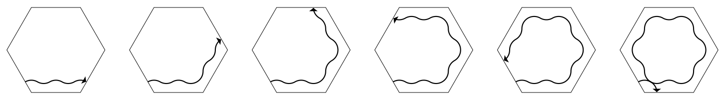

Fig. 6.6: The six canonical turnpaths in a hexagon starting from a fixed side.

Then we have $\omega(a, \rightarrow L, \varphi) = \omega(a, \rightarrow L, \varphi_L)$ and $\omega(a, L\rightarrow, \varphi) = \omega(a, L\rightarrow, \varphi_L)$ for any side a of an f-flatspace L. Hence, $\omega(a, \rightarrow L, \varphi) = \omega(a, \rightarrow L, d)$ and $\omega(a, L\rightarrow, \varphi) = \omega(a, L\rightarrow, d)$ by (6.1). So the multiset φ is as required. □

In the subsequent sections we shall prove Theorem 6.6 for the five possible shapes of L. Although the arguments are quite similar for the different shapes, there are subtle differences. We begin with the case of parallelograms.

6.2. Parallelograms. By the *size* of a convex set L we understand the number of hive triangles contained in L. We will prove the Canonical Turnpath Theorem 6.6 for parallelograms L by induction on the size of L. The induction start is provided by the following lemma.

LEMMA 6.8. *The assertion of the Canonical Turnpath Theorem* 6.6 *is true if* $L = \varrho$ *is a rhombus. More specifically, if* d *is an integral hive flow on* ϱ*, then there is a multiset* φ *of canonical turnpaths compatible with* d*, such that for all* $p \in \Psi_+(\varrho)$ *occuring in* φ *we have* $p \subseteq \mathrm{supp}(d)$.

Proof. The canonical turnpaths in a rhombus ϱ are exactly the eight contributions in $\Psi_+(\varrho) \cup \Psi_0(\varrho)$, see Figure 6.3.

Given an integral hive flow d on ϱ. If $p \subseteq \mathrm{supp}(d)$ for some $p \in \Psi_-(\varrho)$, then $p' \subseteq \mathrm{supp}(d)$ by Lemma 3.4 on antipodal contributions. Since $\sigma(\varrho, p + p') = 0$ it follows that $d - (p + p')$ is a hive flow. So we can successively subtract flows of the form $p + p'$ from d to arrive at a flow decomposition $d = \sum_i m_i(p_i + p'_i) + h$, where $m_i \in \mathbb{N}$, h is a hive flow on ϱ, and $p \not\subseteq \mathrm{supp}(h)$ for all $p \in \Psi_-(\varrho)$. Moreover, $\mathrm{supp}(h) \subseteq \mathrm{supp}(d)$ by construction.

It is straighforward to check that h must be a nonnegative integer linear combination of turnpaths $p \in \Psi_+(\varrho) \cup \Psi_0(\varrho)$ such that $p \subseteq \mathrm{supp}(h)$.

Now we replace the sums $p_i + p'_i$ by sums $n_i + n'_i$ of two neutral slack contributions as follows: we replace 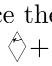 by 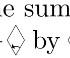, we replace 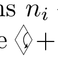 by 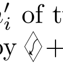, and similarly in the situations rotated by 180°. Since this exchange does not alter the number of turnpaths entering and leaving a side of ϱ, this leads to a multiset of canonical turnpaths of ϱ satisfing the desired requirements. □

The induction step will be based on the following result on *straightening* canonical turnpaths.

PROPOSITION 6.9. *Let* L *be a parallelogram cut into two parallelograms* L_1 *and* L_2 *by a straight line parallel to one of the sides of* L*. Further let* p *be a turnpath going from the side* a *of* L *to the side* b *of* L *such that* p *is either a canonical turnpath of* L_1*, or* p *is obtained by concatenating a canonical turnpath* p_1 *of* L_1 *with a canonical turnpath* p_2 *of* L_2*. Then* p *can be* straightened*, that is, there exists a canonical turnpath of* L *going from* $a_{\rightarrow L}$ *to* $b_{L\rightarrow}$.

Proof. It suffices to check the various cases. Recall the possible canonical turnpaths in a parallelogram from Figure 6.3. Figure 6.7(a) shows how to treat the four possible canonical turnpaths of L_1 going from a side of L to a side of L. Note that only in two of these four cases, the turnpath has to be changed (by "stretching" or moving parallely). Figure 6.7(b) shows how to treat the six possible cases of a concatenation p of a turnpath in L_1 with one in L_2. Only in two of the six cases, the turnpath has to be changed (by "shrinking"). □

(a) The modifications of the four canonical turnpaths of L_1 going from a side of L to a side of L.

(b) The modifications of the six possible cases of a concatenation p of a turnpath in L_1 with one in L_2.

Fig. 6.7: Illustration of the proof of Proposition 6.9.

PROPOSITION 6.10. *The assertion of the Canonical Turnpath Theorem* 6.6 *is true if* L *is a parallelogram.*

Proof. We proceed by induction on the size of L. The induction start is provided by Lemma 6.8. We suppose now that L has size greater than two and we cut L into two parallelograms L_1 and L_2 by a straight line parallel to one the sides of L. The induction hypothesis yields the existence of multisets φ_i of canonical turnpaths of L_i compatible with d restricted to L_i, for $i = 1, 2$.

Let a denote the side shared by L_1 and L_2 and note that $a_{L_1\to} = a_{\to L_2}$. The reader should note that it is not possible for a canonical turnpath in L_i to start and end at a, see Figure 6.3. Using Definition 6.4, the induction hypothesis, and (4.3), we get

$$\omega(a, L_1\!\to, \varphi_1) = \omega(a, L_1\!\to, d) = \omega(a, \to\! L_2, d) = \omega(a, \to\! L_2, \varphi_2).$$

This means that the number of turnpaths p_1 in φ_1 ending at $a_{L_1\to}$ equals the number of turnpaths p_2 in φ_2 starting at $a_{\to L_2}$. It is therefore possible to set up a bijection between the set of turnpaths p_1 in φ_1 ending at $a_{L_1\to}$ with the set of turnpaths p_2 in φ_2 starting at $a_{\to L_2}$, and to concatenate each p_1 with its partner p_2 to obtain a turnpath q in L starting from a side of L and ending at a side of L. However, the turnpath q may not be canonical for L. But now we use Proposition 6.9 to replace q by a canonical turnpath of L starting and ending at the same sides of L as q does.

Similarly, $a_{L_2\to} = a_{\to L_1}$ and we get $\omega(a, L_2\!\to, \varphi_2) = \omega(a, \to\! L_1, \varphi_1)$. As before, we can match and concatenate the turnpaths p_2 in φ_2 ending at $a_{L_2\to}$ with the turnpaths p_1 in φ_1 starting at $a_{\to L_1}$. Again, we use Proposition 6.9 to replace the resulting turnpaths by canonical turnpaths of L without changing the starting and ending side.

We also apply Proposition 6.9 to the turnpaths in φ_1 and φ_2 going from a side of L to a side of L.

After performing these procedures, we obtain a multiset φ of canonical turnpaths of L. Let b be the side of L_1 parallel to a. Then we have by construction

$$\omega(b, \to\! L, \varphi) = \omega(b, \to\! L_1, \varphi_1) = \omega(b, \to\! L_1, d) = \omega(b, \to\! L, d).$$

Similarly for b being the side of L_2 parallel to a. Now let b be a side of L cut by a into line segments b_1 and b_2. Then we have

$$\omega(b, \to\! L, \varphi) = \omega(b_1, \to\! L_1, \varphi) + \omega(b_2, \to\! L_2, \varphi) = \omega(b_1, \to\! L_1, d) + \omega(b_2, \to\! L_2, d) = \omega(b, \to\! L, d).$$

It follows that φ is compatible with d. □

6.3. Trapezoids, pentagons and hexagons. We first treat the case of trapezoids. Again, the strategy is to proceed by induction, cutting the trapezoid into smaller trapezoids or parallelograms. But now, unlike the case of parallelograms before, the cutting has to be done in a certain way in order to ensure the straightening of canonical turnpaths. The following result identifies the critical cases to be avoided. The straightforward proof is similar to the one of Proposition 6.9 and left to the reader, who should consult Figures 6.3–6.4 for the possible canonical turnpaths in a parallelogram or a trapezoid, respectively.

The *height* of a convex set L is defined as the number of its edges on its shortest side.

PROPOSITION 6.11. *Let L be a trapezoid cut into convex sets L_1 and L_2 by a straight line a. Further let p be a turnpath going from the side b of L to the side c of L such that p is either a canonical turnpath of L_1, a canonical turnpath of L_2, or p is obtained by concatenating canonical turnpaths of L_1 with canonical turnpaths L_2 (in any order).*

1. If a is parallel to the longest side of L so that L_1 and L_2 are trapezoids, then p can be straightened, i.e., there exists a canonical turnpath of L going from $b_{\to L}$ to $c_{L\to}$.

2. Suppose that L has the height 1 *and that L is cut by a into a trapezoid (or a triangle) and a rhombus (there are two possibilities to do so). Then p can be straightened unless in the four critical cases depicted in Figure* 6.8*(a)-(b).* □

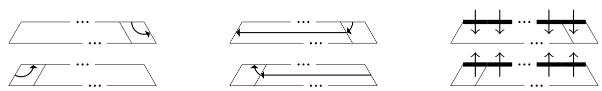

(a) (b) (c)

Fig. 6.8: A trapezoid of height 1 cut into a trapezoid and a parallelogram with the four critical cases of a turnpath p that cannot be straightened.

PROPOSITION 6.12. *The assertion of the Canonical Turnpath Theorem* 6.6 *is true if L is a trapezoid.*

Proof. We make induction on the size of L. The case where L is a hive triangle, which we consider a degenerate trapezoid, is trivial. So suppose that L has size at least three and let d be an integral hive flow on L.

(A) If L has height greater than 1, then we cut L into two trapezoids L_1 and L_2 by a straight line parallel to the longest side of L and apply the induction hypothesis to L_1 and L_2 to obtain multisets φ_i of canonical turnpaths of L_i compatible with d restricted to L_i, for $i = 1, 2$. Proceeding as in the proof of Proposition 6.10 and using Proposition 6.11(1), we construct from φ_1, φ_2 a multiset φ of canonical turnpaths of L satisfying $\omega(a, \to L, \varphi) = \omega(a, \to L, d)$ and $\omega(a, L\to, \varphi) = \omega(a, L\to, d)$ for all sides a of L.

(B) Now suppose that L has height 1. There are two possibilities **t** and **b** of cutting L by a straight line into a trapezoid and a parallelogram, as depicted in the top and bottom row of Figure 6.8.

Choose the **t** version of cutting and apply the induction hypothesis and Proposition 6.10 to L_1 and L_2, respectively, to obtain multisets φ_i of canonical turnpaths of L_i. Then we apply the straightening of Proposition 6.11(2) as before, which succeeds unless we are in one of the critical cases as depicted in the top row of Figure 6.8. For instance, assume that L_2 is a rhombus and the turnpath $p =$  occurs in φ_2 as in Figure 6.8(a). Let k_r denote the edge of L_2 where p starts and a be the corresponding side of L. Then, using Definition 4.15, we have

$$\omega(k_r, \to L_2, d) = \omega(k_r, \to L_2, \varphi_2) \geq 1,$$

hence $\delta(k_r, \to L_2, d) > 0$. Lemma 4.13 implies that $\delta(k, \to L, d) > 0$ for all edges k of a, see Figure 6.8(c). The same conclusion can be drawn when $p =$ 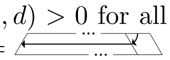 occurs in φ_2.

The clue is now that if we cut L in the other possible way (**b** version, cf. bottom row of Figure 6.8), then no critical case can occur. Indeed, otherwise, by an analogous reasoning as before, we had $\delta(k, L\to, d) > 0$ for all edges k of a, which contradicts $\delta(k, \to L, d) > 0$.

Similarly one shows that if we start with the **b** version of cutting L and a critical case occurs, then cutting with the **t** version succeeds. □

For later use, we note the following observation resulting from the above proof.

Observation 6.13. *Let L be a trapezoid of height 1 and d be an integral hive flow on L. If the multiset φ of canonical turnpaths compatible with d resulting from the proof of Proposition 6.12 contains the turnpath $q =$* 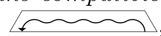*, then there is a rhombus ϱ and a turnedge $p \in \Psi_+(\varrho)$ as in Figure 6.9 such that $p \subseteq \mathrm{supp}(d)$.*

Fig. 6.9: On Observation 6.13.

Proof. Tracing part (B) of the inductive proof of Proposition 6.12 shows that q results from a smaller turnpath $\tilde{q}$ either by stretching to the right or left, or by appending to $\tilde{q}$ a turnedge p in the right or left rhombus ϱ, see Figure 6.10. In the case p is appended, we

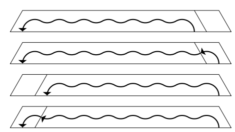

Fig. 6.10: Inductive construction of the turnpath q along the proof of Proposition 6.12. In the first and third case stretching is necessary.

know that $p \subseteq \mathrm{supp}(d)$ by Lemma 6.8. Otherwise, we conclude by the induction hypothesis. □

We settle now the case of pentagons.

Proposition 6.14. *The assertion of the Canonical Turnpath Theorem 6.6 is true if L is a pentagon.*

Proof. We proceed by induction on the size of L, cutting the pentagon by a straight line into a pentagon and a trapezoid, or a parallelogram and a trapezoid. The two critical cases, where a straightening fails, are depicted in Figure 6.11. As in the proof of Proposition 6.12,

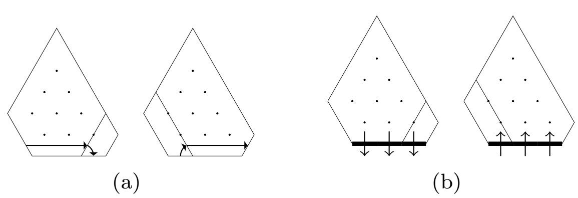

Fig. 6.11: Two ways of cutting a pentagon with the two critical cases of turnpaths that cannot be straightened and the resulting throughputs.

one can show that in one of the two possible ways of cutting L, no critical case can occur. The details are left to the reader. □

The case where L is a hexagon is the simplest one of all and does not require an inductive argument as before.

Proposition 6.15. *The assertion of the Canonical Turnpath Theorem 6.6 is true if L is a hexagon.*

Proof. Let $a_1, \dots, a_6$ denote the six sides of a hexagon L. We write $a_i^- := (a_i)_{\to L}$ and $a_i^+ := (a_i)_{L\to}$ for the entrance and exit edges of L, respectively.

The essential observation is that for any pair (i, j) there exists a canonical turnpath of L going from a_i^- to a_j^+. This is easily verified by looking at Figure 6.6.

Let d be an integral flow on L and put $\mathsf{in}(i) := \omega(a_i, \to L, d)$ and $\mathsf{out}(i) := \omega(a_i, L\to, d)$ for $1 \le i \le 6$. The flow conservation laws imply that $\sum_i \mathsf{in}(i) = \sum_i \mathsf{out}(i)$.

We form a list $\mathcal{L}^-$ of entrance edges in which a_i^- occurs $\mathsf{in}(i)$ many times and we form a list $\mathcal{L}^+$of exit edges in which a_i^+ occurs $\mathsf{out}(i)$ many times. Both lists have the same length. We now connect, for all j, the jth element of $\mathcal{L}^-$ with the jth element of $\mathcal{L}^+$ by a canonical turnpath p_j of L. This is possible by the observation made at the beginning of the proof. The resulting multiset φ of canonical turnpaths of L satisfies $\omega(a_i, \rightarrow L, \varphi) = \omega(a_i, \rightarrow L, d)$ and $\omega(a_i, L\rightarrow, \varphi) = \omega(a_i, L\rightarrow, d)$ by construction. □

6.4. Triangles. We need the following flow propagation lemma.

LEMMA 6.16. *Let L be a trapezoid and d be a hive flow on L.*

1. Let p be the path in Figure 6.12(a) and suppose that $p \subseteq \mathrm{supp}(d)$. Then all the edges of G belonging to the turns in Figure 6.12(b) belong to $\mathrm{supp}(d)$ as well. Moreover, $\mathrm{supp}(d)$ cannot contain the paths $q_1, q_2 \in \Psi_+(\varrho)$ in the shaded rhombi ϱ depicted in Figure 6.12(c).

2. If the path $\tilde{p}$ in Figure 6.12(a') satisfies $\tilde{p} \subseteq \mathrm{supp}(d)$, then a similar conclusion can be drawn, see Figures 6.12(b')-(c').

Proof. 1. The first assertion on $\mathrm{supp}(d)$ follows by successively applying Lemma 3.5 on flow propagation.

The assertion $q_i \not\subseteq \mathrm{supp}(d)$ follows by inspecting the edges of G involved in the paths appearing in the bottom row of the trapezoid in Figure 6.12(c) and noting that $\mathrm{supp}(d)$ cannot contain an edge $k \in E(G)$ and its reverse.

2. The second case is treated similarly. □

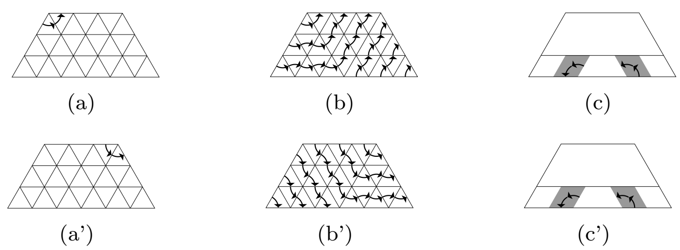

Fig. 6.12: If d is a hive flow and the path on the left figure is contained in $\mathrm{supp}(d)$, then the turns in the middle figure are contained in $\mathrm{supp}(d)$. The two paths on the right figure cannot be contained in $\mathrm{supp}(d)$.

The following proposition completes the proof of the Canonical Turnpath Theorem 6.6 for any shapes of L.

PROPOSITION 6.17. *The assertion of the Canonical Turnpath Theorem 6.6 is true if L is a triangle.*

Proof. Again we proceed by induction on the size of L, the start of a hive triangle being trivial. For the induction step, suppose that d is an integral hive flow on L, and note that there are three ways of cutting L into a trapezoid L_1 of height 1 and a triangle L_2. We choose one as in Figure 6.13(a). The induction hypothesis and Proposition 6.12 yield multisets φ_i compatible with d restricted to L_1 and L_2, respectively. Using Figure 6.2 and Figure 6.4 showing the possible canonical turnpaths in triangles and trapezoids, the reader should verify that the procedure of concatenation and straightening, as explained in the proof of Proposition 6.10, can only fail in the critical case where φ_1 contains a turnpath q as depicted in Figure 6.13(a).

By Observation 6.13 applied to the trapezoid L_1 we may assume that there is rhombus ϱ (shaded in Figure 6.13(b)) and a path $p \in \Psi_+(\varrho)$ such that $p \subseteq \mathrm{supp}(d)$. Now we can apply Lemma 6.16 as depicted in Figure 6.13(b) and conclude that all the turns depicted in this figure are contained in $\mathrm{supp}(d)$. Suppose we are in the left-hand situation of Figure 6.13(b). Then we can decompose L into a trapezoid of height 1 and a triangle by cutting along the right-hand side of L. Lemma 6.16 implies that no critical case can arise, so that in this situation, the procedure of concatenation and straightening works. If we are in the right-

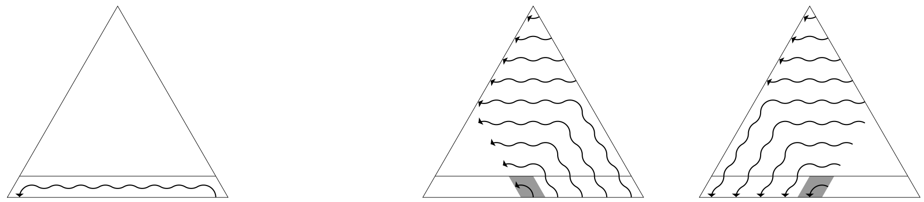

(a) Cutting a triangle into a trapezoid and a triangle with the only critical case of a turnpath q that cannot be straightened.

(b) The turnpath in the shaded rhombus is contained in supp(d), implying flow propagation according to Lemma 6.16.

Fig. 6.13: On the proof of Proposition 6.17.

hand situation of Figure 6.13(b), we can decompose L into a trapezoid of height 1 and a triangle by cutting along the left-hand side of L and argue analogously. □

7. Proof of the Shortest Path Theorem. In this section we prove Theorem 4.8.

7.1. Special rhombi. For the whole subsection we fix a *shortest* s-t-turnpath p in R_f for $f \in B_{\mathbb{Z}}$. Flatness shall always refer to f.

We will show in several steps that the minimal length of p poses severe restrictions on the way p may pass a rhombus. Before doing so, let us verify an even simpler property of p resulting from the minimality. Each turnvertex in R has a *reverse turnvertex* in R that points in the other direction, e.g., the reverse turnvertex of ◊ is ◊. We note that if p contains a turnvertex v touching the boundary of Δ, than p cannot contain the reverse of v (otherwise, p would use s or t more than once).

Here and in the following, statements involving the pictorial description ◊ include the possibility of a rotation by 180°.

PROPOSITION 7.1. *The turnpath p cannot use a turnvertex and its reverse.*

Proof. By way of contradiction, let w be the *first* turnvertex in p whose reverse turnvertex is also used by p. We already noted that w cannot touch the border of Δ. Let v denote the predecessor of w and let ◊ stand for the rhombus that contains both v and w.

Suppose first that v is a clockwise turn: $v =$ ◊. Since the reverse of w is in p, the turnpath p must use ◊ or ◊. But ◊ is excluded because of the minimal choice of w. If ◊ were not flat, then it is easy to check that p could be rerouted via ◊. This contradicts the minimal length of p. So let us assume that ◊ is flat. Then ◊ $\notin E(R_f)$ by construction of R_f and thus ◊ $\in p$. Hence $w =$ ◊, which implies ◊ $\in p$. But since ◊ $\notin E(R_f)$, it follows that ◊ is used by p, in contradiction with the minimal choice of w.

It remains to analyze the case where v is a counterclockwise turn: $v =$ ◊. As before we must have ◊ $\in p$. The existence of counterclockwise turns at acute angles implies that ◊ and ◊ are not flat. Hence p can be rerouted via ◊, in contradiction with the minimal length of p. □

We continue now by analyzing the possible ways p may pass through a rhombus. Note that, due to Proposition 7.1, the turnpath p can cross the diagonal ◊ of a rhombus at most twice.

LEMMA 7.2. *If ◊ is not flat, then its diagonal ◊ is crossed by p at most once.*

Proof. Assume by way of contradiction that both ◊ and ◊ occur in p. Since p cannot use a turnvertex twice, there are only two possibilites:

either ◊ and ◊ are edges of p or ◊ and ◊ are edges of p. (7.1)

In both cases, p can be rerouted resulting in a shorter s-t-turnpath, contradicting the minimal length of p. Note that the rerouting in the second case is possible since ◊ is assumed to be not flat. □

We now focus on the rhombi in which p crosses the diagonal twice.

Definition 7.3. *A rhombus ϱ is called* special *if the turnpath p crosses its diagonal twice. If the crossing is in the same direction, then ϱ is called* confluent*, otherwise, if the crossing is in opposite directions, ϱ is called* contrafluent.

By Lemma 7.2, special rhombi are necessarily flat. Recall the slack contributions of a rhombus ϱ introduced in Definition 3.2.

Proposition 7.4. *In a special rhombus ◇, the turnpath p uses exactly two neutral slack contributions.*

The proof proceeds by several steps.

Lemma 7.5. *In a confluent rhombus ◇ the turnpath p uses at least the two contributions ◇ and ◇.*

Proof. Suppose that both ◇ and ◇ occur in p. Then, as before, there are only the two possibilites of (7.1). Since ◇ is flat, the first case is impossible. □

We can now completely determine how p passes through contrafluent rhombi.

Lemma 7.6. *In a contrafluent rhombus ◇, the turnpath p uses the contributions ◇ and ◇ and no other contributions in this rhombus.*

Proof. Assume first that ◇ and ◇ are in p. Then ◇ and ◇ are both not flat. Hence p can be rerouted via ◇, contradicting the minimal length of p.

We are therefore left with the case where ◇ and ◇ are in p. By construction, ◇ and ◇ are not edges of R_f and thus ◇ and ◇ are both turnedges of p. It remains to show that p uses no other contribution in ◇.

Proposition 7.1 combined with the fact that ◇ is flat easily implies that ◇ and ◇ are the only contributions that p may possibly use. We exclude now these two cases.

Suppose that ◇ occurs in p. Then, as ◇ $\in p$, Lemma 7.2 implies that ◇ is flat. However, this contradicts ◇ $\in p$.

We are left with the case that ◇ occurs in p. Then ◇ and ◇ are both contrafluent. Applying what we have learned so far about contrafluent rhombi, we get the situation depicted in Figure 7.1. Note that the depicted triangle of side length 2 is not only contained

Fig. 7.1: The situation in Lemma 7.6. Left: The arrows represent p. All three overlapping rhombi are contrafluent and thus flat. On the right: The turnvertices that can be used for rerouting.

in a flatspace, but it is a flatspace itself: The reason is that p traverses in counterclockwise direction at the border of flatspaces, see Proposition 4.12. This implies that p can be rerouted as seen in Figure 7.1, which is a contradiction to the minimal length of p. □

Lemma 7.7. *1. Contrafluent rhombi cannot overlap with contrafluent or confluent rhombi.*

2. Confluent rhombi cannot overlap.

Proof. 1. Assume that a contrafluent rhombus ◇ overlaps with a shaded confluent or contrafluent rhombus as in Figure 7.2 (the other cases are similar). The turnpath p uses at

Fig. 7.2: The situation in the proof of Lemma 7.7.

least the turnedges drawn in the left figure, where the directions are irrelevant and hence omitted. Hence the turnvertex in the right figure is used by p. But then, Lemma 7.6 implies that ◇ cannot be contrafluent, contradiction!

2. Let $\Diamond$ be confluent and assume that $\Diamond$ and $\Diamond$ occur in p. Assume that $\Diamond$ is confluent and hence p contains $\Diamond$ and $\Diamond$. Then $\Diamond$ is contrafluent and overlapping with $\Diamond$, which is contradicting part one of this lemma. The same argument works for the other three overlapping cases. ∎

Proof of Proposition 7.4. By Lemma 7.6 it remains to consider the case of a confluent rhombus. We improve on Lemma 7.5. If p would use any additional contribution, then $\Diamond$ would overlap with a confluent or contrafluent rhombus, which is impossible due to Lemma 7.7. ∎

7.2. Rigid and critical rhombi. Recall the polyhedron B of bounded hive flows associated with chosen partitions λ, μ, ν. Again we fix $f \in B_{\mathbb{Z}}$ and we fix a *shortest* s-t-turnpath p in R_f. We set

$$\varepsilon := \max\{t \in \mathbb{R} \mid f + t\pi(p) \in B\}, \quad g := f + \varepsilon\pi(p). \tag{7.2}$$

Then we have $\varepsilon > 0$ by Lemma 4.10 and $g \in B$. For the proof of the Shortest Path Theorem 4.8 it suffices to show that $\varepsilon \geq 1$, since then $f + \pi(p) \in B_{\mathbb{Z}}$.

If all rhombi are f-flat, then there are only two possibilities for p, going directly from the right or bottom entrance edge to the left exit edge. In these two cases we clearly have $\varepsilon \geq 1$.

In the following we suppose that not all rhombi are f-flat. We shall argue indirectly and assume that $\varepsilon < 1$ for the rest of this subsection. After going through numerous detailed case distinctions, describing the possible local situations, we will finally end up with a contradiction, which then finishes the proof of the Shortest Path Theorem 4.8. Our main tools will be Proposition 7.4 on special rhombi and the hexagon equality (3.1). Unfortunately, we see no way of considerably simplifying the tedious arguments.

Definition 7.8. *A rhombus is called* critical *if it is not f-flat, but g-flat. Moreover, we call a rhombus* rigid *if it is both f-flat and g-flat.*

Claim 7.9. *There exists a critical rhombus.*

Proof. Let $S \neq \emptyset$ denote the set of rhombi which are not f-flat and consider the continuous function of $t \in \mathbb{R}$

$$F(t) := \min_{\varrho \in S} \sigma(\varrho, f + t\pi(p)).$$

It is sufficient to show that $F(\varepsilon) = 0$.

By the definition (7.2) of ε we have $F(\varepsilon) \geq 0$. Further, for $\varepsilon < t < 1$ we have $f + t\pi(p) \notin B$. Since the flow $f + t\pi(p)$ satisfies the border capacity constraints by construction of R_f, there is a rhombus ϱ with $\sigma(\varrho, f + t\pi(p)) < 0$. We must have $\varrho \in S$, since otherwise $\sigma(\varrho, p) \geq 0$ (cf. Lemma 4.7), which would lead to the contradiction $\sigma(\varrho, f + t\pi(p)) \geq 0$. We have thus shown that $F(t) < 0$. Since t can be arbitrarily close to ε, we get $F(\varepsilon) \leq 0$. Altogether, we conclude $F(\varepsilon) = 0$. ∎

Claim 7.10. *Each critical rhombus ϱ satisfies $\sigma(\varrho, p) \leq -2$.*

Proof. We have $\sigma(\varrho, f) \geq 1$ and $\sigma(\varrho, f + \varepsilon\pi(p)) = 0$, hence $\sigma(\varrho, p) = -\frac{1}{\varepsilon}\, \sigma(\varrho, f)$. Using $0 < \varepsilon < 1$ we conclude $\sigma(\varrho, p) < -1$. ∎

Lemma 7.11. *A rhombus ϱ is rigid iff it is f-flat and p uses in it only neutral contributions. All special rhombi are rigid.*

Proof. The first assertion follows immediately from Lemma 4.4. The second assertion is a consequence of the first and Proposition 7.4. ∎

For the rest of this subsection, we denote by $\Diamond$ a *first critical rhombus* visited by p (a priori it might not be unique, because critical rhombi could overlap). By Claim 7.10, p uses at least two negative slack contributions in $\Diamond$, cf. Lemma 4.4. In particular, p uses at least one turnvertex among $\Diamond, \Diamond, \Diamond, \Diamond$: let $\Diamond \in \{\Diamond, \Diamond\}$ denote the *first one* used by p in this set. (Rotating with 180° we may assume so without loss of generality.) Further, let $\Diamond$ denote the predecessor of $\Diamond$ in p.

Our goal is to analyze the route of p through $\Diamond$ and nearby rhombi. Narrowing down the possibilities will finally lead to a contradiction.

The situation at the boundary of Δ deserves special treatment and is handled in the following claim.

Claim 7.12. *$\Diamond$ is not at the border of Δ.*

Proof. Assume the contrary. Then p enters Δ once via $\Diamond$, $\Diamond$ or $\Diamond$. If p enters over $\Diamond$, then p must use both $\Diamond$ and $\Diamond$ to satisfy $\sigma(\Diamond, p) \leq -2$. Using both $\Diamond$ and $\Diamond$ is prohibited by Claim 7.1. Hence p uses either $\Diamond$ or $\Diamond$ and p also uses $\Diamond$ or $\Diamond$.

We now make a distinction of cases, each of which leads to a contradiction. Hereby, we heavily rely on Proposition 7.4.

Case $\Diamond \cup \Diamond \subseteq p$: Here $\Diamond$ is special and thus $\Diamond$ must be continued as $\Diamond$ and $\Diamond$ must be continued as $\Diamond$. Thus p leaves and enters the same side of Δ.

Case $\Diamond \cup \Diamond \subseteq p$: This is impossible, because p passes $\Diamond$ twice, but not with two neutral contributions as a special rhombus must do.

Case $\Diamond \cup \Diamond \subseteq p$: Here $\Diamond$ is special and $\Diamond$ must be continued as $\Diamond$ and $\Diamond$ must be continued as $\Diamond$. This enables a rerouting via $\Diamond$, which is a contradiction to the minimal length of p.

Case $\Diamond \cup \Diamond \subseteq p$: Note that $\Diamond \subseteq p$, because $\Diamond \subseteq p$ would be a contradiction, because p cannot leave Δ at $\Diamond$. The fact $\Diamond \subseteq p$ implies that p can be rerouted via $\Diamond$, which is a contradiction to the minimal length of p. ∎

We continue with the analysis of the general situation.

Claim 7.13. *$\Diamond$ is not g-flat.*

Proof. By way of contradiction, assume that $\Diamond$ is g-flat. By Claim 7.12 we know that $\Diamond$ is not at the border of Δ. Then $\Diamond$ and $\Diamond$ exist and are both g-flat by Corollary 3.9 applied to g. So they are either rigid or critical. Since $\Diamond$ is not f-flat, it follows from Corollary 3.9 applied to f that not both $\Diamond$ and $\Diamond$ are rigid, so at least one of them is critical. It remains to exclude the following two cases:

If $\Diamond$ is critical, then the critical rhombus $\Diamond$ is passed by p before $\Diamond$, contradicting the minimal choice of $\Diamond$.

If $\Diamond$ is rigid, then $\Diamond \notin V(R_f)$ by the definition of R_f and hence $\Diamond = \Diamond$ and p passes the critical rhombus $\Diamond$ before $\Diamond$. This again contradicts the minimal choice of $\Diamond$. ∎

Claim 7.14. *The turnpath p goes directly from $\Diamond$ to $\Diamond$, which is the only time that p leaves $\Diamond$ over $\Diamond$. Additionally, p leaves $\Diamond$ exactly once over $\Diamond$. In particular, $\sigma(\Diamond, p) = -2$, $\sigma(\Diamond, f) = 1$ and $\varepsilon = \frac{1}{2}$.*

Proof. We prove the following claims, tacitly using Proposition 7.4 on special rhombi:

(i) If p leaves $\Diamond$ at $\Diamond$, then p goes directly from $\Diamond$ to $\Diamond$. Proof: Otherwise p could be rerouted via $\Diamond$, a contradiction.

(ii) p does not leave $\Diamond$ twice at $\Diamond$. Proof: Otherwise $\Diamond$ would be special and hence p would use both $\Diamond$ and $\Diamond$. Then p could be rerouted via $\Diamond$, a contradiction.

(iii) p does not leave $\Diamond$ twice at $\Diamond$. Proof: Otherwise $\Diamond$ would be special and hence rigid by Lemma 7.11. However, this is prohibited by Claim 7.13.

The fact that p leaves $\Diamond$ over $\Diamond$ at most once and over $\Diamond$ at most once implies $\sigma(\Diamond, p) = \Diamond(p) + \Diamond(p) \geq -2$. On the other hand, since $\Diamond$ is critical, we have $\sigma(\Diamond, p) \leq -2$ by Claim 7.10. Therefore, $\sigma(\Diamond, p) = -2$. Hence $\sigma(\Diamond, f) = -\varepsilon\sigma(\Diamond, p) = 2\varepsilon$, so $\varepsilon = \frac{1}{2}\sigma(\Diamond, f)$ and since $0 < \varepsilon < 1$ we obtain $\sigma(\Diamond, f) = 1$ and $\varepsilon = \frac{1}{2}$. ∎

Despite the fact that, due to Claim 7.14, p enters $\Diamond$ at $\Diamond$ and leaves $\Diamond$ at $\Diamond$, p cannot be rerouted via $\Diamond$ to a shorter s-t-turnpath, because p has minimal length. This can have two reasons, which leads to the following namings (compare (3.2)):

> If $\Diamond$ is f-flat and p uses $\Diamond$, then we say that p *enters nonreroutably*, otherwise p *enters reroutably*. If $\Diamond$ is f-flat and p uses $\Diamond$, then we say that p *leaves nonreroutably*, otherwise p *leaves reroutably*.

An explanation of these namings can be found in the proof of the following claim.

Claim 7.15. *The turnpath p enters nonreroutably or leaves nonreroutably.*

Proof. Assume the contrary, i.e., p enters reroutably and leaves reroutably. Recall that ◊ is critical and hence not f-flat. We make a distinction of four cases.

1. Let ◊ be not f-flat and ◊ not be f-flat. Then p can be rerouted using ◊.

2. Let ◊ be f-flat and ◊ not be f-flat. Then p uses ◊ by our assumption on p at the beginning of the proof. Hence p can be rerouted with ◊.

3. Let ◊ not be f-flat and ◊ be f-flat. Assume that p uses ◊. Then p uses ◊, which is a contradiction to the fact that p leaves reroutably. Hence p uses ◊ and p can be rerouted using ◊.

4. Let ◊ and ◊ be both f-flat. Then p uses ◊ and ◊ by our assumption on p at the beginning of the proof. Hence p can be rerouted with ◊ . □

The possible reroutability of p upon entering or leaving the critical rhombus ◊ gives rise to a distinction of four cases, one of which is dealt with in Claim 7.15. In the rest of this subsection we deal with the other three cases. But first we prove the following auxiliary Claim 7.16.

CLAIM 7.16. ◊ *is not rigid.*

Proof. Assume the contrary. Then, according to Lemma 7.11 and Claim 7.14, p uses ◊. Hence p enters nonreroutably. Corollary 3.9 implies that ◊ and ◊ are both g-flat. The hexagon equality (3.1) and $\sigma(◊, f) = 1$ imply that $\sigma(◊, f) + \sigma(◊, f) = 1$. Integrality of f implies that one of the two shaded rhombi of ◊ and ◊ is critical and the other one is rigid.

The rhombus ◊ cannot be critical since otherwise p would pass the critical ◊ before ◊, which contradicts the choice of ◊ as the first critical rhombus in which p uses turnvertices. Hence ◊ is rigid, which implies that ◊ is not a turnvertex of R_f. Thus p uses ◊ and hence p passes the critical rhombus ◊ before ◊. Again, this is a contradiction.

Finally, if ◊ lies at the border of Δ, then, according to Claim 7.14, p enters and leaves Δ over the same side, which is a contradiction to the minimal length of p. □

CLAIM 7.17. *p enters reroutably.*

Proof. We suppose the contrary, so assume that p uses ◊ and ◊ is f-flat. Since ◊ is not rigid by Claim 7.16, p must use a positive slack contribution in ◊. We claim that only ◊ is possible. This is shown by the following case distinction, leading to contradictions in all three cases.

1. ◊ uses a turnvertex already used by p.

2. ◊ $\in p$ implies that p leaves ◊ over ◊ more than once, which is impossible due to Claim 7.14.

3. ◊ $\in p$ contradicts Proposition 7.1.

The fact ◊ $\in p$ implies that ◊ is special and hence we have ◊ $\in p$. Corollary 3.9 implies that both ◊ and ◊ are f-flat. Moreover, $\sigma(◊, f) = 1$ (see Claim 7.14) and the hexagon equality (3.1) implies $\sigma(◊, f) = 1$. Since ◊ $\notin V(R_f)$ and ◊ $\notin E(R_f)$, p must leave ◊ at ◊ via ◊, see Figure 7.3(a).

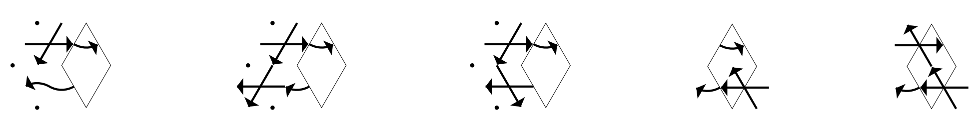

(a) (b) (c) (d) (e)

Fig. 7.3: The situations in the proof of Claims 7.17–7.18. We use straight arrows to depict parts of p which are contained in special rhombi.

(i) If p continues from ◊ to ◊, then ◊ is special, see Figure 7.3(b). Slack computation shows that $\sigma(◊, p) = -2$, which implies with $\varepsilon = \frac{1}{2}$ (Claim 7.14) that $\sigma(◊, g) = 0$ and hence ◊ is critical. But p passes ◊ before ◊, in contradiction with the choice of ◊.

(ii) If p continues from ◊ to ◊, then ◊ is special (see Figure 7.3(c)). Note that ◊ $\in p$ implies that the rhombi ◊ and ◊ are not f-flat. Therefore we can reroute p via ◊, which is in contradiction to the minimal length of p. □

Claim 7.15 and Claim 7.17 imply that p enters reroutably and leaves nonreroutably. It remains to show that this leads to a contradiction. Since ◊ is f-flat and Claim 7.13 ensures that ◊ is not g-flat, p must use a positive slack contribution in ◊.

Claim 7.18. *From the positive slack contributions of p in ◊, only ◊ is possible.*

Proof. The turnedge ◊ uses a turnvertex already used by p. The turnedge ◊ $\subseteq p$ contradicts Proposition 7.1. Now assume that ◊ $\subseteq p$. Then ◊ is special and in particular f-flat. Corollary 3.9 implies that ◊ and ◊ are f-flat. Since p enters reroutably, p uses ◊, which is a contradiction to ◊ $\notin V(R_f)$. □

It follows that ◊ is special. The situation is depicted in Figure 7.3(d). Now ◊ cannot continue to ◊, because, according to Claim 7.14, p leaves ◊ over ◊ only once. Hence p continues to ◊ and ◊ is special, see Figure 7.3(e). But then, p enters nonreroutably, which is a contradiction to Claim 7.17.

This shows that the assumption $\varepsilon < 1$ is absurd. Hence the Shortest Path Theorem 4.8 is finally proved.

8. Proof of the Connectedness Theorem. Recall from Section 2.2 the linear isomorphism $\overline{F}(G) \simeq Z$ mapping flow classes f to their throughput function $E(\Delta) \to \mathbb{R}, k \mapsto \delta(k, f)$. The L_1-norm on Z induces the following *norm* on $\overline{F}(G)$: for $f \in \overline{F}(G)$, we set

$$\|f\| := \textstyle\sum_{k \in E(\Delta)} |\delta(k, f)|.$$

Correspondingly, we define the *distance* between $f, g \in \overline{F}(G)$ as $\mathrm{dist}(f, g) = \|g - f\|$,

In order to prove Theorem 3.12 we have to show that for all $f, g \in P(\lambda, \mu, \nu)_{\mathbb{Z}}$ there exists a finite sequence $f = f_0, f_1, f_2, \ldots, f_\ell = g$ such that each f_{i+1} equals $f_i + c_i$ for some f_i-secure cycle c_i (cf. Proposition 3.8). We will construct this sequence with the additional property that $\mathrm{dist}(f_{i+1}, g) < \mathrm{dist}(f_i, g)$ for all i. To achieve this, it suffices to show the following Proposition 8.1.

Proposition 8.1. *For all distinct $f, g \in P(\lambda, \mu, \nu)_{\mathbb{Z}}$ there exists an f-secure cycle c such that* $\mathrm{dist}(f + c, g) < \mathrm{dist}(f, g)$.

In the rest of this section we prove Proposition 8.1 by explicitly constructing c. We fix $f, g \in P(\lambda, \mu, \nu)_{\mathbb{Z}}$ and set $d := g - f$. We can ensure the distance property $\mathrm{dist}(f + c, g) < \mathrm{dist}(f, g)$ for a proper cycle c by using the following result.

Lemma 8.2. *If more than half of the edges of a proper cycle c on G are contained in* $\mathrm{supp}(d)$, *then we have* $\mathrm{dist}(f + c, g) < \mathrm{dist}(f, g)$.

Proof. Let $K \subseteq E(\Delta)$ be the set of edges of Δ crossed by c. Then

$$\mathrm{dist}(f, g) - \mathrm{dist}(f + c, g) = \sum_{k \in E(\Delta)} \big(|\delta(k, d)| - |\delta(k, d - c)|\big) = \sum_{k \in K} \big(|\delta(k, d)| - |\delta(k, d - c)|\big).$$

But for edges ◊ $\in K$ we easily calculate

$$|\delta(\Diamond, d)| - |\delta(\Diamond, d - c)| = \begin{cases} 1 & \text{if } \mathrm{sgn}\big(\Diamond(c)\big) = \mathrm{sgn}\big(\Diamond(d)\big) \\ -1 & \text{if } \mathrm{sgn}\big(\Diamond(c)\big) \neq \mathrm{sgn}\big(\Diamond(d)\big) \end{cases}.$$

If more than half the edges of c are contained in $\mathrm{supp}(d)$, then also more than half of the summands are 1. The claim follows. □ In the light of Lemma 8.2, we will try to make c use as many edges contained in $\mathrm{supp}(d)$ as possible. Note that since f and g are both capacity achieving, we have that

$$\delta(k, d) = 0 \text{ for all edges } k \text{ at the border of } \Delta. \qquad (\dagger)$$

For the construction of c we distinguish two situations. The first one turns out to be considerably easier.

Situation 1. We assume that d does not cross the sides of f-flatspaces, that is,

$$\delta(k,d) = 0 \text{ for all diagonals } k \text{ of non-}f\text{-flat rhombi} \tag{$*$}$$

Claim 8.3. *In the situation* $(*)$*,* $\mathrm{supp}(d)$ *does not contain a path in* G *consisting of two consecutive clockwise turns.*

Proof. Assume the contrary. We create a contradiction by using three type of arguments: (A) Lemma 3.4 on antipodal pairs, (B) the flow conservation laws, and the fact $(*)$.

The following sequence of pictures shows paths contained in $\mathrm{supp}(d)$ and how rules (A), (B) and $(*)$ imply that additional paths are contained in $\mathrm{supp}(d)$. All rhombi that are known to be f-flat are drawn shaded:

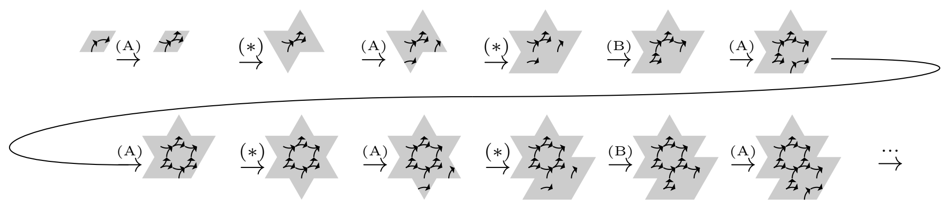

The process of repeatedly applying $\big[$(A), $(*)$, (A), $(*)$, (B), (A)$\big]$ can be continued infinitely while extending the f-flat region to the lower right. This is a contradiction to the finite size of Δ. □

By Lemma 2.1 we have a decomposition $d = \sum_j \alpha_j c_j$ with $\alpha_j > 0$ and cycles c_j in G that are contained in $\mathrm{supp}(d)$. According to $(*)$, each c_j runs in a single f-flatspace and does not cross any f-flatspace border. Let c be any of the cycles c_j and suppose that c runs inside the f-flatspace L. Claim 8.3 implies that c runs in counterclockwise direction. We will next show that L is a hexagon.

Let γ denote the polygon (without self-intersections) obtained from c by linearly interpolating between the successive white vertices of c. Following the white vertices of c (in counterclockwise order) reveals that two consecutive counterclockwise turns lead to an angle of 120° in γ. Further, an alternating sequence of clockwise and counterclockwise turns in c is represented by a line segment in γ. By an elementary geometric argument we see that γ must be a hexagon.

Let $\tilde{c}$ be the counterclockwise cycle *surrounding* c: more specifically, $\tilde{c}$ consists of the clockwise antipodal contributions of all counterclockwise turns in c and, additionally, of the necessary counterclockwise turns in between, as illustrated below:

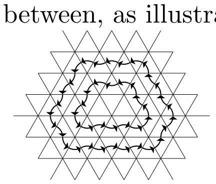

The Flow Propagation Lemma 3.5 implies that all turns of $\tilde{c}$ lying inside L are contained in $\mathrm{supp}(d)$. Hence, by $(*)$, $\tilde{c}$ cannot pass the border of L. Therefore, $\tilde{c}$ either lies completely inside L or completely outside L. If $\tilde{c}$ lies completely inside L, we can form the cycle surrounding $\tilde{c}$ and continue inductively, until we find a cycle $c' \subseteq \mathrm{supp}(d)$ which lies inside L and such that $\tilde{c}'$ lies outside L. Since the polygon γ' corresponding to c' is a hexagon, it follows that L must be a hexagon. Summarizing, we see that the cycle c' runs counterclockwise through the border triangles of a hexagon L. Such c' is clearly f-secure. Moreover, since $c' \subseteq \mathrm{supp}(d)$, we have $\mathrm{dist}(f + c', g) < \mathrm{dist}(f, g)$ by to Lemma 8.2. This proves Proposition 8.1 in Situation 1.

Situation 2. We now treat the case where d has nonzero throughput through some edge k of an f-flatspace L. By (†), k is not at the border of Δ. By Lemma 4.13, we can assume w.l.o.g. that k is an L-entrance edge and $\delta(k, \to L, d) > 0$. Let $p \subseteq \mathrm{supp}(d)$ be a turn in L starting at k.

We will show that the following Algorithm 3 constructs a desired c.

Algorithm 3

Input: $f, g \in P(\lambda,\mu,\nu)_{\mathbb{Z}}$, an edge k such that $\delta(k, \to L, d) > 0$, where $d := g - f$, and a turn p in $\mathrm{supp}(d)$ starting at k.

Output: An f-secure cycle c such that $\mathrm{dist}(f + c, g) < \mathrm{dist}(f, g)$.

1: **while** p does not contain a vertex more than once **do**
2:   **if** one can append to p a turn ϑ contained in $\mathrm{supp}(d)$ such that p after appending does not use a negative contribution in any f-flat rhombus **then**
3:     Append ϑ to p.
4:   **else**
5:     Append a clockwise turn followed by a counterclockwise turn to p.
6:   **end if**
7: **end while**
8: Generate a cycle c from the edges of p by "truncation and concatenation".
9: **return** c.

We postpone the definition of the procedure used in line 8 for the construction of c from p. Later on, we will give a precedence rule to determine what should happen in line 2 when both turns, clockwise and counterclockwise, are possible to append.

What is striking about Algorithm 3 is that it is a priori unclear that line 5 can be executed (without p leaving Δ). We next explain why this is the case.

To ease notation we index the intermediate results that occur during the construction of p by $p_0, p_1, \ldots$, where p_i either has one or two more turns than p_{i-1}, depending on whether in the while loop there has been appended only one turn or (in case of line 5) two turns to p_{i-1}. The paths q_i are defined such that each p_{i+1} is the result of the concatenation of p_i and q_i. We denote by the term *swerve* each q_i that is not a single turn, i.e. those q_i that consist of a clockwise turn followed by a counterclockwise turn. For a swerve q_i we denote by $\varrho(q_i)$ the rhombus in which both turns of q_i lie.

Claim 8.4. *For all i we have the following properties:*

(1) Let  $\in \{$, $\}$ *denote the last turn of p_i and suppose that line 5 is about to be executed. Then* 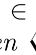 $\in E(\Delta)$ *is not at the border of Δ, which means that $q_i =$*  *can be appended to p_i in line 5 without leaving Δ.*

(2) The first and last edge of each q_i are contained in $\mathrm{supp}(d)$.

(3) Each p_i does not use negative contributions in f-flat rhombi.

(4) The rhombus $\varrho(q_i)$ is f-flat for each swerve q_i.

Before proving Claim 8.4 we start out with a fairly easy lemma that will prove useful.

Lemma 8.5. *Given a walk p in G starting with a turn at a side of an f-flatspace, for some fixed $f \in B$. Further assume that p does not use negative contributions in f-flat rhombi. If the trapezoid*  *consists of two overlapping f-flat rhombi, then p does not end with one of the two turns* 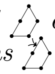 .

Proof. According to the hexagon inequality (3.1), both trapezoids 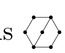 are f-flat. Note that the following three possible cases, which could precede 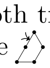, all use negative contributions in f-flat rhombi, which contradicts our assumption:

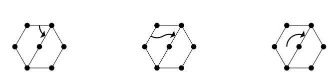

∎

Proof of Claim 8.4. We prove all claims simultaneously by induction on i. If q_{i-1} is a single turn, then q_{i-1} is supported by d and (by definition of the if-clause in Algorithm 3) p_i does not use any negative contributions in f-flat rhombi, which proves (2) and (3) in this case.

It remains to consider the case where q_{i-1} is a swerve, that is, line 5 is about to execute. Let  $\in \{$, $\}$ be the last turn of p_{i-1}. The induction hypothesis (2) ensures $(d) > 0$.

We first show (1). For the sake of contradiction, we assume that the edge  $\in E(\Delta)$ is at the border of Δ. Then, considering (†), it follows that  $\not\subseteq \mathrm{supp}(d)$, but  $\subseteq \mathrm{supp}(d)$.

Thus Algorithm 3 uses line 3 and $q_{i-1} =$ ◊. This is a contradiction to the assumption that line 5 is about to be executed. Hence ◊ $\in E(\Delta)$ is not at the border of Δ. This proves (1).

It remains to show (2), (3) and (4). The fact that line 5 is about to execute can have the following two reasons (a) and (b):

(a) ◊ $\subseteq \mathrm{supp}(d)$, but ◊ cannot be appended to p_{i-1} in line 3.
Then ◊ is f-flat and ◊ $=$ ◊ $\subseteq \mathrm{supp}(d)$ as this turn was appended in line 3. Lemma 8.5 applied to ◊ yields that ◊ is not f-flat. Since ◊ $\subseteq \mathrm{supp}(d)$ we have ◊ $\subseteq \mathrm{supp}(d)$ by Lemma 3.4. Therefore, p_{i-1} can be continued via $q_{i-1} =$ ◊ in line 3, in contradiction to the fact that line 5 is about to execute.

(b) ◊ $\subseteq \mathrm{supp}(d)$, but ◊ cannot be appended to p_{i-1} in line 3.
Then ◊ is f-flat, which shows (4). Lemma 3.4 implies that ◊ $\subseteq \mathrm{supp}(d)$. In line 5, the turns $q_{i-1} =$ ◊ are appended to p_{i-1}, which shows (2). It remains to show that appending q_{i-1} does not result in negative contributions in f-flat rhombi. But if ◊ leads to a negative contribution in an f-flat rhombus, then ◊ is f-flat and if ◊ leads to a negative contribution in an f-flat rhombus, then ◊ is f-flat. In both cases, this contradicts Lemma 8.5, for the f-flat trapezoid ◊ and ◊ , respectively. This shows (3). ∎

We specify now the precedence rule (‡) for breaking ties in line 2 of Algorithm 3.

> If p_{i-1} ends at the diagonal of an f-flat rhombus, then Algorithm 3 appends *counterclockwise* turns; if p_{i-1} ends at the diagonal of a non-f-flat rhombus, then Algorithm 3 appends *clockwise* turns. (‡)

Finally, to fully specify Algorithm 3, we now define how the cycle c is generated from p in line 8: When line 8 is about to execute, then p has used a vertex more than once. Let v denote the first vertex of p that is used more than once. Note that v is a black vertex. Let q denote the q_i that was appended last. We note that q consists of either two edges or four edges. Now we truncate everything of p previous to the first occurence of v and everything after the last occurence of v, thus generating the cycle c. We denote by ϑ the first turn of p that uses v and by ϑ' the turn of c that uses v.

For example, suppose p uses the swerve ◊ and $q =$ ◊. Then $\vartheta =$ ◊ and $\vartheta' =$ ◊. Note that the turn ϑ' is contained in c but not contained in p.

Since p uses no negative contributions in f-flat rhombi, the first assertion of Claim 8.6 below is plausible, but needs proof as c may contain turns that are not contained in p. In fact, we must ensure that no negative contributions exist in c "near v".

Claim 8.6. *(1) The cycle c is f-hive preserving.*
(2) If ϑ' is not used by p, then $\vartheta' \subseteq \mathrm{supp}(d)$.

Let us first show that once Claim 8.6 is shown, we are done.

Proof of Proposition 8.1. We first show that Algorithm 3 produces an f-secure cycle c. Claim 8.6 already tells us that c is f-hive preserving. Assume that c uses both ◊ in a rhombus ◊. Claim 8.4(2) and Claim 8.6(2) imply that at least the second edge of every counterclockwise turn in c is contained in $\mathrm{supp}(d)$. Hence ◊$(d) > 0$ and ◊$(d) > 0$, which implies $\sigma($◊$, d) \le -2$. The fact $\sigma($◊$, f+d) \ge 0$ implies $\sigma($◊$, f) \ge 2$ and hence ◊ is not nearly f-flat. It follows that c is f-secure.

Claim 8.6(2) combined with Claim 8.4(2) also ensures that the only turns in c, that are not contained in $\mathrm{supp}(d)$, are turns of swerves. Hence at least half of the edges of c are contained in $\mathrm{supp}(d)$. This inequality is strict, because c cannot consist of swerves only. Lemma 8.2 implies $\mathrm{dist}(f+c, g) < \mathrm{dist}(f, g)$. ∎

From now on, swerves (e.g. ◊) will be drawn as straight arrows with a filled triangular head, e.g. ◊ . (They are not to be confused with throughput arrows like ◊, which have a different head and are always drawn crossing fat edges.)

Proof of Claim 8.6. Since p uses no negative contributions in f-rhombi by construction, the proof that the cycle c is f-hive preserving breaks down into the following parts:

(neg1) The turn ϑ' is not counterclockwise at the acute angle of an f-flat rhombus.

(neg2) The turn ϑ' is not both clockwise and preceded in c by another clockwise turn such that both turns lie in the same f-flat rhombus.

(neg3) The turn ϑ' is not both clockwise and succeded in c by another clockwise turn such that both turns lie in the same f-flat rhombus.

We also need to prove the following property:

(su) If ϑ' is not used by p, then $\vartheta' \subseteq \mathrm{supp}(d)$.

Recall that q denotes the q_i which was appended last. Three cases can appear: (a) q is a counterclockwise turn, (b) q is a clockwise turn, (c) q is a swerve. All three cases are significantly different and require careful attention to detail. We start with the simplest one, which does not require the precedence rule (‡):

(a) Assume that q is a counterclockwise turn, pictorially $q =$. Considering Algorithm 3, we see that $q \subseteq \mathrm{supp}(d)$. There are two possibilities for ϑ: $\vartheta =$ or $\vartheta =$, because the other four turns lead to a contradiction to the fact that Algorithm 3 stops as soon as p contains a vertex twice.

(a1) Suppose first $\vartheta =$. In this case, we have $\vartheta' = q$. The statement (neg1) holds, because ϑ' is used by p and p uses no negative contributions in f-flat rhombi. We also see that (su) holds in this case, because ϑ' is used by p.

(a2) Suppose now $\vartheta =$. Since $(d) > 0$, it follows that is part of a swerve: The situation of p can be depicted as . By construction, $\vartheta' =$, so ϑ' consists of an edge of q and the last edge of a swerve. Hence $q \subseteq \mathrm{supp}(d)$ and (su) follows in this case. It remains to verify (neg2) and (neg3). Note that Algorithm 3 ensures that the counterclockwise turn q is not a negative contribution in f-flat rhombi and hence the shaded rhombus is not f-flat. This proves (neg3). If we assume the contrary of (neg2), then the path is a negative contribution in an f-flat rhombus. But since swerves lie in f-flat rhombi according to Claim 8.4(4), it follows that the trapezoid is f-flat, which is a contradiction to Lemma 8.5, applied to . This proves (neg2).

(b) Assume that q is a clockwise turn, pictorially $q =$. As in case (a), we have two possibilities: $\vartheta =$ or $\vartheta =$. Since $q \subseteq \mathrm{supp}(d)$, we have $(d) > 0$. If we had $\vartheta =$, then $(d) > 0$, which is a contradiction. Hence $\vartheta =$ and thus $\vartheta' = q$. This proves (su) in this case. The fact that $\vartheta' = q$ is a part of p shows (neg2). It remains to show (neg3). Assume the contrary. Then the rhombus is f-flat. Hence $(d) > 0$ implies $(d) > 0$. The rhombus is not f-flat by Lemma 8.5 applied to . But the precedence rule (‡) of Algorithm 3 implies that p continues from with the counterclockwise turn . This is a contradiction, proving (neg3) in this case.

(c) Assume that q is a swerve, pictorially $q =$. The rhombus which contains the swerve is f-flat by Claim 8.4(4). Since p does not use negative contributions in f-flat rhombi, we get that is not f-flat. The possibilities for ϑ here are , , (note that is ruled out, because is f-flat). We distinguish the following three cases:

(c1) Suppose $\vartheta =$. Here $\vartheta' =$, which is part of q. This proves (su) in this case. The fact that is not f-flat implies (neg1) in this case.

(c2) Suppose $\vartheta =$. Here $\vartheta' =$, which is part of q, which again proves (su) in this case. The fact (neg2) follows because p uses no negative contributions in f-flat rhombi. It remains to show (neg3). Note that $(d) > 0$, because the first edge of is contained in $\mathrm{supp}(d)$. Further, $(d) > 0$, because the second edge of the counterclockwise turn is contained in $\mathrm{supp}(d)$ (this is always the case for counterclockwise turns by construction of p). Hence $\subseteq \mathrm{supp}(d)$. If were not f-flat, then Algorithm 3 would have appended the clockwise turn over appending the swerve . Hence is f-flat. The hexagon equality (3.1) implies that the trapezoid if f-flat. The fact (neg3) follows from Lemma 8.5 applied to .

(c3) Suppose $\vartheta =$. Here $\vartheta' =$, which is a negative contribution in the f-flat rhombus . Hence we need to show that this case leads to a contradiction. Recall that is not f-flat. Clearly, $(d) > 0$ and $(d) > 0$, which implies $(d) > 0$. This means that ϑ is preceded in p by the counterclockwise turn . This is a contradiction to the precedence rule (‡), because Algorithm 3 would have chosen instead of . □

REFERENCES

[AMO93] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.

[BC13] Peter Bürgisser and Felipe Cucker. Condition: The Geometry of Numerical Algorithms. Springer Grundlehren der Mathematischen Wissenschaften. To appear.

[BI09] Peter Bürgisser and Christian Ikenmeyer. A max-flow algorithm for positivity of Littlewood-Richardson coefficients. In *FPSAC 2009, Hagenberg, Austria, DMTCS proc. AK*, pages 267–278, 2009.

[Buc00] Anders S. Buch. The saturation conjecture (after A. Knutson and T. Tao) with an appendix by William Fulton. *Enseign. Math.*, 2(46):43–60, 2000.

[BZ92] Arkady D. Berenstein and Andrei V. Zelevinsky. Triple multiplicities for sl(r + 1) and the spectrum of the exterior algebra of the adjoint representation. *J. Algebraic Comb.*, 1(1):7–22, 1992.

[DLM06] Jesús A. De Loera and Tyrrell B. McAllister. On the computation of Clebsch-Gordan coefficients and the dilation effect. *Experiment. Math.*, 15(1):7–19, 2006.

[FF62] Lester R. Ford and Delbert R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, N.J., U.S.A., 1962.

[Ful97] William Fulton. *Young tableaux*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1997.

[Ful00] William Fulton. Eigenvalues, invariant factors, highest weights, and Schubert calculus. *Bull. Amer. Math. Soc. (N.S.)*, 37(3):209–249 (electronic), 2000.

[GLS93] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*. Springer-Verlag, Berlin, 1993.

[HR95] Uwe Helmke and Joachim Rosenthal. Eigenvalue inequalities and Schubert calculus. *Math. Nachr.*, 171:207–225, 1995.

[Ike08] Christian Ikenmeyer. On the complexity of computing Kronecker coefficients and deciding positivity of Littlewood-Richardson coefficients. Master's thesis, Institute of Mathematics, University of Paderborn, 2008. Online available at `http://math-www.uni-paderborn.de/agpb/work/ikenmeyer_diplom.pdf`.

[Ike12] Christian Ikenmeyer. Small Littlewood-Richardson coefficients. arXiv:1209.1521 [math.RT], 2012.

[Kly98] Alexander A. Klyachko. Stable bundles, representation theory and Hermitian operators. *Selecta Math. (N.S.)*, 4(3):419–445, 1998.

[KT99] Allen Knutson and Terence Tao. The honeycomb model of $\mathrm{GL}_n(\mathbf{C})$ tensor products. I. Proof of the saturation conjecture. *J. Amer. Math. Soc.*, 12(4):1055–1090, 1999.

[KTT04] Ron C. King, Christophe Tollu, and Frédéric Toumazet. Stretched Littlewood-Richardson and Kostka coefficients. In *Symmetry in physics*, volume 34 of *CRM Proc. Lecture Notes*, pages 99–112. Amer. Math. Soc., Providence, RI, 2004.

[KTW04] Allen Knutson, Terence Tao, and Christopher Woodward. The honeycomb model of GL(n) tensor products II: Puzzles determine facets of the Littlewood-Richardson cone. *J. Amer. Math. Soc.*, 17(1):19–48, 2004.

[MS01] Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory. I. An approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526 (electronic), 2001.

[MS05] Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory III: On deciding positivity of Littlewood-Richardson coefficients. cs.ArXive preprint cs.CC/0501076, 2005.

[MS08] Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory. II. Towards explicit obstructions for embeddings among class varieties. *SIAM J. Comput.*, 38(3):1175–1206, 2008.

[Mul11] Ketan D. Mulmuley. On P vs. NP and geometric complexity theory. *J. ACM*, 58(2):Art. 5, 26, 2011.

[Nar06] Hariharan Narayanan. On the complexity of computing Kostka numbers and Littlewood-Richardson coefficients. *J. Algebraic Combin.*, 24(3):347–354, 2006.

[PV05] Igor Pak and Ernesto Vallejo. Combinatorics and geometry of Littlewood-Richardson cones. *Eur. J. Comb.*, 26(6):995–1008, 2005.