

Concave Cocirculations in a Triangular Grid

Alexander V. Karzanov

*Institute for System Analysis
9, Prospect 60 Let Oktyabrya, 117312 Moscow, Russia
Email: sasha@cs.isa.ru*

Abstract

Let G be a planar digraph embedded in the plane such that each bounded face contains three edges and forms an equilateral triangle, and let the union \mathcal{R} of these faces be a convex polygon. We consider the polyhedral cone $\mathcal{B}(G)$ formed by the real-valued functions σ on the set of boundary edges of G with the following property: there exists a concave function c on \mathcal{R} which is affinely linear within each bounded face and satisfies $c(v) - c(u) = \sigma(e)$ for each boundary edge $e = (u, v)$.

Knutson, Tao and Woodward obtained a result on honeycombs which implies that if the polygon \mathcal{R} is a triangle, then the cone $\mathcal{B}(G)$ is described by linear inequalities of Horn's type with respect to so-called *puzzles*, along with obvious linear constraints.

The purpose of this paper is to give an alternative proof of that result, working in terms of discrete concave functions, rather than honeycombs. Our proof is based on a linear programming approach and a non-standard flow model. Moreover, the result is extended to an arbitrary convex polygon \mathcal{R} as above.

Keywords: Triangular lattice, Discrete convex function, Cocirculation, Planar graph, Flow

1 Introduction

Let ξ_1, ξ_2, ξ_3 be three affinely independent vectors in the plane \mathbb{R}^2 whose sum is the zero vector. The triangular lattice generated by ξ_1, ξ_2, ξ_3 is associated with the infinite planar directed graph \mathcal{L} whose vertices are integer combinations of these vectors and whose edges are the ordered pairs (u, v) of vertices such that $v - u \in \{\xi_1, \xi_2, \xi_3\}$. An edge (u, v) is identified with the straight-line segment between u, v oriented from u to v .

Consider a convex region \mathcal{R} in the plane formed by the union of a nonempty finite set of faces (*little triangles*) of \mathcal{L} ; it is a polygon with 3 to 6 sides. We refer to the subgraph $G = (V(G), E(G))$ of \mathcal{L} consisting of the vertices and edges occurring in \mathcal{R} as a *convex (triangular) grid*. The sets of vertices and edges in the *boundary* $b(G)$ of G are denoted by $V_0(G)$ and $E_0(G)$, respectively.

A real-valued function f on the vertices of G is called *discrete concave (convex)* if its piece-wise linear extension c to the region \mathcal{R} is a concave (resp. convex) function (here c is affinely linear within each little triangle of G and coincides with f on $V(G)$). In this paper we prefer to deal with discrete concave functions; the corresponding results for discrete convex functions follow by symmetry.

We are interested in the functions on the set of boundary vertices that can be extended to discrete concave functions on all vertices of G . Instead, one can consider the corresponding functions on edges. More precisely, a function $h : E(G) \rightarrow \mathbb{R}$ is said to be a *circulation* if there exists $f : V(G) \rightarrow \mathbb{R}$ such that $h(e) = f(v) - f(u)$ for each edge $e = (u, v)$. Such an h determines f up to a constant, and we refer to h as a *concave circulation* if f is discrete concave. In these terms, the problem is:

(1.1) Given a function $\sigma : E_0(G) \rightarrow \mathbb{R}$, decide whether σ is extendable to a concave circulation in G .

Clearly the set $\mathcal{B}(G)$ of functions σ admitting such an extension forms a convex cone in $\mathbb{R}^{E_0(G)}$. Two necessary (but far to be sufficient) conditions on σ to belong to $\mathcal{B}(G)$ are obvious: (i) the sum of values of σ , taken with signs + or - depending on the direction of an edge in the boundary circuit, amounts to zero, and (ii) σ is weakly decreasing along each side-path of $b(G)$.

Interesting results have been obtained for the case when the polygon \mathcal{R} spanned by G is a triangle. In this case the boundary of G is the concatenation of three paths B_1, B_2, B_3 forming the sides of \mathcal{R} , where the edges of B_i are parallel to ξ_i . We refer to G as a *3-side grid* and say that G has *size* n if $|B_i| = n$ (where $|P|$ denotes the number of edges of a path P). It turned out that the cone $\mathcal{B}(G)$ for a 3-side grid G also arises in two other interesting models. More precisely, Knutson and Tao [6] showed that for a triple of weakly decreasing n -tuples $(\lambda, \mu, \nu) \in (\mathbb{R}^n)^3$, the following properties are equivalent:

(P1) λ, μ, ν are the spectra of three $n \times n$ Hermitian matrices whose sum is the zero matrix;

(P2) there exists a honeycomb of size n in which the three tuples of semiinfinite edges have the constant coordinates λ, μ, ν (see [6] for a definition);

(P3) let G be the 3-side grid of size n and let σ be the function on $E_0(G)$ taking the value λ_j (resp. μ_j, ν_j) on j th edge of the path B_1 (resp. B_2, B_3); then $\sigma \in \mathcal{B}(G)$.

Note that while the equivalence of (P2) and (P3) is rather transparent (they are related via Fenchel's duality), the equivalence of these to (P1) is quite sophisticated. In the 1960s Horn [4] recursively constructed a finite list of nontrivial necessary conditions on λ, ν, μ to satisfy property (P1) and conjectured the sufficiency of this list (which, in particular, implies that these (λ, μ, ν) 's constitute a polyhedral cone). Horn's conditions are viewed as linear inequalities of the form

$$(1.2) \quad \lambda(I) + \mu(J) + \nu(K) \geq 0$$

for certain subsets I, J, K of $\{1, \dots, n\}$ with $|I| = |J| = |K|$, letting $\alpha(S) := \sum(\alpha_i : i \in S)$ for $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n$ and $S \subseteq \{1, \dots, n\}$. Subsequent efforts of several authors have resulted in a proof of Horn's conjecture; the obtained result is referred in [7] as the "H-R/T/K theorem", abbreviating the names of Helmke, Rosenthal, Totaro, and Klyachko. (A history of studying problems concerning (P1) and related topics are reviewed in [3].) Recently Knutson, Tao and Woodward [7] established a combinatorial existence criterion for honeycombs, obtaining another proof of that theorem, in view of the equivalence between (P2) and (P1). According to their criterion, each Horn's triple (I, J, K) is induced by a *puzzle*, a certain subdivision of a 3-side grid into little triangles and little rhombi endowed with a certain 0,1 labelling on the sides of these pieces. One more method of proof of the H-R/T/K theorem is given by Danilov and Koshevoy [2].

The purpose of this paper is to give a direct proof of the puzzle criterion for the solvability of problem (1.1), without using relationships to honeycombs. We extend the notion of puzzle in a natural way to an arbitrary convex grid G and show that $\sigma : E_0(G) \rightarrow \mathbb{R}$ is extendable to a concave cocirculation if and only if it obeys the linear inequalities of Horn's type determined by puzzles and the above-mentioned obvious linear constraints. Our proof combines a linear programming approach and some combinatorial techniques where a non-standard flow model is involved.

This paper is organized as follows. Section 2 contains basic definitions and facts and states problem (1.1) as a linear program. In Section 3 we explain the notion of puzzle for a convex grid (using a definition somewhat different from, but equivalent to, that in [7]) and formulate the puzzle criterion for the solvability of the problem (Theorem 3.1). The proof of this theorem is given in Section 6, based on a weaker, linear programming, criterion discussed in Sections 4 and on a representation of dual variables by use of a flow in an auxilliary graph, explained in Section 5. The concluding Section 7 discusses some additional aspects and a generalization to convex grids of results in [7] on the puzzles determining facets of the cone $\mathcal{B}(G)$ for a 3-side grid G , where these puzzles are characterized combinatorially and in terms of rigidity.

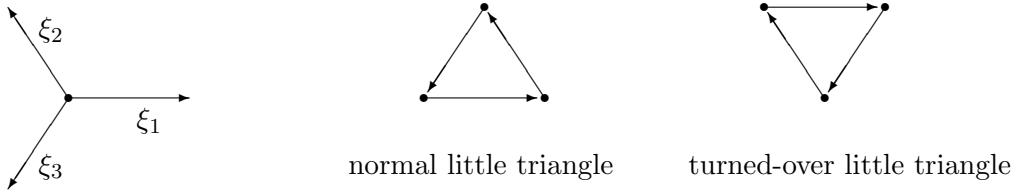
Related and other aspects of discrete convex (concave) functions on triangular grids in the plane and some applications in algebra are discussed in [1].

2 Preliminaries

We start with terminology, notation and conventions. Edges, faces, subgraphs, paths, circuits and other relevant objects in a convex grid G or another graph in question are usually identified

with their *closed* images in the plane. By a path (circuit) we usually mean a simple directed path (circuit) $P = (v_0, e_1, v_1, \dots, e_k, v_k)$, where e_i is the edge (v_{i-1}, v_i) ; it may be abbreviatedly denoted as (e_1, e_2, \dots, e_k) (via edges). A path P with beginning vertex u and end vertex v is called a *u–v path*; P is called *degenerate* if it consists of only one vertex $u = v$. When P forms a straight-line segment in the plane, P is called a *straight path*, or a *line* of the graph. A *k-circuit* is a circuit with k edges.

Problem (1.1) does not depend, in essence, on the choice of lattice generating vectors ξ_1, ξ_2, ξ_3 , and for convenience we fix these vectors as $\xi_1 = (1, 0)$, $\xi_2 = (-1, \sqrt{3})/2$ and $\xi_3 = (-1, -\sqrt{3})/2$. Then the little triangles in G are equilateral triangles of *size* 1. Note that the boundary of any triangle in G (formed by the union of some faces) is a circuit directed clockwise or anticlockwise around the triangle. The little triangle surrounded by a 3-circuit C is denoted by Δ_C . We say that a triangle is *normal* if its boundary circuit is directed anticlockwise, and *turned-over* otherwise.



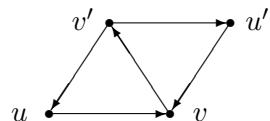
We denote the sets of boundary edges directed anticlockwise and clockwise (around \mathcal{R}) by $E_0^+(G)$ and $E_0^-(G)$, respectively. A maximal straight path in $b(G)$, or a *side-path* of G , whose edges are parallel to ξ_i and belong to $E_0^+(G)$ (resp. $E_0^-(G)$) is denoted by B_i^+ (resp. B_i^-). One may assume that if G is a 3-side grid, then the boundary of G is formed by B_1^+, B_2^+, B_3^+ .

For a function h on $E(G)$, its restriction to the set of boundary edges is called the *border* of h .

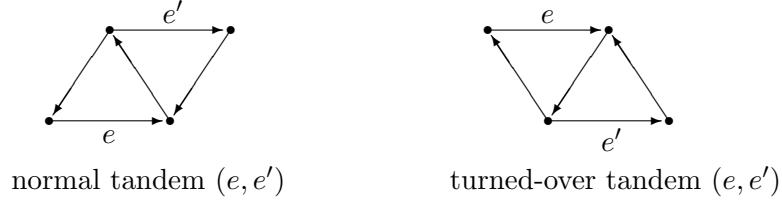
Next we explain how to formulate problem (1.1) as a linear program. Obviously, a function $f : V(G) \rightarrow \mathbb{R}$ is discrete concave if and only if

$$(2.1) \quad f(u) + f(u') \leq f(v) + f(v')$$

holds for each *little rhombus* (the union of two little triangles sharing a common edge) ρ , where u, u' are the *acute* vertices and v, v' are the *obtuse* vertices of ρ :



Clearly $h \in \mathbb{R}^{E(G)}$ is a cocirculation if and only if the sum of its values on each 3-circuit is zero. Linear constraints reflecting the property of a cocirculation h to be concave are derived from (2.1). Let us say that an ordered pair $\tau = (e, e')$ of non-adjacent edges of G is a *tandem* if they occur as opposite sides of a little rhombus ρ and the head of e is an obtuse vertex of ρ (while the other obtuse vertex of ρ is the tail of e'). We distinguish between two sorts of tandems by specifying τ as a *normal* tandem if the little triangle in ρ containing e is normal, and a *turned-over* tandem otherwise. Note that each little rhombus ρ involves two tandems one of which is normal and the other is turned-over. The picture illustrates the case when e, e' are parallel to ξ_1 .



For the cocirculation h generated by a function f on the vertices, (2.1) is just equivalent to the condition $h(e) \geq h(e')$ on the normal tandem (e, e') in the little rhombus ρ . Thus, given $\sigma \in \mathbb{R}^{E_0(G)}$, a concave cocirculation with border σ is a solution $h \in \mathbb{R}^{E(G)}$ of the system:

$$(2.2) \quad h(e) + h(e') + h(e'') = 0, \quad C = (e, e', e'') \in \mathcal{C}(G),$$

$$(2.3) \quad h(e') - h(e) \leq 0, \quad \tau = (e, e') \in \mathcal{T}(G),$$

$$(2.4) \quad h(e) = \sigma(e), \quad e \in E_0(G),$$

where $\mathcal{C}(G)$ is the set of 3-circuits (considered up to cyclically shifting), and $\mathcal{T}(G)$ the set of normal tandems in G . When this system has a solution, we call σ *feasible*.

As mentioned in the Introduction, two necessary conditions on σ to be feasible are obvious. The first one (necessary for the border of any cocirculation) is the *zero-sum condition*:

$$(2.5) \quad \sigma(E_0^+(G)) - \sigma(E_0^-(G)) = 0,$$

The second one is the *monotone condition*:

$$(2.6) \quad \sigma(e_1) \geq \dots \geq \sigma(e_n) \text{ for each straight path } (e_1, \dots, e_n) \text{ in } b(G).$$

Since the set of concave cocirculations on G is described by a finite number of linear constraints, the cone $\mathcal{B}(G)$ formed by all feasible σ 's (the borders of concave cocirculations in G) is polyhedral. To compute the dimension of this cone is easy (cf. [7]).

Statement 2.1 $\dim(\mathcal{B}(G)) = |E_0(G)| - 1$.

Proof. In view of (2.5), $\dim(\mathcal{B}(G)) \leq |E_0(G)| - 1 =: r$. To show the reverse inequality, we first construct a concave cocirculation h for which all tandem inequalities in (2.3) are strict.

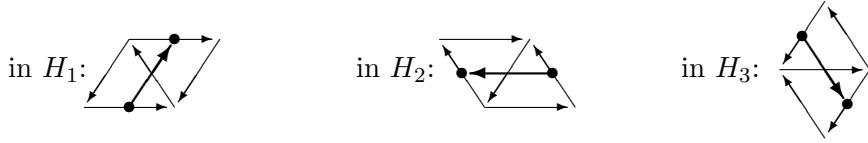
Take a maximal straight $u - v$ path P of G not contained in $b(G)$. Let Z be the set of edges of G that lie in the region on the right from P (when moving from u to v) and are not parallel to P . Define $h_P(e)$ to be 1 if $e \in Z$ and e points toward P , -1 for the other edges e in Z , and 0 for the remaining edges of G . One can check that h_P is a concave cocirculation and that $h(e) > h(e')$ for each tandem (e, e') where e and e' are separated by P . The sum of h_P 's over all such paths P gives the desired concave cocirculation h . Let σ be the border of h .

Now for each boundary vertex v and each edge e , define $h_v(e)$ to be 1 if v is the head of e , -1 if v is the tail of e , and 0 otherwise. Then h_v is a cocirculation; moreover, $h + \frac{1}{2}h_v$ is a concave cocirculation. Let σ_v be the border of h_v . Clearly r borders among these σ_v are linearly independent. This implies that r borders $\sigma + \frac{1}{2}\sigma_v$ of the concave cocirculations $h + \frac{1}{2}h_v$ are linearly independent. ■

3 Theorem

Linear programming suggests a standard way to obtain a solvability criterion for system (2.2)–(2.4). Our aim, however, is to obtain a sharper, combinatorial, characterization for the borders of concave cocirculations on G .

First of all we construct a certain dual digraph H . For each edge $e \in E(G)$, take the middle point v_e on e , making it a vertex of H . For each normal tandem $\tau = (e, e')$, form (straight-line) edge a_τ from v_e to $v_{e'}$, making it an edge of H . Note that when e, e' are parallel to ξ_i , the edge a_τ is *anti-parallel* to ξ_{i-1} , in the sense that a_τ is a parallel translate of the opposite vector $-\xi_{i-1}$. (Hereinafter the corresponding indices are taken modulo 3.) The resulting graph H is the union of three disjoint digraphs H_1, H_2, H_3 , where H_i is induced by the introduced edges connecting points on edges of G parallel to ξ_i . The three types of edges of H are drawn in bold in the picture.



So the maximal paths in H_i are straight, pairwise disjoint and anti-parallel to ξ_{i-1} . If a path P of H begins at v_e and ends at $v_{e'}$, we say that P *leaves* the edge e and *enters* the edge e' (both e, e' concern G), admitting the case of degenerate P . We also say that P leaves (enters) a little triangle Δ if $e \subset \Delta$ (resp. $e' \subset \Delta$).

Definition. A *puzzle* is a pair $\Pi = (\mathcal{F}, \mathcal{P})$ consisting of a set \mathcal{F} of little triangles of G and a set \mathcal{P} of paths of H such that:

- (3.1) (i) the interiors of all triangles in \mathcal{F} and all paths in \mathcal{P} are pairwise disjoint;
- (ii) for each edge e of each normal (resp. turned-over) triangle in \mathcal{F} , there is precisely one path in \mathcal{P} entering (resp. leaving) e ;
- (iii) for each path in \mathcal{P} leaving edge e and entering edge e' , either e belongs to a turned-over triangle in \mathcal{F} or $e \in E_0^+(G)$, and similarly, either e' belongs to a normal triangle in \mathcal{F} or $e \in E_0^-(G)$.

(Degenerate paths $P = v_e$ in \mathcal{P} are admitted. When e is an inner edge of G , such a P serves to “connect” the pair of triangles in \mathcal{F} sharing the edge e . When e is a boundary edge, P “connects” this edge with the triangle in \mathcal{F} containing e .) The *boundary* $b(\Pi)$ of Π is defined to be the set of boundary edges e for which there is a path in \mathcal{P} leaving or entering e . The subsets of edges of $E_0^+(G)$ and $E_0^-(G)$ occurring in $b(\Pi)$ are denoted by $b^+(\Pi)$ and $b^-(\Pi)$, respectively.

The *puzzle criterion* for the solvability of (2.2)–(2.4) is the following.

Theorem 3.1 *Let G be a convex grid, and let $\sigma : E_0(G) \rightarrow \mathbb{R}$ satisfy (2.5) and (2.6). Then a concave cocirculation h in G with $h(e) = \sigma(e)$ for all $e \in E_0(G)$ exists if and only if*

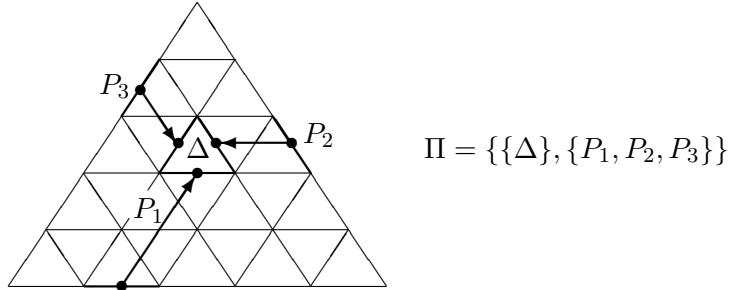
$$(3.2) \quad \sigma(b^+(\Pi)) - \sigma(b^-(\Pi)) \geq 0$$

holds for each puzzle Π .

Thus, the cone $\mathcal{B}(G)$ is described by the *puzzle inequalities* (3.2) and the linear constraints (2.5) and (2.6).

Remark. A puzzle in a 3-side grid G introduced in Knutson et al. [7] is defined to be a diagram D consisting of a subdivision of the big triangle \mathcal{R} into little triangles and little rhombi of G , and of a 0,1 labelling of the edges of G that are sides of these pieces, satisfying the following conditions: (a) the three sides of each little triangle in the subdivision are labelled either 1,1,1 or 0,0,0, and (b) the sides of each little rhombus ρ are labelled 0,1,0,1, in this order clockwise of an acute vertex of ρ . The boundary $b(D)$ of D is defined to be the set of boundary edges labelled 1. There is a natural one-to-one correspondence between the puzzles D of this form and the puzzles $\Pi = (\mathcal{F}, \mathcal{P})$ in the above definition (in the *triangle-path* form) and this correspondence preserves the puzzle boundary: $b(D) = b(\Pi)$. (In this correspondence, \mathcal{F} is set of little triangles labelled 1,1,1, and the edges of H used in the paths of \mathcal{P} are those connecting the sides labelled 1 in the rhombi of D .) The triangle-path form of puzzle is more convenient for us to handle in the proof of Theorem 3.1, which is based on certain path and flow constructions.

To illustrate the theorem, consider a 3-side grid of size n and a puzzle consisting of one triangle Δ and three paths P_1, P_2, P_3 , each P_i connecting Δ with the side-path $B_i^+ = (b_i^1, \dots, b_i^n)$:



Let P_i leave edge $b_i^{r(i)} \in B_i$ and enter edge $e_i \subset \Delta$. Summing up the inequalities in (2.3) for the normal tandems induced by the edges of P_i , we have $\sigma(b_i^{r(i)}) = h(b_i^{r(i)}) \geq h(e_i)$. This together with (2.2) for the 3-circuit (e_1, e_2, e_3) implies that the sum of values of σ on $b_i^{r(i)}$, $i = 1, 2, 3$, is nonnegative. Also $r(1) + r(2) + r(3) = n + 2$. Thus, any feasible $\sigma = (\lambda, \mu, \nu) \in (\mathbb{R}^n)^3$ must obey

$$\lambda_i + \mu_j + \nu_k \geq 0$$

for any choice of i, j, k with $i + j + k = n + 2$. This is the simplest sort of Horn's inequality (1.2).

One can associate with a puzzle $\Pi = (\mathcal{F}, \mathcal{P})$ undirected graph Γ_Π whose vertices correspond to the triangles in \mathcal{F} and the edges in $b(\Pi)$ and where vertices u, v are connected by an edge if and only if there is a path in \mathcal{P} leaving one and entering the other of u, v . One can see that such graphs are determined, up to isomorphism, by the list of cardinalities $|b(\Pi) \cap B|$, where B ranges over the side-paths of G . In particular,

(3.3) the numbers $|\mathcal{F}|$ and $|\mathcal{P}|$ are determined by $b(\Pi)$.

(Instruction: deform G so that each little triangle of G that neither belongs to \mathcal{F} nor meets a path in \mathcal{P} is shrunk into a point, and for each nondegenerate $v_e - v_{e'}$ path in Π , the parallelogram

with opposite sides e, e' is shrunk into the edge e . The resulting graph G' is again a convex grid (possibly degenerate) in which the little triangles one-to-one correspond to those in \mathcal{F} , and the edges to the paths in \mathcal{P} ; also the boundary edges of G' one-to-one correspond to the edges in $b(\Pi)$ when $\mathcal{F} \neq \emptyset$. Moreover, G' depends only on the above-mentioned cardinalities.)

4 Linear Programming Approach

In what follows, speaking of a tandem, we always mean a normal tandem in G . Assign a variable $z(C) \in \mathbb{R}$ to each 3-circuit C of G , a variable $g(\tau) \in \mathbb{R}_+$ to each tandem τ , and a variable $d(e) \in \mathbb{R}$ to each boundary edge e . Then the linear system dual of (2.2)–(2.4) is viewed as

$$(4.1) \quad \sum_{C \in \mathcal{C}(G): e \in C} z(C) - \sum_{\tau=(e,e') \in \mathcal{T}(G)} g(\tau) + \sum_{\tau=(e',e) \in \mathcal{T}(G)} g(\tau) = 0, \quad e \in E(G) - E_0(G),$$

$$(4.2) \quad \sum_{C \in \mathcal{C}(G): e \in C} z(C) - \sum_{\tau=(e,e') \in \mathcal{T}(G)} g(\tau) + \sum_{\tau=(e',e) \in \mathcal{T}(G)} g(\tau) + d(e) = 0, \quad e \in E_0(G).$$

Applying Farkas lemma to (2.2)–(2.4), we obtain an l.p. solvability criterion.

Statement 4.1 *Let $\sigma \in \mathbb{R}^{E_0(G)}$. A concave cocirculation h with border σ exists if and only if*

$$(4.3) \quad \sigma \cdot d \geq 0$$

holds for any $z : \mathcal{C}(G) \rightarrow \mathbb{R}$, $g : \mathcal{T}(G) \rightarrow \mathbb{R}_+$ and $d : E_0(G) \rightarrow \mathbb{R}$ satisfying (4.1) and (4.2). ■

Hereinafter for $a, b \in \mathbb{R}^E$, $a \cdot b$ denotes the inner product $\sum(a(e)b(e) : e \in E)$. We call a triple $K = (z, g, d)$ satisfying (4.1)–(4.2) a *vector configuration*, or, briefly, a *v-configuration*, and regard d as its border.

Statement 4.1 implies that the cone \mathcal{D} of borders of v-configurations (which is convex) is anti-polar to the cone $\mathcal{B}(G)$ of borders of concave cocirculation in G , i.e., $\mathcal{D} := \{d \in \mathbb{R}^{E_0(G)} : \sigma \cdot d \geq 0 \ \forall \sigma \in \mathcal{B}(G)\}$. For a boundary edge e , define $\theta(e) := 1$ if $e \in E_0^+(G)$, and -1 if $e \in E_0^-(G)$. Since the dimension of $\mathcal{B}(G)$ is $|E_0(G)| - 1$ (by Statement 2.1) and $\mathcal{B}(G)$ is contained in the hyperplane θ^\perp orthogonal to θ (by (2.5)), the cone \mathcal{D} is full-dimensional and contains the line $\mathbb{R}\theta$. So the facets of $\mathcal{B}(G)$ one-to-one correspond (by the orthogonality) to the 2-dimensional faces of \mathcal{D} , each being of the form $r_1d + r_2\theta$ ($r_1 \in \mathbb{R}_+, r_2 \in \mathbb{R}$) for a certain $d \in \mathbb{R}^{E_0(G)}$.

For a function (vector) x , let $\text{supp}^+(x)$ and $\text{supp}^-(x)$ denote the *positive* part $\{e : x(e) > 0\}$ and the *negative* part $\{e : x(e) < 0\}$ of the *support* $\text{supp}(x)$ of x , respectively. Since inequality (4.3) is invariant under adding to d any multiple of θ , it suffices to verify this inequality only for the v-configurations $K = (z, g, d)$ satisfying:

$$(4.4) \quad \text{(a) } \text{supp}^+(d) \subseteq E_0^+(G) \text{ and } \text{supp}^-(d) \subseteq E_0^-(G), \text{ and (b) } \text{supp}(d) \neq \emptyset, E_0(G).$$

In what follows, we throughout assume that any v-configuration in question satisfies (a). When (b) takes place too, we call K *proper*.

Let $\Sigma(G)$ be the set of $\sigma \in \mathbb{R}^{E_0(G)}$ satisfying (2.5)–(2.6). Then $\mathcal{B}(G) \subseteq \Sigma(G)$. A v-configuration $K = (z, g, d)$ is called *essential* if d separates $\Sigma(G)$, i.e., $\sigma \cdot d < 0$ for some $\sigma \in \Sigma(G)$.

Consider two v-configurations $K = (z, g, d)$ and $K' = (z', g', d')$. K and K' are called *equivalent* if their borders are proportional, i.e., $d = rd'$ for some $r > 0$. We say that K' *dominates* K if at least one of the following takes place:

(4.5) (i) $\sigma \in \Sigma(G)$ and $\sigma \cdot d < 0$ imply $\sigma \cdot d' < 0$, and there exists $\sigma \in \Sigma(G)$ such that $\sigma \cdot d \geq 0$ but $\sigma \cdot d' < 0$; or

(ii) K is proper and not equivalent to K' , and $K - rK'$ is a v-configuration (subject to (a) in (4.4)) for some $r > 0$.

If K is dominated by some K' , then K is redundant and can be excluded from consideration (as d cannot be facet-determining for $\mathcal{B}(G)$). This is obvious in case (i). And in case (ii), the border $d'' := d - rd'$ of the v-configuration $K'' := K - rK'$ is nonzero and satisfies $(d'')^\perp \cap \mathcal{B}(G) \supseteq d^\perp \cap \mathcal{B}(G)$ and $\text{supp}(d'') \subseteq \text{supp}(d)$. The former inclusion implies that if d^\perp contains a facet F of $\mathcal{B}(G)$, then $(d'')^\perp$ contains F as well. In this case we have $d'' = r_1 d + r_2 \theta$ for some $r_1 > 0$ and $r_2 \in \mathbb{R}$, which contradicts the latter inclusion since $\text{supp}(d) \neq E_0(G)$ and K, K' are not equivalent.

Our method of proof of Theorem 3.1 will consist in examining an arbitrary essential configuration K and attempting to show that K is dominated unless it is equivalent to some “puzzle configuration”. Note that one can consider only rational-valued z, g, d in (4.1)–(4.2). Moreover, by scaling, it suffices to deal with *integer* v-configurations (z, g, d) .

For a boundary edge e of G , let χ^e denote the unit base vector of e in $\mathbb{R}^{E_0(G)}$ (i.e., $\chi^e(a) = 1$ for $a = e$, and 0 otherwise). We will use the following observation:

(4.6) if K is an essential v-configuration with border d , K' is a v-configuration with border d' , and $d' = d - \chi^e + \chi^{e'}$, where e, e' are boundary edges occurring in the same side-path in this order, then K' dominates K .

To see this, let $d'' := \chi^e - \chi^{e'}$. Then $\sigma \cdot d'' \geq 0$ for all $\sigma \in \Sigma(G)$, by (2.6). This and $d = d' + d''$ imply $\sigma \cdot d \geq 0$ for all $\sigma \in \Sigma(G)$ satisfying $\sigma \cdot d' \geq 0$. Take $\sigma_1 \in \mathcal{B}(G)$ such that $\sigma_1(e) > \sigma_1(e')$ (existing by Statement 2.1). Then $\sigma_1 \cdot d' \geq 0$ and $\sigma_1 \cdot d'' > 0$, implying $p := \sigma_1 \cdot d > 0$. Take $\sigma_2 \in \Sigma(G)$ such that $q := \sigma_2 \cdot d < 0$ (existing as K is essential). Define $\sigma := \sigma_2 - \frac{q}{p}\sigma_1$. We have $\sigma \cdot d = \sigma_2 \cdot d - \frac{q}{p}\sigma_1 \cdot d = q - q = 0$ and $\sigma \cdot d' = \sigma \cdot d - \sigma \cdot d'' = -\sigma \cdot d'' = -\sigma_2 \cdot d'' + \frac{q}{p}\sigma_1 \cdot d'' < 0$, yielding (4.5)(i).

5 Flow Model

In the proof of Theorem 3.1 we will take advantage of a representation of a v-configuration $K = (z, g, d)$ in a more combinatorial form. It is described in this section.

For a 3-circuit C , let us interpret $z(C)$ as the *weight* of the little triangle Δ_C surrounded by C . Similarly, $d(e)$ is the weight of a boundary edge e . For each tandem τ , set $g(a_\tau) := g(\tau)$, interpreting it as the *flow* on the edge a_τ of the graph H (introduced in Section 3). The boundary edges and little triangles with nonzero weights are interpreted as “sources” or “sinks” of the flow. We say that a boundary edge e *emits* $d(e)$ (units of) flow if $d(e) > 0$, and *absorbs* $|d(e)|$ flow if $d(e) < 0$. Similarly, a little triangle Δ_C emits $z(C)$ flow (through each of its three sides) if

$z(C) > 0$, and absorbs $|z(C)|$ flow if $z(C) < 0$. Then relations (4.1)–(4.2) turn into the flow balance condition

$$(5.1) \quad \text{div}_g(v) + \sum_{C \in \mathcal{C}(G): v \in \Delta_C} z(C) + \sum_{e \in E_0(G): v \in e} d(e) = 0 \quad \text{for each } v \in V(H),$$

where

$$\text{div}_g(v) := \sum_{u:(u,v) \in E(H)} g(u,v) - \sum_{w:(v,w) \in E(H)} g(v,w).$$

Next, for a path P in H , let $\chi^P \in \mathbb{R}^{E(H)}$ denote the incidence vector of the set of edges of P . Considering g as a function on $E(H)$, applying usual flow decomposition techniques and taking into account (5.1), one can find paths P_1, \dots, P_k in H (possibly including degenerate paths) and positive real *weights* $\alpha_1, \dots, \alpha_k$ of these paths such that:

$$(5.2) \quad g = \alpha_1 \chi^{P_1} + \dots + \alpha_k \chi^{P_k};$$

(5.3) for each edge e of G , the sum of weights of emitting elements containing e is equal to the sum of weights of paths P_i leaving e ; similarly, the sum of absolute values of weights of absorbing elements containing e is equal to the sum of weights of paths P_i entering e .

We call $(P_1, \dots, P_k; \alpha_1, \dots, \alpha_k)$ satisfying (5.2)–(5.3) a *paths decomposition* of g .

When g is integer-valued, there is a decomposition with all weights α_i integer (an *integer paths decomposition*). In this case we define a triple $\mathcal{K} = (\Phi, \mathcal{P}, \iota)$ representing K , in a sense, as follows. Take $d(e)$ copies of each emitting boundary edge e and $z(C)$ copies of each emitting triangle Δ_C , forming family Φ^+ of (unweighted) emitting elements. Take $|d(e)|$ copies of each absorbing boundary edge e and $|z(C)|$ copies of each absorbing triangle Δ_C , forming family Φ^- of absorbing elements. Then Φ is the disjoint union of Φ^+ and Φ^- . Take α_i copies of each path P_i , forming \mathcal{P} . Assign a map $\iota : \mathcal{P} \rightarrow \Phi^+ \times \Phi^-$ so as to satisfy the following property:

(5.4) if $P \in \mathcal{P}$ and $\iota(P) = (\phi, \phi')$, then P leaves ϕ and enters ϕ' ; moreover, for each $\phi \in \Phi^+$ (resp. $\phi \in \Phi^-$) and each edge e in ϕ , there is exactly one path $P \in \mathcal{P}$ such that $\iota(P) = (\phi, \cdot)$ and P leaves e (resp. $\iota(P) = (\cdot, \phi)$ and P enters e).

The existence of such an ι follows from (5.3). When $\iota(P) = (\phi, \phi')$, we say that the path P is *attached* to the elements ϕ and ϕ' . So each triangle in Φ has three attached paths, by one from each of H_1, H_2, H_3 , and each boundary edge in Φ has one attached path.

A converse construction also takes place. More precisely, consider families $\Phi^+, \Phi^-, \mathcal{P}$ consisting of copies of some little triangles and edges from $E_0^+(G)$, of copies of little triangles and edges from $E_0^-(G)$, and of copies of paths in H , respectively. Let Φ be the disjoint union of Φ^+ and Φ^- , and ι a map of \mathcal{P} to $\Phi^+ \times \Phi^-$ satisfying (5.4). We refer to $\mathcal{K} = (\Phi, \mathcal{P}, \iota)$ as a *combinatorial configuration*, or, briefly, a *c-configuration*. Emphasize that we admit some little triangles of G (but not boundary edges) to have copies simultaneously in both Φ^+ and Φ^- . Now

(5.5) for $C \in \mathcal{C}(G)$, define $z(C)$ to be the number of copies of the triangle Δ_C in Φ^+ minus the number of copies of Δ_C in Φ^- ; for $e \in E_0^+(G)$, define $d(e)$ to be the number of copies of e in Φ^+ ; for $e \in E_0^-(G)$, define $d(e)$ to be minus the number of copies of e in Φ^- ; and define $g := \sum \{\chi^P : P \in \mathcal{P}\}$.

Then z, g, d give a v-configuration, denoted by $K(\mathcal{K})$. We formally define border $d(\mathcal{K})$ of \mathcal{K} to be the border of $K(\mathcal{K})$. Also we apply to \mathcal{K} adjectives “proper, essential” if $K(\mathcal{K})$ is such, and similarly for the property of being “equivalent to” or “dominated by” another configuration.

When no little triangle of G has copies simultaneously in both Φ^+, Φ^- , we say that \mathcal{K} is *homogeneous*. In particular, any c-configuration \mathcal{K} obtained from a v-configuration K by the first construction is homogeneous; in this case $K(\mathcal{K}) = K$.

6 Proof of the Theorem

The proof of Theorem 3.1 for a convex grid G falls into three lemmas. By reasonings in Sections 4 and 5, we can deal with combinatorial configurations and, moreover, with those of them that are proper, essential and homogeneous.

Given a c-configuration $\mathcal{K} = (\Phi, \mathcal{P}, \iota)$, we say that a little triangle or a boundary edge of G or a path of H is in \mathcal{K} if at least one copy of this element occurs in \mathcal{K} . Adding to (deleting from) \mathcal{K} such an element means adding (deleting) exactly one copy of it.

We associate with \mathcal{K} undirected (multi)graph $\Gamma_{\mathcal{K}}$ whose vertices are the elements of Φ and whose edges one-to-one correspond to the paths in \mathcal{P} : each path $P \in \mathcal{P}$ generates an edge connecting ϕ and ϕ' when $\iota(P) = (\phi, \phi')$ (it is analogous to the graph Γ_{Π} associated with a puzzle Π , defined in the end of Section 3). The (disjoint) union of \mathcal{K} with another or the same c-configuration \mathcal{K}' is defined in a natural way and denoted by $\mathcal{K} + \mathcal{K}'$ (its associated graph $\Gamma_{\mathcal{K} + \mathcal{K}'}$ is the disjoint union of $\Gamma_{\mathcal{K}}$ and $\Gamma_{\mathcal{K}'}$).

If the interiors of distinct little triangles or edges ϕ, ϕ', ϕ'' of G are intersected by a line of H in this order, we say that ϕ' lies *between* ϕ and ϕ'' .

We call \mathcal{K} *oriented* if all triangles in Φ^- (the absorbing triangles) are normal and all triangles in Φ^+ (the emitting triangles) are turned-over. The first lemma eliminates the non-oriented configurations.

Lemma 6.1 *Let a c-configuration $\mathcal{K} = (\Phi, \mathcal{P}, \iota)$ be proper, essential and homogeneous. There exists a c-configuration \mathcal{K}' such that either \mathcal{K}' dominates \mathcal{K} , or \mathcal{K}' is equivalent to \mathcal{K} and is oriented.*

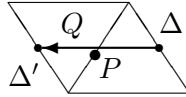
Proof. Since we can consider any homogeneous c-configuration equivalent to \mathcal{K} , one may assume that, among such configurations, \mathcal{K} is chosen so that

(6.1) the number $\eta(\mathcal{K}) := |\Phi| + |\mathcal{P}|$ is as small as possible.

Let us say that a triangle in Φ is *good* if it is either emitting and turned-over, or absorbing and normal. If all triangles are good, \mathcal{K} is already oriented. So assume \mathcal{K} contains at least one bad triangle. Our aim is to show that \mathcal{K} is dominated.

First of all we impose an additional condition on \mathcal{K} . Suppose there is a degenerate path $P \in \mathcal{P}$ attached to a pair of bad triangles $\Delta \in \Phi^+$ and $\Delta' \in \Phi^-$; so Δ, Δ' share an edge e , and P is of the form v_e . Let e be parallel to ξ_i and let a, a' be the edges of Δ, Δ' , respectively, parallel to ξ_{i-1} . Observe that H_{i-1} has path Q (with one edge) leaving v_a and entering $v_{a'}$. When \mathcal{P} contains

a copy of Q attached to the pair (Δ, Δ') as well, we call this pair *dense*. See the picture where $i = 3$.



We assume that, among all homogeneous c-configurations having the same border $d(\mathcal{K})$ and satisfying (6.1), \mathcal{K} is chosen so that

(6.2) the number $\omega(\mathcal{K})$ of dense pairs in \mathcal{K} is maximum.

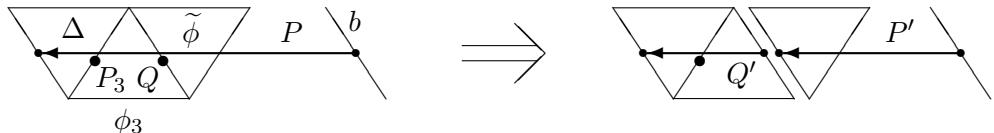
Suppose the graph $\Gamma_{\mathcal{K}}$ associated with \mathcal{K} is not connected. Then \mathcal{K} is the union of two nonempty c-configurations $\mathcal{K}', \mathcal{K}''$, and we have $d(\mathcal{K}) = d(\mathcal{K}') + d(\mathcal{K}'')$ and $\eta(\mathcal{K}) = \eta(\mathcal{K}') + \eta(\mathcal{K}'')$. (6.1) implies that $d(\mathcal{K}') \neq 0$ and \mathcal{K}' is not equivalent to \mathcal{K} . Hence \mathcal{K}' dominates \mathcal{K} , by (4.5)(ii). So one may assume that $\Gamma_{\mathcal{K}}$ is connected. Then each $\phi \in \Phi$ is reachable in $\Gamma_{\mathcal{K}}$ by a path from a vertex representing a boundary edge; let $\rho(\phi)$ denote the minimum number of edges of such a path.

We consider a bad triangle Δ with $\rho(\Delta) =: \bar{\rho}$ minimum and proceed by induction on $\bar{\rho}$. Let $P \in \mathcal{P}$ be a path attached to Δ and to an element $\phi \in \Phi$ with $\rho(\phi) = \bar{\rho} - 1$. Consider two cases.

Case 1. Let $\bar{\rho} = 1$. Then ϕ is (a copy of) a boundary edge b . Assume $b \in E_0^+(G)$; the case $b \in E_0^-(G)$ is symmetric. Then Δ is absorbing and turned-over, and P leaves b and enters Δ . Let for definiteness b be parallel to ξ_2 . For $i = 1, 2, 3$, consider $P_i \in \mathcal{P}$ and $\phi_i \in \Phi^+$ such that P_i is in H_i and $\iota(P_i) = (\phi_i, \Delta)$. Let e_i be the edge of Δ parallel to ξ_i . (So $P_2 = P$ and $\phi_2 = b$.)

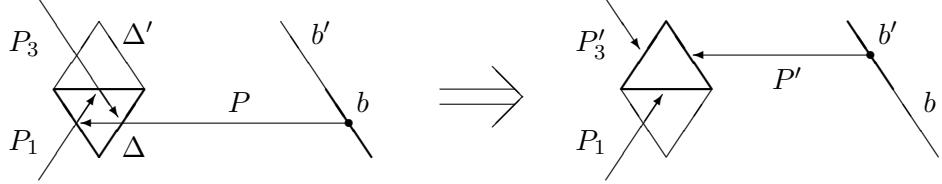
Suppose P_3 is degenerate, i.e., $P_3 = v_{e_3}$. Then ϕ_3 is a normal emitting triangle, and therefore, ϕ_3 is bad. Take path Q in H_2 attached to ϕ_3 , and let $\iota(Q) = (\phi_3, \tilde{\phi})$. The fact that P_3 is degenerate implies that b, ϕ_3, Δ are intersected by a line of H_2 in this order. Hence H_2 has path P' leaving b and entering ϕ and path Q' leaving ϕ_3 and entering Δ . Replace in \mathcal{P} the paths P, Q by P', Q' , making P' attached to $b, \tilde{\phi}$ and making Q' attached to ϕ_3, Δ .

This results in a correct c-configuration \mathcal{K}' with $\eta(\mathcal{K}') = \eta(\mathcal{K})$ in which (ϕ_3, Δ) becomes a dense pair. One can see that if Q is nondegenerate, then such a transformation does not destroy any dense pair of the previous configuration; so $\omega(\mathcal{K}') > \omega(\mathcal{K})$, contradicting (6.2). And if Q is degenerate, then ϕ_3 and $\tilde{\phi}$ share an edge of H_2 , whence $\tilde{\phi}$ is a turned-over absorbing triangle forming a pair of bad triangles with ϕ_3 . The only possible dense pair which could be destroyed by the transformation is just $(\phi_3, \tilde{\phi})$ (when this pair is also connected in \mathcal{K} by the corresponding path in H_1). In this case we have $\omega(\mathcal{K}') \geq \omega(\mathcal{K})$, so the above replacement maintains (6.2). Moreover, the new path leaving b (namely, P') enters a bad triangle (namely, $\tilde{\phi}$) as before and is shorter than P , as illustrated in the picture:



Doing so, we eventually obtain a c-configuration where b is connected with a bad triangle whose attached path in H_3 is nondegenerate.

Thus, we may assume that P_3 is nondegenerate. Then, by the convexity of G , the edge e_1 of Δ does not lie on the boundary of G , and b cannot be the last edge of the side-path B_2^+ . We now transform \mathcal{K} as follows. Let Δ' be the normal triangle of G containing e_1 , and b' the edge of B_2^+ next to b . Then H_2 has path P' leaving b' and entering Δ' and H_3 has path P'_3 leaving ϕ_3 and entering Δ' (P'_3 is a part of the nondegenerate P_3). We replace in \mathcal{K} the edge b by (one emitting copy of) b' , the triangle Δ by (one absorbing copy of) Δ' , and the paths P, P_3 by P', P'_3 , making P' attached to b', Δ' , and making P'_3 attached to ϕ_3, Δ' (while P_1 becomes attached to Δ' instead of Δ) :



This results in a (not necessarily homogeneous) c-configuration \mathcal{K}' with the border $d(\mathcal{K}) - \chi^b + \chi^{b'}$. By (4.6), \mathcal{K} is dominated by \mathcal{K}' .

Case 2. Let $\bar{\rho} > 1$. Assume the bad triangle Δ in question is absorbing (and turned-over); the case of emitting Δ is symmetric. Let for definiteness P be in H_2 , and define P_i, ϕ_i, e_i ($i = 1, 2, 3$) as in Case 1. (So $P = P_2$ and $\phi = \phi_2$.) Since $\rho(\phi) = \bar{\rho} - 1 \geq 1$, ϕ is a good triangle. So ϕ is a turned-over emitting triangle and P is nondegenerate. Arguing as in Case 1, we can impose the condition that P_3 is nondegenerate. This and the convexity of G imply that neither the edge e_1 of Δ nor the edge q of ϕ parallel to ξ_1 lies on the boundary of G . Let Δ' be the normal little triangle of G containing e_1 , and ϕ' the normal triangle containing q . We replace Δ, ϕ in Φ by Δ', ϕ' .

More precisely, when Δ is replaced by Δ' , we accordingly replace the paths P, P_3 attached to Δ by paths P', P'_3 (while P_1 preserves, becoming attached to Δ'). Here P' is the path of H_2 leaving ϕ' and entering Δ' , and P'_3 is the path of H_3 leaving ϕ_3 and entering Δ' (as before, P'_3 is a part of the nondegenerate path P_3). And when replacing ϕ by ϕ' , we should also replace path \tilde{Q} of H_3 attached to ϕ , entering triangle $\tilde{\phi} \in \Phi^-$ say, by path \tilde{Q}' of H_3 leaving ϕ' and entering $\tilde{\phi}$. (\tilde{Q}' exists since ϕ lies between ϕ' and $\tilde{\phi}$.) The path of H_1 attached to ϕ becomes attached to ϕ' . This gives a c-configuration \mathcal{K}' in which the added triangle ϕ' is bad and its rank $\rho(\phi')$ is equal to $\bar{\rho} - 1$.

We have $d(\mathcal{K}') = d(\mathcal{K})$ and $\eta(\mathcal{K}') = \eta(\mathcal{K})$. The latter implies that \mathcal{K}' is homogeneous, i.e., \mathcal{K} has no emitting copy of Δ' or ϕ' . For otherwise, cancelling in \mathcal{K}' one emitting copy and one absorbing copy of the same little triangle of G and properly concatenating their attached paths, we would obtain a configuration with a smaller value of η , contrary to (6.1). Finally, one can see that neither Δ nor ϕ can be involved in dense pairs of \mathcal{K} . Hence no dense pair is destroyed while constructing \mathcal{K}' , implying $\omega(\mathcal{K}') \geq \omega(\mathcal{K})$. Now the result follows by induction on $\bar{\rho}$. ■

Thus, it suffices to consider only oriented configurations.

A puzzle $\Pi = (\mathcal{F}, \mathcal{P})$ generates an oriented c-configuration $(\Phi, \mathcal{P}, \iota)$ in a natural way: Φ^+ is the set of turned-over triangles in \mathcal{F} and edges in $b^+(\Pi)$, Φ^- is the set of normal triangles in \mathcal{F}

and edges in $b^-(\Pi)$, and for each $u - v$ path $P \in \mathcal{P}$, $\iota(P)$ is the pair $(\phi \in \Phi^+, \phi' \in \Phi^-)$ such that the point u is contained in ϕ and the point v is contained in ϕ' . Such a *puzzle c-configuration* is denoted by \mathcal{K}_Π .

The next lemma describes a situation when an oriented configuration \mathcal{K} can be split into a puzzle configuration and another one (and therefore, \mathcal{K} is redundant). Let us say that paths P, P' of H are *crossing* if they are not parallel and their interiors have a point in common, and that P and a little triangle Δ of G are *overlapping* if P meets the interior of Δ :



One can see that the puzzle configurations are precisely those having neither crossing nor overlapping pairs. Given an oriented c-configuration $\mathcal{K} = (\Phi, \mathcal{P}, \iota)$, define its *minimal pre-configuration* $\mathcal{K}^{\min} = (\Psi, \mathcal{P}^{\min}, \hat{\iota})$ as follows. Let Ψ^+ (resp. Ψ^-) be the set of little triangles and boundary edges of G having at least one copy in Φ^+ (resp. Φ^-). Then $\Psi := \Psi^+ \cup \Psi^-$. The set \mathcal{P}^{\min} is formed by taking, for each edge $e \in E(G)$ contained in a member of Ψ^+ , one (inclusion-wise) minimal path in \mathcal{P} with the beginning v_e , taking for each edge $e \in E(G)$ contained in a member of Ψ^- , one minimal path in \mathcal{P} with the end v_e , and ignoring repeated paths if arise. Define $\hat{\iota}$ to be the map attaching a $u - v$ path $P \in \mathcal{P}^{\min}$ to the pair $(\phi \in \Psi^+, \phi' \in \Psi^-)$ such that $u \in \phi$ and $v \in \phi'$ (this pair is unique since K is oriented). Note that \mathcal{K}^{\min} need not be a c-configuration since some triangles (boundary edges) in it may have more than three (resp. one) attached paths.

Lemma 6.2 *Let a c-configuration $\mathcal{K} = (\Phi, \mathcal{P}, \iota)$ be proper and oriented, and let $\mathcal{K}^{\min} = (\Psi, \mathcal{P}^{\min}, \hat{\iota})$ be its minimal pre-configuration. Suppose \mathcal{K}^{\min} contains neither crossing paths nor overlapping a path and a triangle. Then: (a) \mathcal{K}^{\min} is a puzzle c-configuration, and (b) \mathcal{K}^{\min} either is equivalent to \mathcal{K} or dominates \mathcal{K} .*

Proof. From the non-existence of paths in \mathcal{P}^{\min} overlapping triangles in Ψ it easily follows that for each element $\phi \in \Psi^+$ and each edge e in ϕ , there is exactly one path $P \in \mathcal{P}^{\min}$ leaving e , and similarly for each element $\phi' \in \Psi^-$ and each edge e' in ϕ' , there is exactly one path $P' \in \mathcal{P}^{\min}$ entering e' . Hence \mathcal{K}^{\min} is a c-configuration, and now the absence of crossing paths in \mathcal{K} implies that \mathcal{K}^{\min} is a puzzle configuration, yielding (a). Next, one can rearrange the attaching map ι in \mathcal{K} so that \mathcal{K} be represented as the union of \mathcal{K}^{\min} and some c-configuration \mathcal{K}'' . This implies (b), by (4.5)(ii). ■

For $i = 1, 2, 3$, a sequence (ϕ_1, \dots, ϕ_k) of distinct little triangles or edges of G is called an *i-chain* if their interiors are intersected in this order by a path of H_i . If (Δ, Δ') is an *i-chain* of two normal little triangles and there is no normal triangle between them, we say that Δ is the *i-predecessor* of Δ' , and similarly for turned-over triangles.

Our final lemma is the following.

Lemma 6.3 *Let a c-configuration $\mathcal{K} = (\Phi, \mathcal{P}, \iota)$ be proper, essential and oriented. If \mathcal{K} is not equivalent to a puzzle c-configuration, then \mathcal{K} is dominated.*

Proof. Since we can replace \mathcal{K} by any oriented c -configuration equivalent to \mathcal{K} (e.g., by taking the union of r copies of \mathcal{K} for any r), one may assume that, among such configurations, \mathcal{K} is chosen so that:

- (6.3) (i) there are sufficiently many copies of each member of $\Phi \cup \mathcal{P}$;
- (ii) subject to (i), the number $t(\mathcal{K})$ of little triangles of G having copies in Φ is maximum;
- (iii) subject to (i),(ii), the number $p(\mathcal{K})$ of paths of H having copies in \mathcal{K} is maximum.

From (iii) it follows that

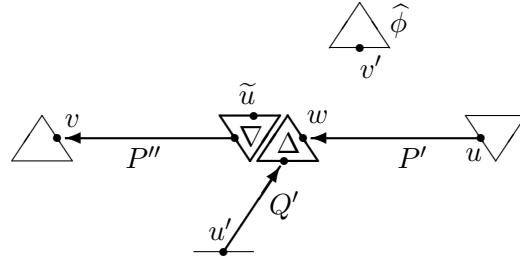
- (6.4) for any (not necessarily distinct) vertices u_1, u_2, u_3, u_4 occurring in a path of H in this order, if \mathcal{P} contains copies of both $u_1 - u_3$ path P and $u_2 - u_4$ path P' , then \mathcal{P} contains copies of both $u_1 - u_4$ path Q and $u_2 - u_3$ path Q' as well, and vice versa.

Indeed, if at least one of Q, Q' is not in \mathcal{P} , we can add Q, Q' to \mathcal{P} and delete P, P' from \mathcal{P} , accordingly correcting the map ι . This increases $p(\mathcal{K})$. (Recall that adding to \mathcal{K} a triangle or a boundary edge of G or a path of H means adding *one* copy of this element, and similarly for deleting an element.) The reverse assertion is proved similarly.

Also we assume that the minimal pre-configuration \mathcal{K}^{\min} contains crossing paths or overlapping a path and a triangle; otherwise the result immediately follows from Lemma 6.2. We show that \mathcal{K} is dominated in both cases.

Case 1. Let \mathcal{K}^{\min} contain crossing a $u - v$ path P and a $u' - v'$ path Q . Assume for definiteness that P is in H_2 and is minimal among the paths of \mathcal{P} beginning at u , and that Q is in H_1 (P is anti-parallel to ξ_1 and Q is anti-parallel to ξ_3); the case when P is minimal among the paths ending at v is symmetric. Observe that the point w where P, Q meet is a vertex of H_2 . Let Δ be the normal little triangle whose edge parallel to ξ_2 contains w as the middle point. Then Δ is not in Φ . For otherwise \mathcal{P} would contain a path from some vertex w' to w (as Δ is absorbing). Applying (6.4) to w', u, w, v or to u, w', w, v , we obtain that \mathcal{P} contains the $u - w$ path, contradicting the minimality of P .

Next we proceed as follows. For $i = 1, 2, 3$, let e_i be the edge of Δ parallel to ξ_i . (So $w = v_{e_2}$.) Take the turned-over little triangle ∇ containing e_3 . Let e'_1, e'_2 be the edges of ∇ parallel to ξ_1, ξ_2 , respectively. Then H_2 has $u - w$ path P' and $v_{e'_2} - v$ path P'' , and H_1 has $u' - v_{e_1}$ path Q' :



Add (one copy of) the triangle Δ to Φ^- , the triangle ∇ to Φ^+ , and the paths P', P'', Q' together with the degenerate path v_{e_3} (“connecting” Δ and ∇ in H_3) to \mathcal{P} . Accordingly delete P, Q from

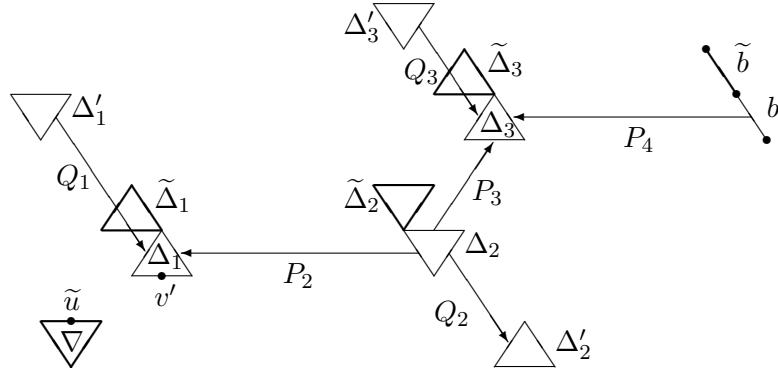
\mathcal{P} . The attachments for the added elements are assigned in a natural way (e.g., $\iota(P') := (\phi, \Delta)$, where ϕ is the element of the old Φ^+ to which P was attached). This increases the parameter t (since Δ is added while the new \mathcal{K} contains a copy of each triangles from the previous \mathcal{K} , by assumption (6.3)(i)). However, \mathcal{K} becomes an “incomplete” configuration since ∇ has no attached path in H_1 , and similarly for element $\tilde{\phi}$ of Φ^- to which Q was attached. We cannot improve \mathcal{K} straightforwardly because the points $\tilde{u} := v_{e'_1}$ and v' do not lie on one line of H_1 .

Our aim is to improve the new \mathcal{K} , without decreasing the current value of t , in order to obtain a correct c-configuration \mathcal{K}' either dominating or equivalent to the initial \mathcal{K} . This will yield the result in the former case and lead to a contradiction with assumption (6.3)(ii) in the latter case.

First of all we iteratively construct a sequence S of alternating members of Φ and \mathcal{P} as follows. Start with $\Delta_1 := \tilde{\phi}$. Let $\Delta_i \in \Phi$ be the last element of the current S . If Δ_i is a boundary edge, halt. Otherwise add P_{i+1}, Δ_{i+1} to S , where P_{i+1} is attached to Δ_i, Δ_{i+1} . More precisely: (a) if i is odd (and Δ_i is a normal triangle), then P_{i+1} is a path of H_2 and $\iota(P_{i+1}) = (\Delta_{i+1}, \Delta_i)$, and (b) if i is even (and Δ_i is a turned-over triangle), then P_{i+1} is a path of H_1 and $\iota(P_{i+1}) = (\Delta_i, \Delta_{i+1})$. Let Δ_{q+1} be the last element of the final S . Clearly the edge $b := \Delta_{q+1}$ belongs to B_2^+ when q is odd, and to B_1^- when q is even.

Assume q is odd; the case of q even is examined analogously. For $i = 1, \dots, q$, let $Q_i \in \mathcal{P}$ be the path of H_3 attached to Δ_i (it enters Δ_i for i odd, and leaves Δ_i for i even). Let Δ'_i be the other element of Φ to which Q_i is attached. We say that the triangle Δ_i is *squeezed* if i is odd and Q_i is degenerate.

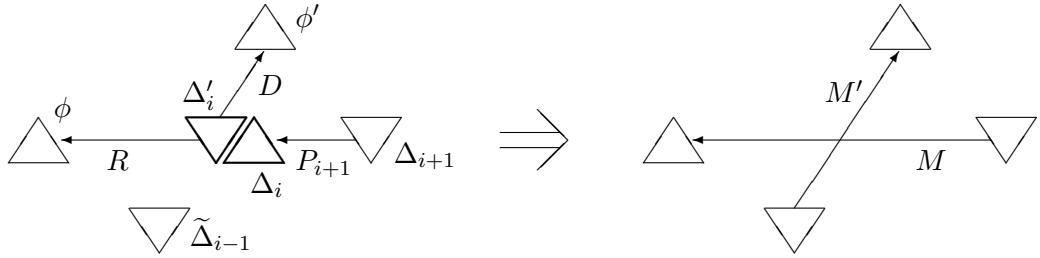
We first explain how to transform \mathcal{K} into the desired correct c-configuration when no Δ_i is squeezed. By the convexity of G (and regardless of the squeezedness of any Δ_i), the line in the plane parallel to ξ_3 and passing the point \tilde{u} in ∇ separates S from B_3^+ (letting B_3^+ be the common vertex of B_2^- and B_1^- when they meet). This implies that S can be shifted by distance 1 in the direction of ξ_2 (approaching B_3^+). More precisely, each triangle Δ_i has 3-predecessor $\tilde{\Delta}_i$ in G , and B_2^+ contains edge \tilde{b} next to b . See the picture where $q = 3$:



These triangles $\tilde{\Delta}_i$ and the elements $\tilde{\Delta}_0 := \nabla$ and $\tilde{\Delta}_{q+1} := \tilde{b}$ are connected in H by paths P'_1, \dots, P'_{q+1} in a natural way: P'_i is the path of H_1 leaving $\tilde{\Delta}_i$ and entering $\tilde{\Delta}_{i+1}$ when i is odd, and the path of H_2 leaving $\tilde{\Delta}_{i+1}$ and entering $\tilde{\Delta}_i$ when i is even. Also there are paths Q'_1, \dots, Q'_q of H_3 such that Q'_i leaves Δ'_i and enters $\tilde{\Delta}_i$ when i is odd (as Δ_i is not squeezed, and therefore, $\tilde{\Delta}_i$ lies between Δ'_i and Δ_i), and Q'_i leaves $\tilde{\Delta}_i$ and enters Δ'_i when i is even.

Add to \mathcal{K} the triangles $\tilde{\Delta}_1, \dots, \tilde{\Delta}_q$, the paths $P'_1, \dots, P'_{q+1}, Q'_1, \dots, Q'_q$ and the edge \tilde{b} , making P'_i attached to $\tilde{\Delta}_{i-1}, \tilde{\Delta}_i$, and making Q'_j attached to $\tilde{\Delta}_j, \Delta'_j$. Accordingly delete from \mathcal{K} the triangles $\Delta_1, \dots, \Delta_q$, the paths $P_2, \dots, P_{q+1}, Q_1, \dots, Q_q$ and the boundary edge b . This results in a correct c-configuration \mathcal{K}' . Moreover, \mathcal{K}' has the border $d(\mathcal{K}) - \chi^b + \chi^{\tilde{b}}$. Therefore, \mathcal{K}' dominates the initial \mathcal{K} , by (4.6).

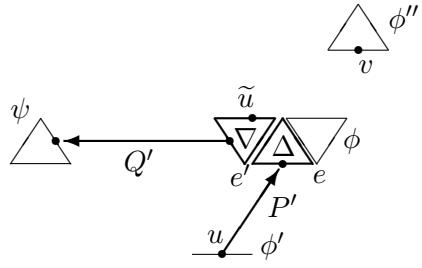
Next suppose there is a squeezed Δ_i (i is odd); let i be minimum among such triangles. Form the triangles $\tilde{\Delta}_0, \dots, \tilde{\Delta}_{i-1}$ and paths $P'_1, \dots, P'_i, Q'_1, \dots, Q'_{i-1}$ as above. Take paths $R, D \in \mathcal{P}$ attached to Δ'_i and belonging to H_2 and H_1 , respectively. Let ϕ, ϕ' be the other (normal) triangles to which R, D are attached, respectively. Since Δ_i is squeezed, $(\Delta_{i+1}, \Delta_i, \Delta'_i, \phi)$ is a 2-chain and $(\tilde{\Delta}_{i-1}, \Delta'_i, \phi')$ is a 1-chain. See the picture:



Let M be the path of H_2 leaving Δ_{i+1} and entering ϕ , and M' the path of H_1 leaving $\tilde{\Delta}_{i-1}$ and entering ϕ' . We add to \mathcal{K} the triangles $\tilde{\Delta}_1, \dots, \tilde{\Delta}_{i-1}$ and the paths $P'_1, \dots, P'_i, Q'_1, \dots, Q'_{i-1}, M, M'$ and accordingly delete the triangles $\Delta_1, \dots, \Delta_i$ and Δ'_i and the paths $P_2, \dots, P_{i+1}, Q_1, \dots, Q_i, R, D$. (Note that if P_i is degenerate, then $\tilde{\Delta}_{i-1}$ and Δ'_i are copies of the same triangle of G ; we consider them as different objects one of which is added and the other is deleted.) The resulting \mathcal{K}' is a correct c-configuration with the same border $d(\mathcal{K})$. But $t(\mathcal{K}') > t(\mathcal{K})$ (as Δ was added, while deleting the above triangles does not affect t , by (6.3)(i)). This contradicts (6.3)(ii).

Case 2. Let \mathcal{K}^{\min} contain overlapping a path P and a triangle ϕ . One may assume that P is a $u - v$ path of H_1 and that P is minimal among the paths in \mathcal{P} beginning at u . Let $\iota(P) = (\phi', \phi'')$. Then ϕ lies between ϕ' and ϕ'' . Notice that there is no normal (absorbing) triangle $\tilde{\phi} \in \Phi$ between ϕ' and ϕ'' . For if such a $\tilde{\phi}$ exists, then the end vertex w of the path of H_1 attached to $\tilde{\phi}$ is an intermediate vertex of P . But then \mathcal{P} contains the $u - w$ path (by (6.4)), contrary to the minimality of P . So ϕ is a turned-over (emitting) triangle.

Take path Q of H_2 attached to ϕ ; let $\iota(Q) = (\phi, \psi)$. Since the absorbing element ψ cannot be a normal triangle lying between ϕ' and ϕ'' (by the argument above), the path Q is nondegenerate. Let e be the edge of ϕ parallel to ξ_2 , and Δ the normal little triangle of G containing e . Then Δ lies between ϕ' and ϕ'' ; let P' be the path of H_1 leaving ϕ' and entering Δ . Note that \mathcal{K} contains no copy of Δ (again by the argument above). Next, let e' be the edge of Δ parallel to ξ_3 , and ∇ the turned-over triangle of G containing e' . Then ∇ lies between ϕ and ψ (as Q is nondegenerate); let Q' be the path of H_2 leaving ∇ and entering ψ . See the picture:



Now add to \mathcal{K} the triangles Δ, ∇ , the paths P', Q' , the degenerate path v_e (“connecting” ϕ and Δ in H_2) and the degenerate path $v_{e'}$ (“connecting” ∇ and Δ in H_3), assigning the attachments for them in an obvious way. Accordingly delete from \mathcal{K} the paths P, Q . This results in an “incomplete” c-configuration, but having a larger value of t , in which ∇ and ϕ'' have no attached paths in H_1 . (It cannot be improved straightforwardly since v and the middle point \tilde{u} of the edge of ∇ parallel to ξ_1 do not lie on one line of H_1). So we have a situation as in Case 1 and proceed in a similar way to transform \mathcal{K} into a correct c-configuration \mathcal{K}' either dominating the initial \mathcal{K} or being equivalent to \mathcal{K} but having a larger value of t .

This completes the proof of the lemma. ■

By Lemmas 6.1 and 6.3, every non-dominated proper essential configuration is equivalent to a puzzle configuration. This implies Theorem 3.1, in view of explanations in Sections 4,5.

Remark. Analysing the proof of Lemma 6.3, one sees that, in fact, a slightly sharper version of this lemma is obtained. It reads (taking into account assumption (6.3)(ii) and the construction of the minimal pre-configuration \mathcal{K}^{\min}):

(6.5) if a c-configuration \mathcal{K} is proper, essential and oriented and if \mathcal{K} is not dominated, then \mathcal{K} is equivalent to a puzzle configuration \mathcal{K}_Π such that the set of triangles of the puzzle Π includes all little triangles of G having copies in \mathcal{K} .

7 Concluding Remarks

We conclude this paper with several remarks.

First, for a cocirculation h in G and a tandem $\tau = (e, e')$, call $\delta_h(\tau) := h(e) - h(e')$ the *discrepancy* of h at τ . So h is concave if the discrepancy at each tandem is nonnegative. A more general problem is: (*) find a cocirculation h having a given border σ and obeying prescribed lower bounds c on the discrepancies: $\delta_h(\tau) \geq c(\tau)$ for each $\tau \in \mathcal{T}(G)$. This is reduced to the case of zero bounds when c comes up from another cocirculation g in G . More precisely, let $c(\tau) := \delta_g(\tau)$ for each tandem τ . Re-define the required border by $\sigma'(e) := \sigma(e) - g(e)$ for each boundary edge e . Then h' is a concave cocirculation with border σ' if and only if $h := h' + g$ is a cocirculation with border σ satisfying the lower bound c on the discrepancies. Thus, the corresponding changes in the puzzle inequalities (3.2) and in the monotone condition (2.6) give a solvability criterion for problem (*) with a cocirculation-induced c .

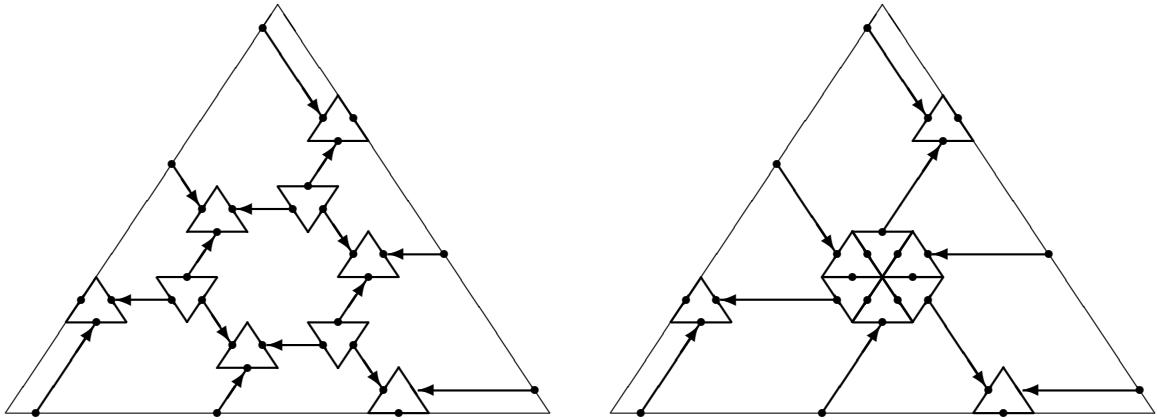


Figure 1: Two puzzles with equal boundaries

In particular, the puzzle criterion modified in this way works when all tandem discrepancies are required to be greater than or equal to a prescribed constant $\alpha \in \mathbb{R}$. This is because there exists a cocirculation g in G where the discrepancy at each tandem is exactly α . (Such a g is constructed easily: assuming w.l.o.g. that G is a 3-side grid of size n , put $g(e_i) := (k - 2i + 1)\alpha$ ($i = 1, \dots, k$) for each maximal straight path (e_1, \dots, e_k) in G .)

Second, from the sharper version of Lemma 6.3 given in (6.5) one derives that each puzzle Π determining a facet of the cone $\mathcal{B}(G)$ is (uniquely) determined by its boundary $b(\Pi)$.

Indeed, suppose Π_1, Π_2 are two different puzzles with $b(\Pi_1) = b(\Pi_2)$. Let \mathcal{K}_i stand for the c-configuration induced by Π_i ; one may assume that \mathcal{K}_i is proper and essential. Then $\mathcal{K} := \mathcal{K}_1 + \mathcal{K}_2$ is an oriented c-configuration equivalent to \mathcal{K}_i . Assume \mathcal{K} is not dominated and take a puzzle Π as in (6.5). We have $b(\mathcal{K}_\Pi) = b(\mathcal{K}_i)$, so the number q of triangles in Π is equal to the number q_1 of triangles in Π_1 , by (3.3). On the other hand, the fact that Π_1 and Π_2 are different implies that \mathcal{K} involves more little triangles of G compared with \mathcal{K}_1 . This implies $q > q_1$, by (6.5); a contradiction.

Different puzzles with equal boundaries do exist. An example for a 3-side grid is shown in Fig. 1.

A puzzle determined by its boundary is called *rigid*. Knutson, Tao and Woodward proved that in the case of 3-side grids the facet-determining puzzles are exactly the rigid ones. They also obtained a combinatorial characterization for the facet-determining puzzles, implying that such puzzles are recognizable in polynomial time.

Theorem 7.1 [7] *Let $\Pi = (\mathcal{F}, \mathcal{P})$ be a puzzle in a 3-side grid G such that \mathcal{F} is nonempty and different from the set of all little triangles of G . The following are equivalent:*

- (i) Π determines a facet of $\mathcal{B}(G)$;
- (ii) Π is rigid;
- (iii) Π admits no gentle circuits.

To explain the notion of gentle path/circuit, let R be the set of little rhombi of G that are split by a path in \mathcal{P} into two parallelograms. Let G_0 be the subgraph of G induced by the edges separating either a triangle in \mathcal{F} and a rhombus in R (the *tp-edges*), or a rhombus in R and a little triangle contained in no member of $\mathcal{F} \cup R$ (the *pn-edges*). Re-orient each tp-edge (resp. pn-edge) e so that the triangle of \mathcal{F} (resp. the rhombus in R) containing e lie on the right. A path or circuit P of G_0 is called *gentle* if, when moving along P from an edge to the next edge, the angle of turn is either 0° or 60° , never 120° . For example, the circuit surrounding the hexagon formed by the six central triangles in the right puzzle in Fig. 1 is gentle.

One can show that Theorem 7.1 remains valid for an arbitrary convex grid G . (Implication (i) \rightarrow (ii) has already been shown. The method of proof of (ii) \rightarrow (iii) and (iii) \rightarrow (i) given in [7] is applicable to an arbitrary convex grid, as it, in essence, does not depend on the shape of the convex region \mathcal{R} spanned by G . Roughly speaking, the proof of (ii) \rightarrow (iii) relies on a local transformation of a puzzle Π having a gentle circuit C . It creates another puzzle with the same boundary by re-arranging Π only within the 1-neighbourhood of C (being the union of little triangles and rhombi sharing common edges with C). The proof of (iii) \rightarrow (i) uses the function on the tp- and pn-edges whose value on an edge e is defined to be the number of all maximal gentle paths with the first edge e . When Π has no gentle circuits, this function (regardless of the shape of \mathcal{R}) is well-defined and it can easily be transformed into a concave cocirculation h_0 in G for which the tandem inequality is strict on each little rhombus separated by a tp- or pn-edge. Using h_0 , it is routine to construct $|E_0(G)| - 2$ concave cocirculations whose borders are linearly independent and orthogonal to the border of \mathcal{K}_Π .) We omit details of the proof here.

It is not difficult to check that any puzzle Π with $\mathcal{F} = \emptyset$ and $|\mathcal{P}| = 1$ is facet-determining as well (such a puzzle can arise when \mathcal{R} has ≥ 4 sides). When $\mathcal{F} = \emptyset$ and $|\mathcal{P}| \geq 2$, Π is already not facet-determining as it is the union of two disjoint puzzles.

Third, a result of Knutson and Tao [6] on integral honeycombs implies that a feasible integer-valued function σ on the boundary edges of a convex grid G is extendable to an integer concave cocirculation. In [5] one shows that a sharper property takes place: a concave cocirculation h in a convex grid G can be turned into an integer concave cocirculation preserving the values of h on all boundary edges e with $h(e) \in \mathbb{Z}$ and on each edge occurring in a little triangle where h is integral on the three edges.

Acknowledgement. I am thankful to Vladimir Danilov and Gleb Koshevoy for stimulating discussions on discrete concave functions and related topics.

References

- [1] V.I. Danilov and G.A. Koshevoy, Diskretnaya vypuklost' i nil'potentnye operatory, *Izvestiya Rossijskoi Akademii Nauk* **67** (1) (2003) 3–20, in Russian. (Title translation: Discrete convexity and nilpotent operators.)

- [2] V.I. Danilov and G.A. Koshevoy, Diskretnaya vognutost' i Ermitovy matritsy, in: *Trudy Matematicheskogo Instituta imeni Steklova*, vol. 241, Moscow, 2003, pp. 68–89, in Russian. (Title translation: Discrete concavity and Hermitian matrices.)
- [3] W. Fulton, Eigenvalues, invariant factors, highest weights, and Schubert calculus, *Bull. Amer. Math. Soc.* **37** (2000) 209–249.
- [4] A. Horn, Eigenvalues of sums of Hermitian matrices, *Pacific J. Math.* **12** (1962) 225–241.
- [5] A.V. Karzanov, Integer concave cocirculations and honeycombs, in: *Integer Programming and Combinatorial Optimization* (10th International IPCO Conference, NY, USA, June 2004, Proceedings) (*Lecture Notes in Computer Sci.* **3064**, Springer, 2004), pp. 368–387.
- [6] A. Knutson and T. Tao, The honeycomb model of $GL_n(\mathbb{C})$ tensor products I: Proof of the saturation conjecture, *J. Amer. Math. Soc.* **12** (1999) 1055–1090.
- [7] A. Knutson, T. Tao, and C. Woodward, The honeycomb model of $GL_n(\mathbb{C})$ tensor products II: Puzzles determine facets of the Littlewood-Richardson cone, *Preprint*, 2001; to appear in *J. Amer. Math. Soc.*