J Algebr Comb (2010) 32: 497–531
DOI 10.1007/s10801-010-0224-x

On maximal weakly separated set-systems

Vladimir I. Danilov · Alexander V. Karzanov · Gleb A. Koshevoy

Received: 10 November 2009 / Accepted: 23 March 2010 / Published online: 21 April 2010

Abstract For a permutation $\omega \in S_n$, Leclerc and Zelevinsky in Am. Math. Soc. Transl., Ser. 2 **181**, 85–108 (1998) introduced the concept of an *ω-chamber weakly separated collection* of subsets of $\{1, 2, \dots, n\}$ and conjectured that all inclusionwise maximal collections of this sort have the same cardinality $\ell(\omega) + n + 1$, where $\ell(\omega)$ is the length of ω. We answer this conjecture affirmatively and present a generalization and additional results.

Keywords Weakly separated sets · Rhombus tiling · Generalized tiling · Weak Bruhat order · Cluster algebras

1 Introduction

For a positive integer n, let $[n]$ denote the ordered set of elements $1, 2, \dots, n$. We deal with two binary relations on subsets of $[n]$.

(1.1) For $A, B \subseteq [n]$, we write:

(i) $A \lessdot B$ if $B - A$ is nonempty and if $i < j$ holds for any $i \in A - B$ and $j \in B - A$ (where $A' - B'$ stands for the set difference $\{i' \colon A' \ni i' \notin B'\}$);

V.I. Danilov · G.A. Koshevoy
Central Institute of Economics and Mathematics of the RAS, 47, Nakhimovskii Prospect, 117418 Moscow, Russia

V.I. Danilov
e-mail: danilov@cemi.rssi.ru

G.A. Koshevoy
e-mail: koshevoy@cemi.rssi.ru

A.V. Karzanov (✉)
Institute for System Analysis of the RAS, 9, Prospect 60 Let Oktyabrya, 117312 Moscow, Russia
e-mail: sasha@cs.isa.ru

(ii) $A \rhd B$ if both $A - B$ and $B - A$ are nonempty and if the set $B - A$ can be (uniquely) expressed as a disjoint union $B' \sqcup B''$ of nonempty subsets so that $B' \lessdot A - B \lessdot B''$.

Note that these relations need not be transitive in general. For example, $13 \lessdot 23 \lessdot 24$ but $13 \not\lessdot 24$; similarly, $346 \rhd 256 \rhd 157$ but $346 \not\rhd 157$, where for brevity we write $i \dots j$ instead of $\{i\} \cup \dots \cup \{j\}$.

Definition 1 Sets $A, B \subseteq [n]$ are called *weakly separated* (from each other) if either $A \lessdot B$, or $B \lessdot A$, or $A \rhd B$ and $|A| \geq B$, or $B \rhd A$ and $|B| \geq |A|$, or $A = B$. A collection $\mathcal{C} \subseteq 2^{[n]}$ is called weakly separated if any two of its members are weakly separated.

We will usually abbreviate the term "weakly separated collection" to "ws-collection". When a set X is weakly separated from a set Y (from all sets in a collection $\mathcal{C}$), we write $X \,\text{ws}\, Y$ (resp. $X \,\text{ws}\, \mathcal{C}$).

Definition 2 Let ω be a permutation on $[n]$. A subset $X \subset [n]$ is called an ω-*chamber set* if $i \in X$, $j < i$ and $\omega(j) < \omega(i)$ imply $j \in X$. A ws-collection $\mathcal{C} \subseteq 2^{[n]}$ is called an ω-*chamber ws-collection* if all members of $\mathcal{C}$ are ω-chamber sets.

These notions were introduced by Leclerc and Zelevinsky in [8] where their importance is demonstrated, in particular, in connection with the problem of characterizing quasicommuting quantum flag minors of a generic q-matrix. (Note that [8] deals with a relation $\prec$ which is somewhat different from $\lessdot$; nevertheless, Definition 1 is consistent with the corresponding definition in [8]. The term "ω-chamber" for a set X is motivated by the fact that such an X corresponds to a face, or *chamber*, in the pseudo-line arrangement related to some reduced word for ω; see [1].)

Let $\ell(\omega)$ denote the *length* of ω, i.e. the number of pairs $i < j$ such that $\omega(j) < \omega(i)$ (*inversions*). It is shown in [8] that the cardinality $|\mathcal{C}|$ of any ω-chamber ws-collection $\mathcal{C} \subseteq 2^{[n]}$ does not exceed $\ell(\omega) + n + 1$ and is conjectured (Conjecture 1.5 there) that this bound is achieved by *any* (inclusionwise) maximal collection among these:

(C) For any permutation ω on $[n]$, $|\mathcal{C}| = \ell(\omega) + n + 1$ holds for all maximal ω-chamber ws-collections $\mathcal{C} \subseteq 2^{[n]}$.

The main purpose of this paper is to answer this conjecture.

Theorem A *Statement* (C) *is valid.*

The *longest* permutation ω_0 on $[n]$ (defined by $i \mapsto n - i + 1$) is of especial interest, many results in [8] are devoted just to this case, and (C) with $\omega = \omega_0$ has been open so far as well. Since $\omega_0(j) > \omega_0(i)$ for any $j < i$, no "chamber conditions" are imposed in this case in essence, i.e. the set of ω_0-chamber ws-collections consists of all ws-collections. The length of ω_0 is equal to $\binom{n}{2}$, so the above upper bound turns into $\binom{n+1}{2} + 1$. Then the assertion in the above theorem is specified as follows.

Theorem B *All maximal ws-collections in* $2^{[n]}$ *have the same cardinality* $\binom{n+1}{2}+1$.

We refer to a ws-collection of this cardinality as a *largest* one and denote the set of these collections by $\mathbf{W}_n$. An important instance is the set $\mathcal{I}_n$ of all intervals $[p..q] := \{p, p+1, \ldots, q\}$ in $[n]$, including the "empty interval" $\emptyset$. One can see that for a ws-collection $\mathcal{C}$, the collection $\{[n]-X : X \in \mathcal{C}\}$ is weakly separated as well; it is called the *complementary ws-collection* of $\mathcal{C}$ and denoted by co-$\mathcal{C}$. Therefore, co-$\mathcal{I}_n$, the set of *co-intervals* in $[n]$, is also a largest ws-collection. In [8] it is shown that $\mathbf{W}_n$ is preserved under so-called *weak raising flips* (which transform one collection into another) and is conjectured (Conjecture 1.8 there) that in the poset structure on $\mathbf{W}_n$ induced by such flips, $\mathcal{I}_n$ and co-$\mathcal{I}_n$ are the unique minimal and unique maximal elements, respectively. That conjecture was affirmatively answered in [4].

We will show that Theorem A can be obtained relatively easily from Theorem B. The proof of the latter theorem is more intricate and takes the most part of this paper. The breakthrough step on this way consists of showing the following lattice property for a largest ws-collection $\mathcal{C}$: the partial order $\prec^*$ on $\mathcal{C}$ given by $A \prec^* B \Longleftrightarrow A \lessdot B$ & $|A| \le |B|$ forms a lattice. To prove this and some other intermediate results we will essentially use results and constructions from our previous work [4].

The main result in [4] shows the coincidence of four classes of collections: (i) the set of *semi-normal bases* of tropical Plücker functions on $2^{[n]}$; (ii) the set of *spectra* of certain collections of n curves on a disc in the plane, called *proper wirings* (which generalize commutation classes of pseudo-line arrangements); (iii) the set $\mathbf{ST}_n$ of *spectra* of so-called *generalized tilings* on an n-zonogon (a $2n$-gon representable as the Minkowsky sum of n generic line segments in the plane); and (iv) the set $\mathbf{W}_n$. Objects mentioned in (i) and (ii) are beyond our consideration in this paper (for definitions, see [4]), but we will extensively use the generalized tiling model and rely on the equality $\mathbf{ST}_n = \mathbf{W}_n$. Our goal is to show the following property:

Any ws-collection can be extended to the spectrum of some generalized tiling, whence Theorem B will immediately follow. Due to this property, generalized tilings give a geometric model for ws-collections.

Roughly speaking, a generalized tiling, briefly called a *g-tiling*, is a certain generalization of the notion of a *rhombus tiling*. While the latter is a subdivision of an n-zonogon Z into rhombi, the former is a cover of Z with rhombi that may overlap in a certain way. (It should be noted that rhombus tilings have been well studied, and one of important properties of these is that their "spectra" turn out to be exactly the maximal strongly separated collections, where $\mathcal{C} \subseteq 2^{[n]}$ is called *strongly separated* if any two members of $\mathcal{C}$ obey relation $\lessdot$ as in (1.1)(i). Leclerc and Zelevinsky explored such collections in [8] in parallel with ws-collections. In particular, they established a counterpart of Theorem A saying that the cardinality of any maximal strongly separated ω-chamber collection is exactly $\ell(\omega)+n+1$. For a wider discussion and related topics, see also [1, 3, 5–7, 9].)

This paper is organized as follows. Section 2 explains how to reduce Theorem A to Theorem B. Then we start proving the latter theorem; the whole proof lasts throughout Sects. 3–6. Section 3 recalls the definitions of a g-tiling and its spectrum and gives a review of properties of these objects established in [4] that are important for us. Sections 4 proves Theorem B under the assumption of validity of the above-mentioned lattice property of largest ws-collections. Then there begins a rather long way of proving the latter property stated in Theorem 4.1. In Sect. 5 we associate to a g-tiling T a certain acyclic directed graph $\Gamma = \Gamma_T$ whose vertex set is the spectrum $\mathfrak{S}_T$ of T (forming a largest ws-collection) and show that the natural partial order $\prec_\Gamma$ induced by Γ forms a lattice, which is not difficult. Section 6 is devoted to proving the crucial property that $\prec_\Gamma$ coincides with the partial order $\prec^*$ on $\mathfrak{S}_T$, thus yielding Theorem 4.1 and completing the proof of Theorem B; this is apparently the most sophisticated part of the paper where special combinatorial techniques of handling g-tilings are elaborated. The concluding Sect. 7 presents additional results and finishes with a generalization of Theorem A. This generalization (Theorem A′) deals with two permutations ω', ω on $[n]$ such that each inversion of ω' is an inversion of ω (in this case the pair (ω', ω) is said to obey the *weak Bruhat relation*). It asserts that all maximal ws-collections whose members X are ω-chamber sets and simultaneously satisfy the condition: $i \in X$ & $j > i$ & $\omega'(j) < \omega'(i) \Longrightarrow j \in X$, have the same cardinality, namely, $\ell(\omega) - \ell(\omega') + n + 1$. When ω' is the identical permutation $i \mapsto i$, this turns into Theorem A.

In the following, for a set $X \subset [n]$, distinct elements $i, \dots, j \in [n] - X$ and an element $k \in X$, we usually abbreviate $X \cup \{i\} \cup \dots \cup \{j\}$ as $Xi \dots j$, and $X - \{k\}$ as $X - k$.

2 Maximal ω-chamber ws-collections

In this section we explain how to derive Theorem A from Theorem B. The proof given here is direct and relatively short, though rather technical. Another proof, which is more geometric and appeals to properties of tilings, will be seen from a discussion in Sect. 7. Let ω be a permutation on $[n]$.

For $k = 0, \dots, n$, let I^k_ω denote the set $\omega^{-1}[k] = \{i : \omega(i) \in [k]\}$, called k-th *ideal* for ω (it is an ideal of the linear order on $[n]$ given by: $i \prec j$ if $\omega(i) < \omega(j)$). We will use the following auxiliary collection

$$\mathcal{C}^0 = \mathcal{C}^0_\omega := \left\{ I^k_\omega \cap [j..n] : 1 \le j \le \omega^{-1}(k),\ 0 \le k \le n \right\}, \tag{2.1}$$

where possible repeated sets are ignored and where $I^0_\omega := \emptyset$. The role of this collection is emphasized by the following

Theorem 2.1 *Let $X \subset [n]$ and $X \notin \mathcal{C}^0$. The following properties are equivalent*:

(i) X ⓦⓢ $\mathcal{C}^0$ (*i.e. X is weakly separated from all sets in $\mathcal{C}^0$*);
(ii) *X is an ω-chamber set.*

Due to this property, we call $\mathcal{C}^0$ the (canonical) *ω-checker*. (We shall explain in Sect. 7 that $\mathcal{C}^0$ is chosen to be the spectrum of a special tiling and that there are other tilings whose spectra can be taken as a checker in place of $\mathcal{C}^0$ in Theorem 2.1; see Corollary 7.2.) It is easy to verify that: (a) $\mathcal{C}^0$ is a ws-collection; (b) its subcollection

$$\mathcal{I}_\omega := \left\{ I^0_\omega, I^1_\omega, \dots, I^n_\omega \right\} \tag{2.2}$$

consists of ω-chamber sets; and (c) any member of $\mathcal{C}^0 - \mathcal{I}_\omega$ is not an ω-chamber set.

Relying on Theorems B and 2.1, we can prove Theorem A as follows. Given an ω-chamber ws-collection $\mathcal{C} \subset 2^{[n]}$, consider $\mathcal{C}' := \mathcal{C} \cup \mathcal{C}^0$. By Theorem 2.1, $\mathcal{C}'$ is a ws-collection. Also $\mathcal{C} \cap \mathcal{C}^0 \subseteq \mathcal{I}_\omega$, in view of (c) above. Extend $\mathcal{C}'$ to a largest ws-collection $\mathcal{D}$, which is possible by Theorem B. Let $\mathcal{D}' := (\mathcal{D} - \mathcal{C}_0) \cup \mathcal{I}_\omega$. Then $\mathcal{D}'$ includes $\mathcal{C}$ and is an ω-chamber ws-collection by Theorem 2.1. Since the cardinality of $\mathcal{D}'$ is always the same (as it is equal to $\binom{n+1}{2} + 1 - |\mathcal{C}^0 - \mathcal{I}_\omega|$), $\mathcal{D}'$ is a largest ω-chamber ws-collection, and now Theorem A follows from the fact that the upper size bound $\ell(\omega) + n + 1$ is achieved by some ω-chamber weakly (or even strongly) separated collection.

The rest of this section is devoted to proving Theorem 2.1. The proof of implication (i) $\Rightarrow$ (ii) falls into three lemmas. Let $X \subset [n]$ be such that $X \notin \mathcal{C}^0$ and $X \circledS \mathcal{C}^0$, and let $k := |X|$ and $Y := I_\omega^k$.

Lemma 2.2 *Neither $Y \lessdot X$ nor $Y \rhd X$ can take place.*

Proof Suppose $Y \lessdot X$ or $Y \rhd X$. Let k' be the maximum number such that either $Y' \lessdot X$ or $Y' \rhd X$, where $Y' := I_\omega^{k'}$. Then $k \le k'$ and $Y \subseteq Y'$. Define

$$\Delta := Y' - X \quad \text{and} \quad \Delta' := \{i \in X - Y' : \Delta \lessdot \{i\}\}.$$

Then $\Delta, \Delta' \neq \emptyset$ and $|\Delta| \ge |\Delta'|$. The maximality of k' implies that $|\Delta'| = 1$ and that the unique element of Δ', say, a, is exactly $\omega^{-1}(k'+1)$ (in all other cases either $k'+1$ fits as well, or X is not weakly separated from $I_\omega^{k'+1}$).

Let b be the *maximal* element in Δ. Then $a > b$ and the element $\tilde{k} := \omega(b)$ is at most k'. We assert that there is no $d \in X$ such that $d < b$. To see this, consider the sets $I_\omega^{\tilde{k}}$ and $Z := I_\omega^{\tilde{k}} \cap [b..n]$. Then $Z \in \mathcal{C}^0$ and $Z \subseteq I_\omega^{\tilde{k}} \subseteq Y'$. Therefore, $a \in X - Z$. Also $b \in Z - X$. Moreover, $Z - X = \{b\}$, by the maximality of b. Now if X contains an element $d < b$, then we have $|X| > |Z|$ (in view of $a, d \in X - Z$ and $|Z - X| = 1$) and $Z \rhd X$ (in view of $d < b < a$), which contradicts $X \circledS \mathcal{C}^0$.

Thus, all elements of X are greater than b. This and $X - Y' = \{a\}$ imply that the set $U := Y' \cap [b..n]$ satisfies $X - U = \{a\}$ and $U - X = \{b\}$. Then X coincides with the set $I_\omega^{k'+1} \cap [b+1..n]$. But the latter set belongs to $\mathcal{C}^0$ (since $\omega^{-1}(k'+1) = a > b$). So X is a member of $\mathcal{C}^0$; a contradiction. □

Lemma 2.3 *$X \rhd Y$ cannot take place.*

Proof Suppose $X \rhd Y$. Take the maximal k' such that the set $Y' := I_\omega^{k'}$ satisfies $X - Y' \neq \emptyset$. Then $k' \ge k$ and $|X - Y'| = 1$. Since $k \le k'$ implies $Y \subseteq Y'$, we have $X \rhd Y'$. Also $|Y' - X| \ge |Y - X| \ge 2$. Then $|X| < |Y'|$, contradicting $X \circledS Y'$. □

In view of $X \circledS Y$, Lemmas 2.2 and 2.3 imply that only the case $X \lessdot Y$ is possible.

Lemma 2.4 *Let $X \lessdot Y$. Then X is an ω-chamber set.*

Proof Suppose that there exist $i \in X$ and $j \notin X$ such that $j < i$ and $\omega(j) < \omega(i)$. Consider possible cases.

Case 1: $i \in Y$. Then $\omega(j) < \omega(i)$ implies that $j \in Y$. Take $d \in X - Y$ (existing since $|X| = |Y|$). Then $d < j$ (since $X \lessdot Y$ and $j \in Y - X$). Let $Y' := I_\omega^{\omega(j)}$. We have $j \in Y'$, $i \notin Y'$ and $d \notin Y'$. This together with $X \text{ ws } Y'$ and $d < j < i$ implies $Y' \rhd X$. But $|X| = |Y| > |Y'|$ (in view of $\omega(j) < \omega(i) \le k$); a contradiction.

Case 2: $i, j \notin Y$. Then $\omega(i), \omega(j) > k$. Take $a \in Y - X$. Since $X \lessdot Y$, we have $a > i$. Also $\omega(a) \le k$. Let $Y' := I_\omega^{\omega(j)}$. Then $|Y'| > |Y|$ (in view of $\omega(j) > k$). Also $a, j \in Y' - X$ and $i \in X - Y'$. Therefore, $X \rhd Y'$, contradicting $|X| = |Y| < |Y'|$.

Finally, the case with $i \notin Y$ and $j \in Y$ is impossible since $i > j$ and $X \lessdot Y$. □

Thus, (i) ⇒ (ii) in Theorem 2.1 is proven. Now we prove the other direction.

Lemma 2.5 *Let $X \subset [n]$ be an ω-chamber set. Then $X \text{ ws } \mathcal{C}^0$.*

Proof Consider an arbitrary set $Y = I_\omega^k \cap [j..n]$ in $\mathcal{C}^0$ (where $j \le \omega^{-1}(k)$). One may assume that both $X - Y$ and $Y - X$ are nonempty. Let $a \in Y - X$ and $b \in X - Y$. We assert that $a > b$ (whence $X \lessdot Y$ follows).

Indeed, $a \in Y$ implies $\omega(a) \le k$. If $b \notin I_\omega^k$, then $\omega(b) > k \ge \omega(a)$. In case $a < b$ we would have $a \in X$, by the ω-chamberness of X. Therefore, $a > b$, as required.

Now suppose $b \in I_\omega^k$. Then $b \in I_\omega^k - [j..n]$, and therefore, $b < j$. Since $j \le a$, we again obtain $a > b$. □

This completes the proof of Theorem 2.1, reducing Theorem A to Theorem B.

3 Generalized tilings and their properties

As mentioned in the Introduction, the proof of Theorem B will essentially rely on results on generalized tilings from [4]. This section starts with definitions of such objects and their spectra. Then we review properties of generalized tilings that will be important for us later: Subsect. 3.2 describes rather easy consequences from the defining axioms and Subsect. 3.3 is devoted to less evident properties.

3.1 Generalized tilings

Tiling diagrams that we deal with live within a zonogon, which is defined as follows.

In the upper half-plane $\mathbb{R} \times \mathbb{R}_+$, take n non-colinear vectors $\xi_1, \dots, \xi_n$ so that:

(3.1) (i) $\xi_1, \dots, \xi_n$ follow in this order clockwise around $(0, 0)$, and
(ii) all integer combinations of these vectors are different.

Then the set

$$Z = Z_n := \{\lambda_1\xi_1 + \dots + \lambda_n\xi_n : \lambda_i \in \mathbb{R},\ 0 \le \lambda_i \le 1,\ i = 1, \dots, n\}$$

is a $2n$-gon. Moreover, Z is a *zonogon*, as it is the sum of n line segments $\{\lambda\xi_i : 1 \le \lambda \le 1\}$, $i = 1, \dots, n$. Also it is the image by a linear projection π of the solid cube

$\mathrm{conv}(2^{[n]})$ into the plane $\mathbb{R}^2$, defined by $\pi(x) := x_1\xi_1 + \cdots + x_n\xi_n$. The boundary $\mathrm{bd}(Z)$ of Z consists of two parts: the *left boundary* $\mathrm{lbd}(Z)$ formed by the points (vertices) $z_i^\ell := \xi_1 + \cdots + \xi_i$ ($i = 0, \ldots, n$) connected by the line segments $z_{i-1}^\ell z_i^\ell := z_{i-1}^\ell + \{\lambda\xi_i : 0 \le \lambda \le 1\}$, and the *right boundary* $\mathrm{rbd}(Z)$ formed by the points $z_i^r := \xi_{i+1} + \cdots + \xi_n$ ($i = 0, \ldots, n$) connected by the line segments $z_i^r z_{i-1}^r$. So $z_0^\ell = z_n^r$ is the minimal vertex of Z, denoted as z_0, and $z_n^\ell = z_0^r$ is the maximal vertex, denoted as z_n. We direct each segment $z_{i-1}^\ell z_i^\ell$ from z_{i-1}^ℓ to z_i^ℓ and direct each segment $z_i^r z_{i-1}^r$ from z_i^r to z_{i-1}^r.

When it is not confusing, a subset $X \subseteq [n]$ is identified with the corresponding vertex of the n-cube and with the point $\sum_{i\in X}\xi_i$ in the zonogon Z (and we will usually use capital letters when we are going to emphasize that a vertex (or a point) is considered as a set). Due to (3.1)(ii), all such points in Z are different.

By a *(pure) tiling diagram* we mean a subdivision T of Z into *tiles*, each being a parallelogram of the form $X + \{\lambda\xi_i + \lambda'\xi_j : 0 \le \lambda, \lambda' \le 1\}$ for some $i < j$ and some subset $X \subset [n]$ (regarded as a point in Z); so the tiles are pairwise non-overlapping (have no common interior points) and their union is Z. A tile τ determined by X, i, j is called an *ij-tile* at X and denoted by $\tau(X; i, j)$. According to a natural visualization of τ, its vertices X, Xi, Xj, Xij are called the *bottom, left, right, top* vertices of τ and denoted by $b(\tau)$, $\ell(\tau)$, $r(\tau)$, $t(\tau)$, respectively. The edge from $b(\tau)$ to $\ell(\tau)$ is denoted by $\mathrm{bl}(\tau)$, and the other three edges of τ are denoted as $\mathrm{br}(\tau)$, $\mathrm{lt}(\tau)$, $\mathrm{rt}(\tau)$ in a similar way.

In fact, it is not important for our purposes which set of base vectors ξ_i is chosen, subject to (3.1). (In works on a similar subsect, it is most often when the ξ_i are assumed to have equal Euclidean norms; in this case each tile forms a rhombus and T is usually referred to as a *rhombus tiling*.) However, to simplify technical details and visualization, it will be convenient for us to assume that these vectors always have *unit height*, i.e. each ξ_i is of the form $(a_i, 1)$. Then each tile becomes a parallelogram of height 2. Accordingly, we say that: a point (subset) $Y \subseteq [n]$ is of *height* $|Y|$; the set of vertices of tiles in T having height h forms h-th *level*; and a point Y *lies to the right* of a point Y' if $|Y| = |Y'|$ and $\sum_{i\in Y}\xi_i \ge \sum_{i\in Y'}\xi_i$.

In a *generalized tiling*, or a *g-tiling*, some tiles may overlap. It is a collection T of tiles $\tau(X; i, j)$ which is partitioned into two subcollections T^w and T^b, of *white* and *black* tiles, respectively, obeying axioms (T1)–(T4) below. When $T^b = \emptyset$, T becomes a pure tiling.

We associate to T the directed graph $G_T = (V_T, E_T)$ whose vertices and edges are, respectively, the points and line segments occurring as vertices and sides in the tiles of T (not counting multiplicities). An edge connecting vertices X and Xi is directed from the former to the latter; such an edge (parallel to ξ_i) is called an edge with *label* i, or an *i-edge* ([4] uses the term "color" rather than "label"). For a vertex $v \in V_T$, the set of edges incident with v is denoted by $E_T(v)$, and the set of tiles having a vertex at v is denoted by $F_T(v)$.

(T1) Each boundary edge of Z belongs to exactly one tile. Each edge in E_T not contained in $\mathrm{bd}(Z)$ belongs to exactly two tiles. All tiles in T are different, in the sense that no two coincide in the plane.

(T2) Any two white tiles having a common edge do not overlap, i.e. they have no common interior point. If a white tile and a black tile share an edge, then these tiles do overlap. No two black tiles share an edge.

See the picture; here all edges are directed up and the black tiles are drawn in bold.

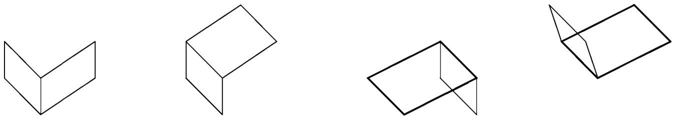

(T3) Let τ be a black tile. None of $b(\tau), t(\tau)$ is a vertex of another black tile. All edges in $E_T(b(\tau))$ *leave* $b(\tau)$, i.e. they are directed from $b(\tau)$. All edges in $E_T(t(\tau))$ *enter* $t(\tau)$, i.e. they are directed to $t(\tau)$.

We refer to a vertex $v \in V_T$ as *terminal* if v is the bottom or top vertex of some black tile. A nonterminal vertex v is called *ordinary* if all tiles in $F_T(v)$ are white, and *mixed* otherwise (i.e. v is the left or right vertex of some black tile). Note that a mixed vertex may belong, as the left or right vertex, to several black tiles.

Each tile $\tau \in T$ corresponds to a square in the solid cube conv($2^{[n]}$), denoted by $\sigma(\tau)$: if $\tau = \tau(X; i, j)$ then $\sigma(\tau)$ is the convex hull of the points X, Xi, Xj, Xij in the cube (so $\pi(\sigma(\tau)) = \tau$). (T1) implies that the interiors of these squares are pairwise disjoint and that $\bigcup(\sigma(\tau): \tau \in T)$ forms a 2-dimensional surface, denoted by D_T, whose boundary is the preimage by π of the boundary of Z. The last axiom is:

(T4) D_T is a disc, in the sense that it is homeomorphic to $\{x \in \mathbb{R}^2: x_1^2 + x_2^2 \le 1\}$.

The *reversed* g-tiling T^{rev} of a g-tiling T is formed by replacing each tile $\tau(X; i, j)$ of T by the tile $\tau([n] - Xij; i, j)$ (or, roughly speaking, by changing the orientation of all edges in E_T, in particular, in bd(Z)). Clearly (T1)–(T4) remain valid for T^{rev}.

The *spectrum* of a g-tiling T is the collection $\mathfrak{S}_T$ of (the subsets of $[n]$ represented by) *nonterminal* vertices in G_T. Figure 1 illustrates an example of g-tilings; here the unique black tile is drawn by thick lines and the terminal vertices are indicated by black rhombi.

The following result on g-tilings is of most importance for us.

Theorem 3.1 ([4]) *The spectrum $\mathfrak{S}_T$ of any generalized tiling T forms a largest ws-collection. Conversely, for any largest ws-collection $\mathcal{C} \subseteq 2^{[n]}$, there exists a gen-*

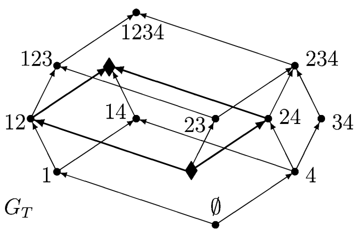

$$\mathfrak{S}_T = \{\emptyset, 1, 4, 12, 14, 23, 24, 34, 123, 234, 1234\}$$

Fig. 1 A g-tiling instance for $n = 4$

eralized tiling T on Z_n such that $\mathfrak{S}_T = \mathcal{C}$. (Moreover, such a T is unique and there is an efficient procedure to construct T from $\mathcal{C}$.)

In what follows, when it is not confusing, we may speak of a vertex or edge of G_T as a vertex or edge of T. The map σ of the tiles in T to squares in conv$(2^{[n]})$ is extended, in a natural way, to the vertices, edges, subgraphs or other objects in G_T. Note that the embedding of $\sigma(G_T)$ in the disc D_T is *planar* (unlike G_T and Z, in general), i.e. any two edges of $\sigma(G_T)$ can intersect only at their end points. It is convenient to assume that the clockwise orientations on Z and D_T are agreeable, in the sense that the image by σ of the boundary cycle $(z_0, z_1^\ell, \dots, z_n^\ell, z_1^r, \dots, z_n^r = z_0)$ is oriented clockwise around the interior of D_T. Then the orientations on a tile $\tau \in T$ and on the square $\sigma(\tau)$ are consistent when τ is white, and different when τ is black.

3.2 Elementary properties of generalized tilings

The properties of g-tilings reviewed in this subsection can be obtained rather easily from the above axioms; see [4] for more explanations. Let T be a g-tiling on $Z = Z_n$.

1. Let us say that the edges of T occurring in black tiles (as side edges) are *black*, and the other edges of T are *white*. For a vertex v and two edges $e, e' \in E_T(v)$, let $\Theta(e, e')$ denote the cone (with angle $< \pi$) in the plane pointed at v and generated by these edges (ignoring their directions). When another edge $e'' \in E_T(v)$ (a tile $\tau \in F_T(v)$) is contained in $\Theta(e, e')$, we say that e'' (resp. τ) lies *between* e and e''. When these e, e' are edges of a tile τ, we also write $\Theta(\tau; v)$ for $\Theta(e, e')$ (the conic hull of τ at v), and denote by $\theta(\tau, v)$ the angle of this cone taken with sign $+$ if τ is white, and sign $-$ if τ is black. The sum $\sum(\theta(\tau, v) \colon \tau \in F_T(v))$ is denoted by $\rho(v)$ and called the *full angle* at v. Terminal vertices of T behave as follows.

Corollary 3.2 *Let v be a terminal vertex belonging to a black ij-tile τ. Then*:

- (i) *v is not connected by edge with any other terminal vertex of T (in particular, $E_T(v)$ contains exactly two black edges, namely, those belonging to τ)*;
- (ii) *$E_T(v)$ contains at least one white edge and all such edges e, as well as all tiles in $F_T(v)$, lie between the two black edges in $E_T(v)$ (so e is a q-edge with $i < q < j$)*;
- (iii) $\rho(v) = 0$;
- (iv) *v does not belong to the boundary of Z (so each boundary edge e of Z, as well as the tile containing e, is white)*.

Note that (ii) implies that

(3.2) if a black tile τ and a white tile τ' share an edge and if v is their common nonterminal vertex (which is either left or right in both τ, τ'), then τ is contained in the cone $\Theta(\tau'; v)$.

Using this and applying Euler's formula to the planar graph $\sigma(G_T)$ on D_T, one can specify the full angles at nonterminal vertices.

Corollary 3.3 *Let $v \in V_T$ be a nonterminal vertex.*

(i) *If* v *belongs to* $\mathrm{bd}(Z)$, *then* $\rho(v)$ *is equal to the* (*positive*) *angle between the boundary edges incident to* v.
(ii) *If* v *is inner* (*i.e. not in* $\mathrm{bd}(Z)$), *then* $\rho(v) = 2\pi$.

2. Using (3.2) and Corollary 3.3, one can obtain the following useful (though rather lengthy) description of the local structure of edges and tiles at nonterminal vertices.

Corollary 3.4 *Let* v *be a nonterminal* (*ordinary or mixed*) *vertex of* T *different from* z_0, z_n. *Let* $e_1, \dots, e_p$ *be the sequence of edges leaving* v *and ordered clockwise around* v (*or by increasing their labels*), *and* $e'_1, \dots, e_{p'}$ *the sequence of edges entering* v *and ordered counterclockwise around* v (*or by decreasing their labels*). *Then there are integers* $r, r' \ge 0$ *such that*:

(i) $r + r' < \min\{p, p'\}$, *the edges* $e_{r+1}, \dots, e_{p-r'}$ *and* $e'_{r+1}, \dots, e'_{p'-r'}$ *are white, the other edges in* $E_T(v)$ *are black*, $r = 0$ *if* $v \in \mathrm{lbd}(Z)$, *and* $r' = 0$ *if* $v \in \mathrm{rbd}(Z)$;
(ii) *for* $q = r+1, \dots, p-r'-1$, *the edges* e_q, e_{q+1} *are spanned by a white tile* (*so such tiles have the bottom at* v *and lie between* e_{r+1} *and* $e_{p-r'}$);
(iii) *for* $q = r+1, \dots, p'-r'-1$, *the edges* e'_q, e'_{q+1} *are spanned by a white tile* τ (*so such tiles have the top at* v *and lie between* e'_{r+1} *and* $e'_{p'-r'}$);
(iv) *unless* $v \in \mathrm{lbd}(Z)$, *each of the pairs* $\{e_1, e'_{r+1}\}, \{e_2, e'_r\}, \dots, \{e_{r+1}, e'_1\}$ *is spanned by a white tile, and each of the pairs* $\{e_1, e'_r\}, \{e_2, e'_{r-1}\}, \dots, \{e_r, e'_1\}$ *is spanned by a black tile* (*all tiles have the right vertex at* v);
(v) *unless* $v \in \mathrm{rbd}(Z)$, *each of the pairs* $\{e_p, e'_{p'-r'}\}, \{e_{p-1}, e'_{p'-r'+1}\}, \dots, \{e_{p-r'}, e'_{p'}\}$ *is spanned by a white tile, and each of the pairs* $\{e_p, e'_{p'-r'+1}\}, \{e_{p-1}, e'_{p'-r'+2}\}, \dots, \{e_{p-r'+1}, e'_{p'}\}$ *is spanned by a black tile* (*all tiles have the left vertex at* v).

In particular, (a) *there is at least one white edge leaving* v *and at least one white edge entering* v; (b) *the tiles in* (ii)–(v) *give a full list of tiles in* $F_T(v)$; *and* (c) *any two tiles* $\tau, \tau' \in F_T(v)$ *with* $r(\tau) = \ell(\tau') = v$ *do not overlap* (*have no common interior point*).

Also: *for* $v = z_0, z_n$, *all edges in* $E_T(v)$ *are white and the pairs of consecutive edges are spanned by white tiles.*

(When v is ordinary, we have $r = r' = 0$.) The case with $p = 4$, $p' = 5$, $r = 2$, $r' = 1$ is illustrated in the picture; here the black edges are drawn in bold and the thin (bold) arcs indicate the pairs of edges spanned by white (resp. black) tiles.

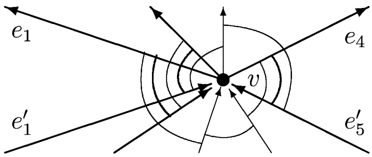

3. In view of (3.1)(ii), the graph $G_T = (V_T, E_T)$ is *graded* for each label $i \in [n]$, which means that for any closed path P in G_T, the amounts of forward i-edges and backward i-edges in P are equal. In particular, this easily implies that

(3.3) if four vertices and four edges of G_T form a (non-directed) cycle, then they are the vertices and edges of a tile (not necessarily contained in T).

Hereinafter, a path in a directed graph is meant to be a sequence $P = (\tilde{v}_0, \tilde{e}_1, \tilde{v}_1, \dots, \tilde{e}_r, \tilde{v}_r)$ in which each $\tilde{e}_p$ is an edge connecting vertices $\tilde{v}_{p-1}, \tilde{v}_p$; an edge $\tilde{e}_p$ is called *forward* if it is directed from $\tilde{v}_{p-1}$ to $\tilde{v}_p$ (denoted as $\tilde{e}_p = (\tilde{v}_{p-1}, \tilde{v}_p)$), and *backward* otherwise (when $\tilde{e}_p = (\tilde{v}_p, \tilde{v}_{p-1})$). When $v_0 = v_r$ and $r > 0$, P is a *closed path*, or a *cycle*. The path P is called *directed* if all its edges are forward, and *simple* if all vertices $v_0, \dots, v_r$ are different. P^{rev} denotes the reversed path $(\tilde{v}_r, \tilde{e}_r, \tilde{v}_{r-1}, \dots, \tilde{e}_1, \tilde{v}_0)$. Sometimes we will denote a path without explicitly indicating its edges: $P = (\tilde{v}_0, \tilde{v}_1, \dots, \tilde{v}_r)$. A directed graph is called *acyclic* if it has no directed cycles.

3.3 Strips, contractions, expansions, and others

In this subsection we describe additional, more involved, results and constructions concerning g-tilings that are given in [4] and will be needed to prove Theorem B. They are described in Propositions 3.5–3.11 below.

3.3.1 Strips in T

Definition 3 Let $i \in [n]$. An *i-strip* (or a *dual i-path*) in T is a maximal alternating sequence $Q = (e_0, \tau_1, e_1, \dots, \tau_r, e_r)$ of edges and tiles of T such that: (a) $\tau_1, \dots, \tau_r$ are different tiles, each being an ij- or ji-tile for some j, and (b) for $p = 1, \dots, r$, e_{p-1} and e_p are the opposite i-edges of τ_p.

(Recall that speaking of an $i'j'$-tile, we assume that $i' < j'$.) In view of axiom (T1), Q is determined uniquely (up to reversing it, and up to shifting cyclically when $e_0 = e_r$) by any of its edges or tiles. For $p = 1, \dots, r$, let $e_p = (v_p, v'_p)$. Define the *right boundary* of Q to be the (not necessarily directed) path $R_Q = (v_0, a_1, v_1, \dots, a_r, v_r)$, where a_p is the edge of τ_p connecting v_{p-1} and v_p. Similarly, the *left boundary* of Q is the path $L_Q = (v'_0, a'_1, v'_1, \dots, a'_r, v'_r)$, where a'_p is the edge of τ_p connecting v'_{p-1} and v'_p. Then two edges a_p, a'_p have the same label. Considering R_Q and using the fact that G_T is graded, one shows that

(3.4) Q cannot be cyclic, i.e. the edges e_0 and e_r are different.

In view of the maximality of Q, (3.4) implies that one of e_0, e_r belongs to the left boundary, and the other to the right boundary of the zonogon Z; we assume that Q is directed so that $e_0 \in \text{lbd}(Z)$ (justifying the assignment of the right and left boundaries for Q). Properties of strips are described in the following

Proposition 3.5 *For each $i \in [n]$, there is exactly one i-strip, say Q_i. It contains all i-edges of T, begins with the edge $z^\ell_{i-1}z^\ell_i$ of* lbd(Z) *and ends with the edge $z^r_i z^r_{i-1}$ of* rbd(Z). *Furthermore, each of R_{Q_i} and L_{Q_i} is a simple path, L_{Q_i} is disjoint from R_{Q_i} and is obtained by shifting R_{Q_i} by the vector ξ_i. An edge of R_{Q_i} is forward if and only if it belongs to either a white $i*$-tile or a black $*i$-tile in Q_i, and similarly for the edges of L_{Q_i}.*

3.3.2 Strip contractions

Let $i \in [n]$. Partition T into three subsets T_i^0, T_i^-, T_i^+, where T_i^0 consists of all $i*$- and $*i$-tiles, T_i^- consists of the tiles $\tau(X; i', j')$ with $i', j' \neq i$ and $i \notin X$, and T_i^+ consists of the tiles $\tau(X; i', j')$ with $i', j' \neq i$ and $i \in X$. Then T_i^0 is the set of tiles occurring in the i-strip Q_i, and the tiles in T_i^- are vertex disjoint from those in T_i^+.

Definition 4 The *i-contraction* of T is the collection T/i of tiles obtained by removing T_i^0, keeping the members of T_i^-, and replacing each $\tau(X; i', j') \in T_i^+$ by $\tau(X - i; i', j')$. The black/white coloring of tiles in T/i is inherited from T.

So the tiles of T/i live within the zonogon generated by the vectors ξ_q for $q \in [n] - i$. Clearly if we remove from the disc D_T the interiors of the edges and squares in $\sigma(Q_i)$, then we obtain two closed simply connected regions, one containing the squares $\sigma(\tau)$ for all $\tau \in T_i^-$, denoted as $D_{T_i^-}$, and the other containing $\sigma(\tau)$ for all $\tau \in T_i^+$, denoted as $D_{T_i^+}$. Then $D_{T/i}$ is the union of $D_{T_i^-}$ and $D_{T_i^+} - \epsilon_i$, where ϵ_i is i-th unit base vector in $\mathbb{R}^{[n]}$. In other words, $D_{T_i^+}$ is shifted by $-\epsilon_i$ and the path $\sigma(L_{Q_i})$ in it (the left boundary of $\sigma(Q_i)$) merges with the path $\sigma(R_{Q_i})$ in $D_{T_i^-}$. In general, $D_{T_i^-}$ and $D_{T_i^+} - \epsilon_i$ may intersect at some other points, and for this reason, $D_{T/i}$ need not be a disc. Nevertheless, $D_{T/i}$ is shown to be a disc in two important special cases: $i = n$ and $i = 1$; moreover, the following property holds.

Proposition 3.6 *The n-contraction T/n of T is a (feasible) g-tiling on the zonogon Z_{n-1} generated by the vectors $\xi_1, \dots, \xi_{n-1}$. Similarly, the 1-contraction $T/1$ is a g-tiling on the $(n-1)$-zonogon generated by the vectors $\xi_2, \dots, \xi_n$.*

(If wished, labels $2, \dots, n$ for $T/1$ can be renamed as $1', \dots, (n-1)'$.) We will use the n- and 1-contraction operations in Sects. 4 and 6.

3.3.3 Legal paths and strip expansions

Next we describe the n-expansion and 1-expansion operations; they are converse, in a sense, to the n-contraction and 1-contraction ones, respectively. We start with introducing the operation for n.

The *n-expansion* operation applies to a g-tiling T on the zonogon $Z = Z_{n-1}$ generated by $\xi_1, \dots, \xi_{n-1}$ and to a simple (not necessarily directed) path P in the graph G_T beginning at the minimal vertex z_0 and ending at the maximal vertex z_{n-1}^ℓ of Z. Then $\sigma(P)$ subdivides the disc D_T into two simply connected closed regions D', D'' such that: $D' \cup D'' = D_T$, $D' \cap D'' = \sigma(P)$, D' contains $\sigma(\mathrm{lbd}(Z))$, and D'' contains $\sigma(\mathrm{rbd}(Z))$. Let $T' := \{\tau \in T : \sigma(\tau) \subset D'\}$ and $T'' := T - T'$. We disconnect D', D'' along $\sigma(P)$ by shifting D'' by the vector ϵ_n, and then connect them by adding the corresponding strip of $*n$-tiles.

More precisely, we construct a collection $\tilde{T}$ of tiles on the zonogon Z_n, called the *n-expansion of T along P*, as follows. The tiles of T' are kept and each tile $\tau(X; i, j) \in T''$ is replaced by $\tau(Xn; i, j)$; the white/black coloring on these tiles is inherited. For each edge $e=(X,Xi)$ of P, we add tile $\tau(X;i,n)$, making it white if e is forward, and black if e is backward in P. The resulting $\tilde{T}$ need not be a g-tiling in general; for this reason, we impose additional conditions on P.

Definition 5 P as above is called an *n-legal path* if it satisfies the following three conditions:

(i) all vertices of P are nonterminal;
(ii) P contains no pair of consecutive backward edges;
(iii) for an i-edge e and a j-edge e' such that e, e' are consecutive edges occurring in this order in P: if e is forward and e' is backward in P, then $i>j$, and if e is backward and e' is forward, then $i<j$.

In view of (ii), P is represented as the concatenation of $P_1,\dots,P_{n-1}$, where P_h is the maximal subpath of P whose edges connect levels $h-1$ and h (i.e. are of the form (X,Xi) with $|X|=h-1$). In view of (iii), each path P_h has planar embedding in Z; it either contains only one edge, or is viewed as a *zigzag* path going *from left to right*. The first and last vertices of these subpaths are called *critical* vertices of P. The importance of legal paths is seen from the following

Proposition 3.7 *The n-expansion of T along P is a (feasible) g-tiling on the zonogon Z_n if and only if P is an n-legal path.*

Under the n-expansion operation, the path P generates the n-strip Q_n of the resulting g-tiling $\tilde{T}$; more precisely, the right boundary of Q_n is the reversed path P^{rev} to P, and the left boundary of Q_n is obtained by shifting P^{rev} by ξ_n. A possible fragment of P consisting of three consecutive edges e, e', e'' forming a zigzag and the corresponding fragment in Q_n (with two white tiles created from e, e'' and one black tile created from e') are illustrated in the picture; here the shifted e, e', e'' are indicated with tildes.

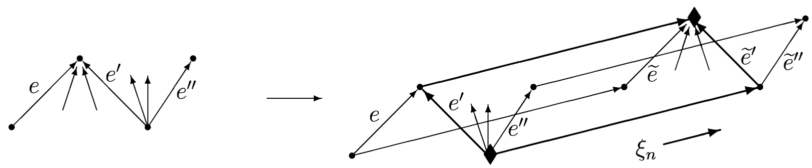

The n-contraction operation applied to $\tilde{T}$ returns the initial T. A relationship between n-contractions and n-expansions is described in the following

Proposition 3.8 *The correspondence $(T,P)\mapsto\tilde{T}$, where T is a g-tiling on Z_{n-1}, P is an n-legal path for T, and $\tilde{T}$ is the n-expansion of T along P, gives a bijection between the set of such pairs (T,P) and the set of g-tilings on Z_n.*

In its turn, the 1-*expansion operation* applies to a g-tiling T on the zonogon Z generated by the vectors $\xi_2,\dots,\xi_n$ (so we deal with labels $2,\dots,n$) and to a simple path P in G_T from the minimal vertex to the maximal vertex of Z; it produces a g-tiling $\tilde{T}$ on Z_n. This is equivalent to applying the n-expansion operation in the mirror-reflected situation: when label i is renamed as label $n-i+1$ (and accordingly a tile $\tau(X;i,j)$ in T is replaced by the tile $\tau(\{k: n-k+1\in X\}; n-j+1, n-i+1)$, preserving the basic vectors $\xi_1,\dots,\xi_n$). The corresponding "1-analogues" of the above results on n-expansions are as follows.

Proposition 3.9 (i) *The* 1*-expansion* $\tilde{T}$ *of* T *along* P *is a g-tiling on* Z_n *if and only if* P *is a* 1-legal path, *which is defined as in Definition* 5 *with the only difference that each subpath* P_h *of* P (*formed by the edges connecting levels* $h-1$ *and* h) *either contains only one edge*, *or is a zigzag path going* from right to left.

(ii) *The* 1*-contraction operation applied to* $\tilde{T}$ *returns the initial* T.

(iii) *The correspondence* $(T,P)\mapsto\tilde{T}$, *where* T *is a g-tiling on the zonogon generated by* $\xi_2,\dots,\xi_n$, P *is a* 1*-legal path for* T, *and* $\tilde{T}$ *is the* 1*-expansion of* T *along* P, *gives a bijection between the set of such pairs* (T,P) *and the set of g-tilings on* Z_n.

3.3.4 Principal trees

Let T be a g-tiling on $Z=Z_n$. We distinguish between two sorts of white edges e of G_T by saying that e is *fully white* if both of its end vertices are nonterminal, and *semi-white* if one end vertex is terminal. (Recall that an edge e of G_T is called white if no black tile contains e (as a side edge); the case when both ends of e are terminal is impossible, cf. Corollary 3.2(i).) In particular, all boundary edges of Z are fully white.

The following result on structural features of the set of white edges is obtained by using Corollary 3.4.

Proposition 3.10 *For* $h=1,\dots,n$, *let* H_h *denote the subgraph of* G_T *induced by the set of white edges connecting levels* $h-1$ *and* h (*i.e. of the form* (X,Xi) *with* $|X|=h-1$). *Then* H_h *is a forest. Furthermore*:

- (i) *there exists a component* (*a maximal tree*) K_h *of* H_h *that contains all fully white edges of* H_p (*in particular, the boundary edges* $z^\ell_{h-1}z^\ell_h$ *and* $z^r_{n-h+1}z^r_{n-h}$) *and no other edges*; *moreover,* K_h *has planar embedding in* Z;
- (ii) *any other component* K' *of* H_h *contains exactly one terminal vertex* v; *this* K' *is a star at* v *whose edges are the* (*semi-*)*white edges incident to* v.

It follows that the subgraph $G^{\mathrm{fw}}=G^{\mathrm{fw}}_T$ of G_T induced by the fully white edges has planar embedding in Z. We refer to K_h in (i) of the proposition as the *principal tree* in H_h. The common vertices of two neighboring principal trees K_h, K_{h+1} will play an important role later; we call them *critical vertices* for T in level h.

3.3.5 Additional facts

Two more useful facts (which look so natural but their proofs are not straightforward) concern relations between vertices and edges in G_T and tiles in T.

Proposition 3.11 (i) *Any two nonterminal vertices of the form X, Xi in G_T are connected by edge*. (Such an edge need not exist when some of X, Xi is terminal.)

(ii) *If four* nonterminal *vertices are connected by four edges forming a cycle in G_T, then there is a tile in T having these vertices and edges*. (Cf. (3.3).)

4 Proof of Theorem B

In this section we explain how to obtain Theorem B under the assumption of validity of the following statement, which is interesting in its own right. Recall that for sets $A, B \subseteq [n]$, we write $A \prec^* B$ if $A \lessdot B$ and $|A| \le |B|$.

Theorem 4.1 *Let $\mathcal{C}$ be a largest ws-collection. Then the partial order on $\mathcal{C}$ determined by $\prec^*$ forms a lattice.*

Here the fact that $(\mathcal{C}, \prec^*)$ is a poset (partially ordered set) is an immediate consequence of the following simple, but important, property established in [8], which describes a situation when the relation $\lessdot$ becomes transitive:

(4.1) for sets $A, A', A'' \subseteq [n]$, if $A \lessdot A' \lessdot A''$, A ws A'' and $|A| \le |A'| \le |A''|$, then $A \lessdot A''$.

Theorem 4.1 will be proved in Sects. 5 and 6. Relying on this, the method of proving that any ws-collection $\mathcal{C} \subseteq 2^{[n]}$ is contained in a largest ws-collection in $2^{[n]}$ (yielding Theorem B) is roughly as follows. We first reduce $\mathcal{C}$ to a ws-collection of subsets of $[n-1]$ and then extend the latter to a maximal ws-collection $\mathcal{C}'$ in $2^{[n-1]}$. One may assume by induction on n that $\mathcal{C}'$ is a largest ws-collection; so, by Theorem 3.1, $\mathcal{C}'$ is the spectrum of some g-tiling T' on the zonogon Z_{n-1}. Theorem 4.1 applied to $\mathcal{C}'$ is then used to show the existence of an appropriate n-legal path P for T'. Applying the n-expansion operation to (T', P) (as described in Subsect. 3.3.3), we obtain a g-tiling T on Z_n whose spectrum contains $\mathcal{C}$, whence the result follows.

We need two lemmas (where we write $X \underline{\lessdot} Y$ if either $X \lessdot Y$ or $X = Y$).

Lemma 4.2 (i) *Let $A \underline{\lessdot} C$ and $B \underline{\lessdot} C$. Then either $C \subset A \cup B$ or $A \cup B \underline{\lessdot} C$.*

(ii) *Let $C \underline{\lessdot} A$ and $C \underline{\lessdot} B$. Then $C \underline{\lessdot} A \cup B$.*

Proof (i) If $C \subset A \cup B$ or $A \cup B \subseteq C$, we are done. So assume that both $C - (A \cup B)$ and $(A \cup B) - C$ are nonempty and consider an element c in the former and an element x in the latter of these sets. One may assume that $x \in A$. Since $x \notin C$, $c \notin A$, and $A \underline{\lessdot} C$, we have $x < c$. This implies $A \cup B \underline{\lessdot} C$.

(ii) If $C \subseteq A \cup B$, we are done. So assume this is not the case, and let $c \in C - (A \cup B)$. Let x be an element of the set $(A \cup B) - C$ (which is, obviously, nonempty). One may assume that $x \in A - C$. Then $c \in C - A$ and $C \underline{\lessdot} A$ imply $c < x$, as required. □

Definition 6 Let $\mathcal{L}, \mathcal{R} \subseteq 2^{[n']}$. We call $(\mathcal{L}, \mathcal{R})$ a *left–right pair*, or, briefly, an *lr-pair*, if $\mathcal{L} \cup \mathcal{R}$ is a ws-collection and

(LR): $L \underline{\lessdot} R$ holds for any $L \in \mathcal{L}$ and $R \in \mathcal{R}$ with $|L| \le |R|$.

Lemma 4.3 *Let* $(\mathcal{L}, \mathcal{R})$ *be an lr-pair in* $2^{[n']}$.

(i) *Suppose that there are* $X \subseteq [n']$ *and* $i, j, k \in X$ *such that*: $i < j < k$, *the sets* $X - k, X - j$ *belong to* $\mathcal{L}$, *and the sets* $X - j, X - i$ *belong to* $\mathcal{R}$. *Then* $(\mathcal{L} \cup \{X\}, \mathcal{R})$ *is an lr-pair as well.*

(ii) *Symmetrically, suppose that there are* $X \subseteq [n']$ *and* $i, j, k \notin X$ *such that*: $i < j < k$, *the sets* Xi, Xj *belong to* $\mathcal{L}$, *and the sets* Xj, Xk *belong to* $\mathcal{R}$. *Then* $(\mathcal{L}, \mathcal{R} \cup \{X\})$ *is an lr-pair as well.*

Proof To see (i), we first show that X is weakly separated from any member Y of $\mathcal{L} \cup \mathcal{R}$.

Suppose that $|X| = |Y| + 1$. If $Y \in \mathcal{L}$, we argue as follows. Since $X - i, X - j \in \mathcal{R}$ and $|X - i| = |X - j| = |Y|$, we have $Y \precsim X - i$ and $Y \precsim X - j$. By (ii) in Lemma 4.2, we obtain that $Y \precsim (X - i) \cup (X - j) = X$. In case $Y \in \mathcal{R}$, we have $X - j, X - k \precsim Y$, and now (i) in Lemma 4.2 implies that either $Y \lessdot X$ or $X \precsim Y$.

Now suppose that X is not weakly separated from some $Y \in \mathcal{L} \cup \mathcal{R}$ with $|X| \neq |Y| + 1$. Three cases are possible.

(1) Let $|X| < |Y|$. Then one easily shows that there are $a, c \in Y - X$ and $b \in X - Y$ such that $a < b < c$; cf. Lemma 3.8 in [8]. The element b belongs to some set X' among $X - i, X - j, X - k$. Then $b \in X' - Y$ and $a, c \in Y - X'$, implying $X' \rhd Y$ (since X' (ws) Y). But $|X'| < |Y|$; a contradiction.

(2) Let $|X| = |Y|$. Then there are $a, c \in Y - X$ and $b, d \in X - Y$ such that either $a < b < c < d$ or $a > b > c > d$, by the same lemma in [8]. But both b, d belong to at least one set X' among $X - i, X - j, X - k$. So X', Y are not weakly separated; a contradiction.

(3) Let $|X| > |Y| + 1$. Then $a < b < c$ for some $a, c \in X - Y$ and $b \in Y - X$. Both a, c belong to some set X' among $X - i, X - j, X - k$. Then $Y \rhd X'$. But $|X'| = |X| - 1 > |Y|$; a contradiction.

Thus, $\mathcal{L} \cup \mathcal{R} \cup \{X\}$ is a ws-collection. It remains to check that $X \precsim R$ for any $R \in \mathcal{R}$ with $|R| \geq |X|$ (then $(\mathcal{L} \cup \{X\}, \mathcal{R})$ is an lr-pair). Since $X - j \in \mathcal{L}$ and $|X - j| < |R|$, we have $X - j \lessdot R$. Similarly, $X - k \lessdot R$. So, by Lemma 4.2(i), $X \precsim R$, as required. (The case $R \subset X$ is impossible since $|R| \geq |X|$.) This yields (i).

Validity of (ii) follows from (i) applied to the complementary lr-pair $(\{[n'] - R : R \in \mathcal{R}\}, \{[n'] - L : L \in \mathcal{L}\})$. □

Now we start proving Theorem B. Let $\mathcal{C} \subseteq 2^{[n]}$ be a ws-collection. The goal is to show that $\mathcal{C}$ is contained in a largest ws-collection in $2^{[n]}$. We use induction on n.

Form the collections $\mathcal{L} := \{X \subseteq [n-1] : X \in \mathcal{C}\}$ and $\mathcal{R} := \{X \subseteq [n-1] : Xn \in \mathcal{C}\}$. By easy observations in [8, Sect. 3], $\mathcal{L} \cup \mathcal{R}$ is a ws-collection and, furthermore, $(\mathcal{L}, \mathcal{R})$ is an lr-pair. Let us extend $(\mathcal{L}, \mathcal{R})$ to a maximal lr-pair $(\bar{\mathcal{L}}, \bar{\mathcal{R}})$ in $2^{[n-1]}$, i.e. $\mathcal{L} \subseteq \bar{\mathcal{L}}$, $\mathcal{R} \subseteq \bar{\mathcal{R}}$, and neither $\bar{\mathcal{L}}$ nor $\bar{\mathcal{R}}$ can be further extended. In particular, $\bar{\mathcal{L}}$ contains the intervals $[i]$ and $\bar{\mathcal{R}}$ contains the intervals $[i..n-1]$ for each i (including the empty interval).

By induction, there exists a largest ws-collection $\mathcal{C}' \subseteq 2^{[n-1]}$ containing $\bar{\mathcal{L}} \cup \bar{\mathcal{R}}$. By Theorem 3.1, $\mathcal{C}'$ is the spectrum $\mathfrak{S}_{T'}$ of some g-tiling T' on the zonogon Z_{n-1}. By Theorem 4.1, the poset $\mathcal{P}$ determined by $\mathcal{C}'$ and $\prec^*$ is a lattice.

For $h=0,\dots,n-1$, let $\mathcal{C}'_h, \bar{\mathcal{L}}_h, \bar{\mathcal{R}}_h$ consist of the sets X with $|X|=h$ in $\mathcal{C}', \bar{\mathcal{L}}, \bar{\mathcal{R}}$, respectively. Let $C_h \in \mathcal{C}'$ be a maximal element in the poset $\mathcal{P}$ such that

(4.2) $C_h \preceq^* R$ for all $R \in \bar{\mathcal{R}}_h \cup \dots \cup \bar{\mathcal{R}}_{n-1}$.

Since $\mathcal{P}$ is a lattice, C_h exists and is unique, and we have:

(4.3) (i) $L \preceq^* C_h$ for all $L \in \bar{\mathcal{L}}_0 \cup \dots \cup \bar{\mathcal{L}}_h$;
(ii) $C_0 \preceq^* C_1 \preceq^* \dots \preceq^* C_n$,

where (i) follows from condition (LR) in the definition of lr-pairs. Note that for each h, both sets $\bar{\mathcal{L}}_h$ and $\bar{\mathcal{R}}_h$ are nonempty, as the former contains the interval $[h]$ (viz. the vertex z^ℓ_h of Γ') and the latter contains $[n-h..n-1]$ (viz. the vertex z^r_{n-h}). Also for any $L \in \bar{\mathcal{L}}_h$ and $R \in \bar{\mathcal{R}}_h$, the facts that $L \prec^* C_h \prec^* R$ (cf. (4.2), (4.3)) and $|L|=|R|=h$ imply that $|C_h|=h$, whence $C_h \in \mathcal{C}'_h$.

If $X \in \mathcal{C}'_h$ and $X \prec^* C_h$ (resp. $C_h \prec^* X$), then X must belong to $\bar{\mathcal{L}}$ (resp. $\bar{\mathcal{R}}$); this follows from the maximality of $(\bar{\mathcal{L}}, \bar{\mathcal{R}})$ and relations (4.2) and (4.3)(i) (in view of the transitivity of $\prec^*$). Also the maximality of $(\bar{\mathcal{L}}, \bar{\mathcal{R}})$ implies that C_h belongs to both $\bar{\mathcal{L}}_h$ and $\bar{\mathcal{R}}_h$. We assert that C_h is a critical vertex (in level h) of the graph $G_{T'}$, for each h; see the definition in Subsect. 3.3.4.

To see this, take a white edge e leaving the vertex C_h in $G_{T'}$ (unless $h=n-1$); it exists by Corollary 3.4. Suppose that the end vertex X of e is terminal. Then X is the top vertex $t(\tau)$ of some black tile $\tau \in T'$; in particular, X is not in $\mathfrak{S}_{T'} = \mathcal{C}'$. Since e is white, there are white tiles $\tau', \tau'' \in F_\tau(X)$ with $r(\tau') = \ell(\tau'') = C_h$. Then: (a) the vertices $X' := \ell(\tau')$, C_h and $X'' := r(\tau'')$ are of the form $X-k, X-j, X-i$, respectively, for some $i<j<k$; (b) X' belongs to $\bar{\mathcal{L}}_h$ (since X' is nonterminal and, obviously, $X' \prec^* C_h$); and (c) X'' belongs to $\bar{\mathcal{R}}_h$ (since $C_h \prec^* X''$). But then, by Lemma 4.3(i), $\bar{\mathcal{L}}$ can be increased by adding the new element X, contrary to the maximality of $(\bar{\mathcal{L}}, \bar{\mathcal{R}})$. So the edge e is fully white. In a similar fashion (using (ii) in Lemma 4.3), one shows that C_h has an entering fully white edge (unless $h=0$). Thus, C_h is critical, as required.

Finally, by Proposition 3.10, each pair of critical vertices C_{h-1}, C_h is connected by a path P_h in the principal tree K_h. We assert that P_h either has only one edge or is a zigzag path going from left to right. Indeed, if this is not so, then P_h is a zigzag path going from right to left, i.e. P_h is of the form $(C_{h-1}=X_1, Y_1, \dots, X_k, Y_k = C_h)$ with $k \ge 2$, the vertices X_p (resp. Y_p) are in level $h-1$ (resp. h), and for labels i_p of (forward) edges (X_p, Y_p) and labels j_p of (backward) edges (Y_p, X_{p+1}), one holds $\dots > i_p < j_p > i_{p+1} < \dots$. Then $X_k \prec^* X_{k-1} \prec^* \dots \prec^* X_1$ (in view of $X_p - X_{p+1} = \{j_p\}$ and $X_{p+1} - X_p = \{i_p\}$). This and $C_{h-1} \prec^* C_h$ imply $X_{k-1} \prec^* Y_k$, which is impossible since $X_{k-1} - Y_k = \{j_{k-1}\}$, $Y_k - X_{k-1} = \{i_{k-1}, i_k\}$ and $i_{k-1}, i_k < j_{k-1}$.

It follows that the concatenation of the paths $P_1, \dots, P_{n-1}$ gives an n-legal path P in $G_{T'}$. By Proposition 3.7, the n-expansion of T' along P is a feasible g-tiling T on the zonogon Z_n, and now it is straightforward to check that the initial collection $\mathcal{C}$ is contained in the spectrum $\mathfrak{S}_T$ of T (which is a largest ws-collection, by Theorem 3.1).

This completes the proof of Theorem B (provided validity of Theorem 4.1).

5 The auxiliary graph

The proof of Theorem 4.1 is divided into two stages. In order to prove the theorem, we are forced to include into consideration an additional combinatorial object. This is a certain acyclic directed graph Γ_T associated to a g-tiling T on the zonogon Z_n; it is different from the graph G_T and its vertex set is the spectrum $\mathfrak{S}_T$ of T. In this section we define Γ_T, called the *auxiliary graph* for T, and show that the natural partial order on $\mathfrak{S}_T$ determined by Γ_T is a lattice; this is the first, easier, stage of the proof of Theorem 4.1. The second, crucial, stage will be given in the next section; it consists of proving that the partial order of Γ_T coincides with $\prec^*$ (within $\mathfrak{S}_T$). Then Theorem 4.1 will follow from the fact that the largest ws-collections in $2^{[n]}$ are exactly the spectra of g-tilings on Z_n, i.e. from Theorem 3.1.

Construction of $\Gamma = \Gamma_T$ Given a g-tiling T on Z_n, the vertex set of Γ is $\mathfrak{S}_T$. The edge set of Γ consists of two subsets: the set $E^{\rm asc}$ of fully white edges of G_T (the edge set of $G^{\rm fw}$), and the set $E^{\rm hor}$ of edges corresponding to the "horizontal diagonals" of white tiles, namely, for each $\tau \in T^{\rm w}$, we form edge e_τ going from $\ell(\tau)$ to $r(\tau)$. An edge in $E^{\rm asc}$ is called *ascending* (it goes from some level h to the next level $h+1$, having the form (X, Xi) with $|X| = h$). An edge in $E^{\rm hor}$ is called *horizontal* (as it connects vertices of the same height).

In particular, Γ contains $\mathrm{bd}(Z_n)$ (since all boundary edges are fully white). Figure 2 compares the graph G_T drawn in Fig. 1 and the graph Γ_T for the same T; here the ascending edges of Γ are indicated by ordinary lines (which should be directed up), and the horizontal edges by double lines or arcs (which should be directed from left to right).

In what follows we write $\prec_{G'}$ for the natural partial order on the vertices of an acyclic directed graph G', i.e. $x \prec_{G'} y$ if vertices x, y are connected in G' by a directed path from x to y.

Since each ascending edge of Γ goes from one level to the next one and each horizontal edge goes from left to right, the graph Γ is acyclic; so $\prec_\Gamma$ is a partial order on $\mathfrak{S}_T$. Moreover, Γ possesses the following nice property.

Proposition 5.1 *The poset* $(\mathfrak{S}_T, \prec_{\Gamma_T})$ *is a lattice.*

This is provided by the following lattice property of acyclic planar graphs; it is pointed out in [2] (p. 32, Exer. 7(a)), but for the completeness of our description, we give a proof.

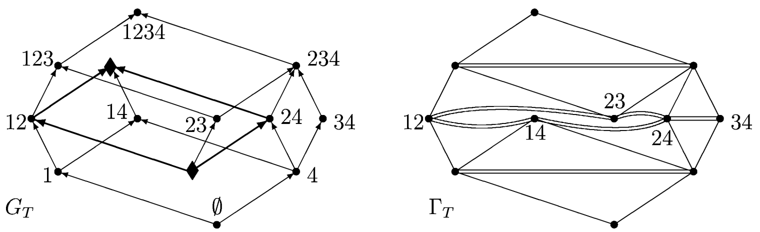

Fig. 2 Graphs G_T and Γ_T

Lemma 5.2 *Let $G' = (V', E')$ be a planar acyclic directed graph embedded in the plane. Suppose that the partial order $\mathcal{P} = (V', \prec_{G'})$ has a unique minimal element s and a unique maximal element t and that both s, t are contained in (the boundary of) the same face of G'. Then $\mathcal{P}$ is a lattice.*

Proof Let $\prec$ stand for $\prec_{G'}$. One may assume that both s, t belong to the outer (unbounded) face of G'.

Consider two vertices $x, y \in V'$. Let L be the set of maximal elements in $\{v \in V' : v \preceq x, y\}$, and U the set of minimal elements in $\{v \in V' : v \succeq x, y\}$. Since $\mathcal{P}$ has unique minimal and maximal elements, both L, U are nonempty. We have to show that $|L| = |U| = 1$. Below by a path we mean a directed path.

Suppose, for a contradiction, that L contains two distinct elements a, b. Take four paths in G' connecting $\{a, b\}$ to $\{x, y\}$: a path P_x from a to x, a path P_y from a to y, a path P'_x from b to x, and a path P'_y from b to y. Then P_x meets P_y only at the vertex a and is disjoint from P'_y (otherwise at least one of the lower bounds a, b for x, y is not maximal). Similarly, P'_x meets P'_y only at b and is disjoint from P_y. Also we may assume that $P_x \cap P'_x = \{x\}$ (otherwise we could replace x by another common point x' of P_x, P'_x; then $a, b \in L(x', y)$ and x' is "closer" to a, b than x). Similarly, we may assume that $P_y \cap P'_y = \{y\}$. Let R be the closed region (homeomorphic to a disc) surrounded by P_x, P_y, P'_x, P'_y. The fact that s, t belong to the outer face easily implies that they lie outside R.

Take a path Q_a from s to a and a path Q_b from s to b. For any vertex $v \neq a$ of Q_a, relation $v \prec a$ implies that v belongs to neither $P_x \cup P_y$ (otherwise $a \prec v$ would hold) nor $P'_x \cup P'_y$ (otherwise $b \preceq v \prec a$ would hold). Therefore, Q_a meets R only at a. Similarly, $Q_b \cap R = \{b\}$.

Let v be the last common vertex of Q_a and Q_b. Take the part Q' of Q_a from v to a, and the part Q'' of Q_b from v to b. Then $Q' \cap Q'' = \{v\}$. Observe that R is contained either in the closed region R_1 surrounded by Q', Q'', P_x, P'_x, or in the closed region R_2 surrounded by Q', Q'', P_y, P'_y. One may assume that $R \subset R_1$ (this case is illustrated in the picture below). Then y is an interior point in R_1. Obviously, $t \notin R_1$. Now since $y \prec t$, there exists a path from y to t in G'. This path meets the boundary of R_1 at some vertex z. But if z occurs in $P_x \cup P'_x$, then $y \prec z \preceq x$, and if z occurs in Q' (in Q''), then $y \prec z \preceq a$ (resp. $y \prec z \preceq b$); a contradiction.

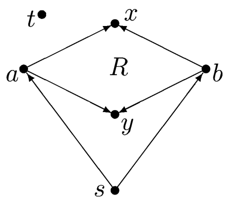

Thus, $|L| = 1$. The equality $|U| = 1$ is obtained by reversing the edges of G'. □

Now we prove Proposition 5.1 as follows. Consider the image $\sigma(\Gamma)$ of Γ on the disc D_T, where the image $\sigma(e)$ of the horizontal edge e drawn in a white tile τ is naturally defined to be the corresponding directed diagonal of the square $\sigma(\tau)$. Since the embedding of $\sigma(G_T)$ in D_T is planar, so is the embedding of $\sigma(\Gamma)$. Also:

(i) $\sigma(\mathrm{bd}(Z))$ is the boundary of D_T; (ii) each boundary vertex of Z lies on a directed path from z_0 to z_n in G_T, which belongs to Γ as well; and (iii) Γ is acyclic. Next, if a nonterminal vertex of T does not belong to the left (right) boundary of Z, then v is the right (resp. left) vertex of some white tile, as is seen from (iv) (resp. (v)) in Corollary 3.4. This implies that

(5.1) for $v \in \mathfrak{S}_T$, if v is not in $\mathrm{lbd}(Z)$ (not in $\mathrm{rbd}(Z)$), then there exists a horizontal edge in Γ entering (resp. leaving) v.

Thus, z_0 and z_n are the unique minimal and maximal vertices in Γ. Applying Lemma 5.2 to $\sigma(\Gamma)$, we conclude that Γ determines a lattice, as required.

6 Equality of two posets

The goal of this section is to show the following crucial property of the auxiliary graph $\Gamma = \Gamma_T$ introduced in the previous section.

Theorem 6.1 (Auxiliary Theorem) *For a g-tiling T on Z_n, the partial orders on $\mathfrak{S}_T$ given by $\prec^*$ and by $\prec_\Gamma$ are equal.*

This together with Proposition 5.1 will imply Theorem 4.1, thus completing the whole proof of the main results (Theorem B and, further, Theorem A) of this paper, as is explained in Sect. 4. The proof of Theorem 6.1 is given throughout this section. We keep notation from Sect. 5.

The fact that $\prec_\Gamma$ implies $\prec^*$ is easy. Indeed, since each edge $e = (A, B)$ of Γ is either of the form (X, Xi) (when e is ascending) or (Xi, Xj) with $i < j$ (when e is horizontal), we have $A \lessdot B$ and $|A| \le |B|$, whence $A \prec^* B$. Then for any $C, D \in \mathfrak{S}_T$ satisfying $C \prec_\Gamma D$, the relation $C \prec^* D$ is obtained by considering a directed path from C to D in Γ and using the transitivity property (4.1).

It remains to show the converse implication, i.e. the following

Proposition 6.2 *For a g-tiling T on $Z = Z_n$, let two sets (nonterminal vertices) $A, B \in \mathfrak{S}_T$ satisfy $A \prec^* B$. Then $A \prec_\Gamma B$, i.e. the graph $\Gamma = \Gamma_T$ contains a directed path from A to B.*

The rest of this section is devoted to proving this proposition, which is rather long and appeals to results on contractions and expansions mentioned in Subsects. 3.3.3, 3.3.4. Moreover, we need to conduct a more meticulous analysis of structural features of the graphs G_T and Γ, and of the action of contraction operations on Γ; this is done in Subsects. 6.1–6.3. Using these, we then prove the desired implication in Subsect. 6.4.

6.1 Refined properties of the graphs G^{fw} and Γ

We start with one fact which immediately follows from the planarity of principal trees K_h defined in Subsect. 3.3.4.

(6.1) the edges of K_h can be (uniquely) ordered as $e_1, \dots, e_p$ so that for $1 \le q < q' \le p$, the edge $e_{q'}$ lies to the right of e_q (in particular, $e_1 = z^\ell_{h-1} z^\ell_h$ and $e_p = z^r_{n-h+1} z^r_{n-h}$); equivalently, consecutive edges e_q, e_{q+1}, with labels i_q, i_{q+1}, respectively, either leave a common vertex and satisfy $i_q < i_{q+1}$, or enter a common vertex and satisfy $i_q > i_{q+1}$.

We denote the sequence of edges of K_h in this order by E_h. Also we denote the sequence of vertices of K_h occurring in level $h-1$ (level h) and ordered from left to right by V_h^{low} (resp. V_h^{up}).

Recall that the common vertices of two neighboring principal trees K_h, K_{h+1} are called critical vertices in level h. Let U_h denote the sequence of these vertices ordered from left to right:

$$U_h := V_h^{\mathrm{up}} \cap V_{h+1}^{\mathrm{low}}.$$

The picture below illustrates an example of neighboring principal trees K_h, K_{h+1}; here the critical vertices in level h are indicated by circles.

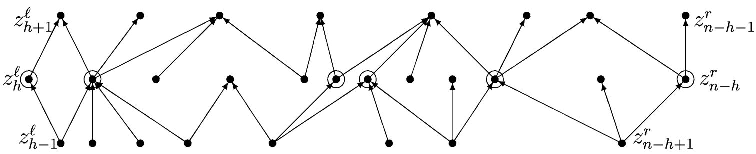

We need to explore the structure of G_T and Γ in a neighborhood of level h in more details. For vertices x, y of K_h, let $P_h(x, y)$ denote the (unique) path from x to y in K_h; in this path the vertices in levels $h-1$ and h alternate. When two consecutive edges of $P_h(x, y)$ enter their common vertex, say, w (lying in level h), we call w a $\wedge$-*vertex* in this path; otherwise (when e, e' leave w) we call w a $\vee$-*vertex*. Also we denote by $K_h(x, y)$ the minimal subtree of K_h containing x, y and all edges incident to *intermediate* vertices of $P_h(x, y)$.

Consider two consecutive critical vertices u, v in level h, where v is the immediate successor of u in U_h. Then the subtrees $K_h(u, v)$ and $K_{h+1}(u, v)$ intersect exactly at the vertices u, v. In particular, the concatenation of $P_{h+1}(u, v)$ and $P_h^{\mathrm{rev}}(u, v)$ forms a simple cycle, denoted by $C(u, v) = C_h(u, v)$, in the graph G^{fw} induced by the fully white edges (forming the set E^{asc}). Define:

$\Omega(u, v) = \Omega_h(u, v)$ to be the closed region in Z surrounded by $C(u, v)$;
$\Omega^*(u, v) = \Omega^*_h(u, v)$ to be the closed region in the disc D_T surrounded by $\sigma(C(u, v))$;
$T(u, v) = T_h(u, v)$ to be the set of tiles $\tau \in T$ such that $\sigma(\tau)$ lies in $\Omega^*(u, v)$.

(For example, in the graph G_T drawn in Figs. 1, 2, the vertices 12 and 24 are consecutive critical vertices in level 2 and the cycle $C(12, 24)$ passes 12, 123, 23, 234, 24, 4, 14, 1, 12.) Clearly each tile in T belongs to exactly one set $T_h(u, v)$.

Let $\mathcal{C}$ be the set of cycles $C_h(u, v)$ among all h, u, v as above, and let $\tilde{G}$ be the subgraph of G^{fw} that is the union of these cycles; this subgraph has planar embedding in Z because G^{fw} does so. Observe that each boundary edge of Z belongs to exactly one cycle in $\mathcal{C}$ and that any other edge of $\tilde{G}$ belongs to exactly two such cycles. It follows that

(6.2) the regions $\Omega(\cdot,\cdot)$ give a subdivision of Z and are exactly the faces of the graph $\tilde{G}$; similarly, the regions $\Omega^*(\cdot,\cdot)$ give a subdivision of D_T and are the faces of $\sigma(\tilde{G})$; the face structures of the planar graphs $\tilde{G}$ and $\sigma(\tilde{G})$ are isomorphic (more precisely, the restriction of σ to $\tilde{G}$ can be extended to a homeomorphism of Z to D_T which maps each $\Omega_h(u,v)$ onto $\Omega_h^*(u,v)$).

Hereinafter, speaking of a face of a planar graph, we mean an inner (bounded) face.

Any vertex v of a cycle $C(u,v)=C_h(u,v)$ belongs to level $h-1$, h or $h+1$, and we call v a *peak* in $C(u,v)$ if it has height $\ne n$, i.e. when v is either a $\vee$-vertex of $P_h(u,v)$ or a $\wedge$-vertex of $P_{h+1}(u,v)$. Also we distinguish between two sorts of edges e in $(K_h(u,v)\cup K_{h+1}(u,v))-C(u,v)$, by saying that e is an *inward pendant* edge w.r.t. $C(u,v)$ if it lies in $\Omega(u,v)$, and an *outward pendant* edge otherwise. (In fact, outward pendant edges will not be important for us in the following.) See the picture where the peaks are indicated by symbol $\otimes$, the inward pendant edges by ι, and the outward pendant edges by o.

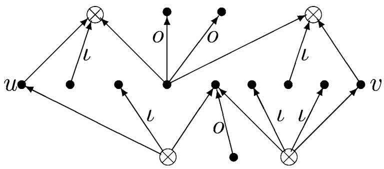

Clearly each edge in $G^{\mathrm{fw}}-\tilde{G}$ is an inward pendant edge of exactly one cycle in $\mathcal{C}$.

Let $\mathfrak{S}(u,v)$ be the set of nonterminal vertices x such that $\sigma(x)\in\Omega^*(u,v)$ and x *is not a peak in* $C(u,v)$. The next lemma describes a number of important properties. Hereinafter Γ^h denotes the subgraph of Γ induced by the (horizontal) edges in level h.

Lemma 6.3 *For* h,u,v *as above*:

(i) *the fully white edges* e *such that* $e\notin C(u,v)$ *and* $\sigma(e)\subset\Omega^*(u,v)$ *are exactly the inward pendant edges for* $C(u,v)$;
(ii) *all tiles in* $T(u,v)$ *are of the same height* h.
(iii) $\mathfrak{S}(u,v)$ *is exactly the set of vertices that are contained in directed paths from* u *to* v *in* Γ^h.

Proof Let Q be the graph whose vertices are the tiles in T and whose edges correspond to the pairs τ,τ' of tiles that have a common edge not in $\tilde{G}$. One can see that any two tiles in the same set $T_h(u,v)$ are connected by a path in Q. On the other hand, if two tiles occur in different sets $T_h(u,v)$ and $T_{h'}(u',v')$, then these tiles cannot be connected by a path in Q. Therefore, the connected components of Q correspond to the sets $T_h(u,v)$. Considering a pair e,e' of consecutive edges in a cycle $C(u,v)$ (which are fully white) and applying Corollary 3.4 to the common vertex w of e,e', we observe that the set $F_T(w)$ of tiles at w is partitioned into two subsets $F^1(w),F^2(w)$ such that: (a) the interior of each tile τ in $F^1(w)$ meets $\Omega(u,v)$ (in particular, τ lies between e and e' when w is a peak in $C(u,v)$), whereas the interior of each tile in $F^2(w)$ is disjoint from $\Omega(u,v)$; (b) the tiles in $F^1(w)$ belong to a path in Q; and (c) each inward pendant edge at w w.r.t. $C(u,v)$ (if any) belongs to some tile in $F^1(w)$. It follows that all tiles in the set $\mathcal{F} := \bigcup(F^1(w) \colon w \in C(u,v))$ belong to the same component of Q. Then all squares in $\sigma(\mathcal{F})$ are contained in one face of $\sigma(\tilde{G})$, and at the same time, they cover the cycle $\sigma(C(u,v))$. This is possible only if $\mathcal{F} \subseteq T(u,v)$. Now (i) easily follows.

Next, observe that for any vertex w of $C(u,v)$, the set $F^1(w)$ as above contains a tile of height h. Therefore, in order to obtain (ii), it suffices to show that any two tiles $\tau, \tau' \in T(u,v)$ sharing an edge e have the same height. This is obvious when τ, τ' have either the same top or the same bottom (in particular, when one of these tiles is black). Suppose this is not the case. Then both tiles are white, and the edge e connects the left or right vertex of one of them to the right or left vertex of the other. So both ends of e are nonterminal and e is fully white. By (i), e is an inward pendant edge for $C(u,v)$ and one end x of e is a peak. Therefore, both τ, τ' belong to the set $F^1(x)$ and lie between the two edges of $C(u,v)$ incident to x. But both edges either enter x or leave x (since x is a peak), implying that x must be either the top or the bottom of both τ, τ'; a contradiction. Thus, (ii) is valid.

Finally, to see (iii), consider a vertex $x \in \mathfrak{S}(u,v)$ different from u. Then x is not in lbd(Z); for otherwise we would have $u = z^\ell_h$ and $x \in \{z^\ell_{h-1}, z^\ell_{h+1}\}$, implying that x is a peak in $C(u,v)$. Therefore (cf. Corollary 3.4(iv)), there exists a white tile τ such that $r(\tau) = x$. Both x, τ have the same height. Suppose that $\tau \notin T(u,v)$. Then $x \in C(u,v)$ and x is of height h (since x is not a peak in $C(u,v)$). So τ is of height h as well, and in view of (ii), τ belongs to some collection $T_h(u',v') \ne T_h(u,v)$. One can see that the latter is possible only if $v' = u$, implying $x = u$; a contradiction.

Hence, τ belongs to $T(u,v)$ and has height h. Take the vertex $y := \ell(\tau)$. Then y is nonterminal, $\sigma(y) \in \Omega^*(u,v)$, and there is a horizontal edge from y to x in Γ. Also y is not a peak in $C(u,v)$ (since y is of height h). So $y \in \mathfrak{S}(u,v)$. Apply a similar procedure to y, and so on. Eventually, we reach the vertex u, obtaining a directed path in Γ^h going from u to the initial vertex x. A directed path from x to v is constructed in a similar way.

Conversely, let P be a directed path from u to v in Γ^h. The fact that all vertices of P belong to $\mathfrak{S}(u,v)$ is easily shown by considering the sequence of white tiles corresponding to the edges of P and using the fact that all these tiles have height h.

Thus, (iii) is valid and the lemma is proven. □

Let the sequence U_h consist of (critical) vertices $u_0 = z^\ell_h, u_1, \dots, u_{r-1}, u_r = z^r_{n-h}$. We abbreviate $T_h(u_{p-1}, u_p)$ as $T_h(p)$, and denote by $G^h(p)$ the subgraph of G_T whose image under σ lies in $\Omega^*_h(u_{p-1}, u_p)$. By (ii) in Lemma 6.3, $T_h(1), \dots, T_h(r)$ give a partition of the set of tiles of height h.

In its turn, (iii) in this lemma shows that the graph Γ^h is represented as the concatenation of $\Gamma^h(1), \dots, \Gamma^h(r)$, where each graph $\Gamma^h(p)$ is the union of (horizontal) directed paths from u_{p-1} to u_p in Γ. We refer to $\Gamma^h(p)$ as p-th *hammock of* Γ^h (or in level h) *beginning at* u_{p-1} *and ending at* u_p, and similarly for the subgraph $\sigma(\Gamma^h(p))$ of $\sigma(\Gamma^h)$. The fact that $\sigma(\Gamma^h(p))$ is the union of directed paths from $\sigma(u_{p-1})$ to $\sigma(u_p)$ easily implies that

(6.3) the boundary of each face F of the hammock $\sigma(\Gamma^h(p))$ is formed by two directed paths with the same beginning x and the same end y;

we say that the face F *begins at x and ends at y*. The *extended hammock* $\bar{\Gamma}^h(p)$ is constructed by adding to Γ^h the cycle $C(u_{p-1}, u_p)$ and the inward pendant edges for it (all added edges are ascending in Γ); this is just the subgraph of Γ whose image under σ is contained in $\Omega^*(u_{p-1}, u_p)$.

Applying Lemma 5.2 to the planar graph $\sigma(\Gamma^h)$, we obtain that

(6.4) the partial order $(\mathfrak{S}^h, \prec_{\Gamma^h})$, where $\mathfrak{S}^h := \{X \in \mathfrak{S}_T : |X| = h\}$, is a lattice with the minimal element z_h^ℓ and the maximal element z_{n-h}^r; similarly, for each $p = 1, \dots, r$, $(\mathfrak{S}(u_{p-1}, u_p), \prec_{\Gamma^h(p)})$ is a lattice with the minimal element u_{p-1} and the maximal element u_p.

An example of $\sigma(\Gamma^h)$ with $r = 5$ is drawn in the picture; here all edges are directed from left to right.

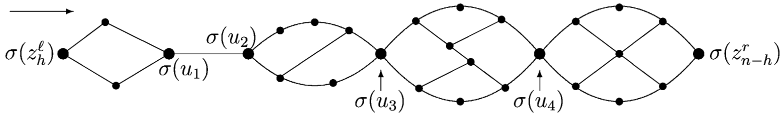

We call a hammock $\Gamma^h(p)$ *trivial* if it has only one edge (which goes from u_{p-1} to u_p). In this case $T_h(p)$ consists of a single white tile τ such that both $b(\tau), t(\tau)$ are nonterminal, $\ell(\tau) = u_{p-1}$ and $r(\tau) = u_p$ (then $\bar{\Gamma}^h(p)$ is formed by the four edges of τ and the horizontal edge from $\ell(\tau)$ to $r(\tau)$).

6.2 Nontrivial hammocks

Next we describe the structure of a *nontrivial* hammock $\Gamma^h(p)$. For a white tile in $T_h(p)$ (which, obviously, exists), at least one of its bottom and top vertices is terminal (for otherwise all edges of this tile are fully white, implying that its left and right vertices are critical). Then $|T_h(p)| \ge 2$ and the set $T_h^b(p)$ of black tiles in $T_h(p)$ is nonempty. We are going to show a one-to-one correspondence between the black tiles and the faces of $\sigma(\Gamma^h(p))$.

Given a black tile $\tau \in T_h^b(p)$, consider the sequence $x_0, \dots, x_k$ of the end vertices of the edges $e_0, \dots, e_k$ leaving $b(\tau)$ and ordered from left to right (i.e. by increasing their labels), and the sequence $y_0, \dots, y_{k'}$ of the beginning vertices of the edges $e'_0, \dots, e'_{k'}$ entering $t(\tau)$ and ordered from left to right. Then $e_0, e_k, e'_0, e'_{k'}$ are the edges of τ, the other edges $e_q, e'_{q'}$ are semi-white, $x_0 = y_0 = \ell(\tau)$ and $x_k = y_{k'} = r(\tau)$. Also each pair e_{q-1}, e_q belongs to a white tile τ_q, each pair $e'_{q'-1}, e'_{q'}$ belongs to a white tile $\tau'_{q'}$, and there are no other tiles having a vertex at $b(\tau)$ or $t(\tau)$, except for τ. For two consecutive tiles τ_q, τ_{q+1} (resp. $\tau'_{q'}, \tau'_{q'+1}$), we have $r(\tau_q) = \ell(\tau_{q+1}) = x_q$ (resp. $r(\tau'_{q'}) = \ell(\tau'_{q'+1}) = y_{q'}$). Therefore, the sequence $(x_0, \dots, x_k)$ gives a directed path in Γ, denoted by γ_τ, and similarly, $(y_0, \dots, y_{k'})$ gives a directed path in Γ, denoted by β_τ. Both paths go from $\ell(\tau)$ to $r(\tau)$ and have no other common vertices (since $x_q = y_{q'}$ for some $0 < q < k$ and $0 < q' < k'$ would lead to a contradiction with (3.3)).

We denote $\beta_\tau \cup \gamma_\tau$ by ζ_τ, regarding it both as a graph and as the simple cycle in which the edges of γ_τ are forward. The closed region in D_T surrounded by $\sigma(\zeta_\tau)$ (which is a disc) is denoted by ρ_τ. We call β_τ and γ_τ the *lower* and *upper* paths in

ζ_τ, respectively, and similarly for the paths $\sigma(\gamma_\tau)$ and $\sigma(\beta_\tau)$ in $\sigma(\zeta_\tau)$ (a motivation will be clearer later).

Any white tile in $T_h(p)$ has the bottom or top in common with some black tile. This implies that the graph $\Gamma^h(p)$ is exactly the union of cycles ζ_τ over $\tau \in T_h^b(p)$. Moreover, each edge e of $\Gamma^h(p)$ with $\sigma(e)$ not in the boundary of $\sigma(\Gamma^h(p))$ belongs to two cycles as above (since such an e is the diagonal of a white tile in which both the bottom and top vertices are terminal). These facts are strengthened as follows.

Lemma 6.4 *The regions ρ_τ, $\tau \in T_h^b(p)$, are exactly the faces of the graph $\sigma(\Gamma^h(p))$.*

Proof For such a τ, form the region R_τ in D_T to be the union of the square $\sigma(\tau)$, the triangles (half-squares) with the vertices $\sigma(\ell(\tau')), \sigma(r(\tau')), \sigma(b(\tau'))$ over all white tiles $\tau' \in T$ having the common bottom with τ, and the triangles (half-squares) with the vertices $\sigma(\ell(\tau')), \sigma(r(\tau')), \sigma(t(\tau'))$ over all white tiles $\tau' \in T$ having the common top with τ. One can see that R_τ is a disc and its boundary is just $\sigma(\zeta_\tau)$. So $R_\tau = \rho_\tau$. Obviously, the regions R_τ, $\tau \in T_h^b(p)$, have pairwise disjoint interiors. □

Thus, the faces of $\sigma(\Gamma^h(p))$ are generated by the black tiles in $T_h(p)$; each face ρ_τ contains $\sigma(\tau)$, begins at $\ell(\tau)$ and ends at $r(\tau)$. Figure 3 illustrates an example with two black tiles: the subgraph $G^h(p)$ of G_T and the extended hammock for it.

A further refinement shows that the pairwise intersections of cycles ζ_τ are poor.

Lemma 6.5 *For distinct $\tau, \tau' \in T_h^b(p)$, let $\zeta_\tau \cap \zeta_{\tau'} \neq \emptyset$. Then the intersection of these cycles is contained in the upper path of one of them and in the lower path of the other, and it consists of a single vertex or a single edge.*

Proof Suppose that γ_τ and $\gamma_{\tau'}$ have a common vertex w and this vertex is intermediate in both paths. Assuming that the label of the edge $(b(\tau), w)$ is smaller than the label of the edge $(b(\tau'), w)$, take the tile $\tilde\tau \in F_T(b(\tau))$ with $r(\tilde\tau) = w$ and the tile $\tilde\tau' \in F_T(b(\tau'))$ with $\ell(\tilde\tau') = w$ (which exist since $w \neq \ell(\tau)$ and $w \neq r(\tau')$). Then $\tilde\tau$ and $\tilde\tau'$ overlap, contrary to Corollary 3.4(c). Similarly, β_τ and $\beta_{\tau'}$ cannot intersect at an intermediate vertex of both paths.

Now suppose that $\gamma_\tau \cap \beta_{\tau'}$ contains two vertices w, w'. Then $w, w', b(\tau), t(\tau')$ are connected by the four edges $e := (b(\tau), w)$, $e' := (b(\tau), w')$, $f := (w, t(\tau'))$, and

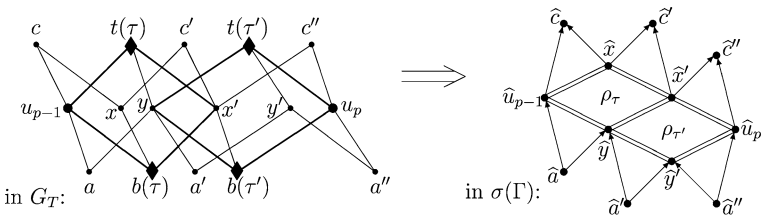

Fig. 3 *On the left*: an instance of $G^h(p)$. Here $T_h(p)$ consists of two black tiles τ, τ' and seven white tiles. The corresponding cycle $C(u_{p-1}, u_p)$ contains all nonterminal vertices (there is no inward pendant edge). *On the right*: the extended hammock $\sigma(\bar\Gamma^h(p))$. Here the hammock $\sigma(\Gamma^h(p))$ is indicated by double lines (the edges should be directed to the right), and for a nonterminal vertex $*$ of G_T, we write $\hat{*}$ for $\sigma(*)$

$f' := (w', t(\tau'))$. By (3.3), these edges are spanned by a tile $\tilde{\tau}$. We assert that $\tilde{\tau}$ is a white tile in T (whence w, w' are connected by a horizontal edge in Γ).

Indeed, if this is not so, then the edges e, e' are not consecutive in $E_T(b(\tau))$ and f, f' are not consecutive in $E_T(t(\tau'))$. Note that at least one of the edges e, e', say, e, is white (for otherwise $\tilde{\tau} = \tau$, implying that the black tiles τ, τ' have a common terminal vertex, namely, $t(\tau) = t(\tau')$). Since f, f are not consecutive, there is a (unique) white tile $\hat{\tau} \in F_T(t(\tau'))$ lying between f and f' and containing f but not f'. Since e, f' are parallel, the white edge e lies between f and the edge of $\hat{\tau}$ connecting $b(\hat{\tau})$ and w (these three edges are incident to w). This leads to a contradiction with Corollary 3.4.

Thus, $\tilde{\tau}$ is a white tile in T, and now the lemma easily follows. □

6.3 Contractions on Γ

Our final step in preparation to proving Proposition 6.2 is to explain how the graph Γ changes under the n- and 1-contraction operations. We use terminology and notation from Subsect. 3.3. A majority of our analysis is devoted to the n-contraction operation that reduces a g-tiling T on $Z = Z_n$ to the g-tiling $T' := T/n$ on Z_{n-1}.

It is more convenient to consider the reversed n-strip $Q = (e_0, \tau_1, e_1, \dots, \tau_r, e_r)$, i.e. e_0 is the n-edge $z^r_n z^r_{n-1}$ on rbd(Z) and e_r is the n-edge $z^\ell_{n-1} z^\ell_n$ on lbd(Z). Let $R_Q = (v_0, a_1, v_1, \dots, a_r, v_r)$ be the right boundary, and $L_Q = (v'_0, a'_1, v'_1, \dots, a'_r, v'_r)$ the left boundary of Q, i.e. $v'_0 = z_0$, $v_r = z_n$, and $e_q = (v'_q, v_q)$ for each q.

Since n is the maximal label, if an n-edge e belongs to a tile τ, then e is either br(τ) or lt(τ). This implies (in view of $e_0 = \mathrm{br}(\tau_1)$) that for consecutive tiles τ_q, τ_{q+1} in Q, one holds: if both tiles are white then $e_q = \mathrm{lt}(\tau_q) = \mathrm{br}(\tau_{q+1})$; if τ_q is black then $e_q = \mathrm{br}(\tau_q) = \mathrm{br}(\tau_{q+1})$; and if τ_{q+1} is black then $e_q = \mathrm{lt}(\tau_q) = \mathrm{lt}(\tau_{q+1})$. So the height of τ_{q+1} is greater by one than the height of τ_q if both tiles are white, and the heights are equal otherwise. In particular, the tile height is weakly increasing along Q and grows from 1 to $n-1$. For $h = 1, \dots, n-1$, let $Q^h = (e_{d(h)-1}, \tau_{d(h)}, e_{d(h)}, \dots, \tau_{f(h)}, e_{f(h)})$ be the maximal part of Q in which all tiles are of height h; we call it h-th *fragment* of Q.

Recall that from the viewpoint of D_T, the n-contraction operation acts as follows. The interior of $\sigma(Q)$ is removed from D_T, forming two closed simply connected regions D^r, D^ℓ, where D^r (the "right" region) contains $\sigma(R_Q)$ and the rest of $\sigma(\mathrm{rbd}(Z))$, and D^ℓ (the "left" region) contains $\sigma(L_Q)$ and the rest of $\sigma(\mathrm{lbd}(Z))$. The region D^r is shifted by $-\epsilon_n$ (where ϵ_n is n-th unit base vector in $\mathbb{R}^{[n]}$) and the path $\sigma(R_Q) - \epsilon_n$ merges with $\sigma(L_Q)$. From the viewpoint of Z, the tiles occurring in Q vanish and the tiles $\tau \in T$ with $\sigma(\tau) \subset D^r$ are shifted by $-\xi_n$; in other words, each vertex X (regarded as a set) containing the element n turns into the vertex $X - n$ of the resulting tiling T' on Z_{n-1}. (Recall that X contains n if and only if $\sigma(X)$ is in D^r.) Each vertex v_q of R_Q merges with the vertex v'_q of L_Q. The path L_Q no longer contains terminal vertices (so all edges in it are now fully white) and becomes an n-legal path for T'. Any zigzag subpath in this path goes from left to right, and

(6.5) for $h = 1, \dots, n-2$, $v'_{f(h)=d(h+1)-1}$ is the critical vertex of L_Q in level h for T'.

Consider h-th fragment $Q^h = (e_{d-1}, \tau_d, e_d, \ldots, \tau_f, e_f)$ (letting $d := d(h)$ and $f := f(h)$). It produces $(f-d)/2+1$ horizontal and four ascending edges in $\Gamma = \Gamma_T$ (note that $f-d$ is even). More precisely, each tile τ_q with $q-d$ even is white and its diagonal makes the horizontal edge $g_q := (v'_q, v_{q-1})$ in Γ. Also τ_d contributes the ascending edges $e_{d-1} = \mathrm{br}(\tau_d) = (v'_{d-1}, v_{d-1})$ and $a'_d = \mathrm{bl}(\tau_q) = (v'_{d-1}, v'_d)$. In its turn, τ_f contributes the ascending edges $e_f = \mathrm{lt}(\tau_f) = (v'_f, v_f)$ and $a_f = \mathrm{rt}(\tau_f) = (v_{f-1}, v_f)$ to Γ. Let $\mathcal{E}^h$ be the set of edges in Γ produced by Q^h. Then $\mathcal{E}^h \cap \mathcal{E}^{h+1} = \{e_f\}$ $(= \{e_{d(h+1)-1}\})$.

Under the n-contraction operation, Γ is transformed into the graph $\Gamma' := \Gamma_{T'}$. The transformation concerns only the sets $\mathcal{E}^h$ and is obvious: all horizontal edges of $\mathcal{E}^h$ disappear (as all tiles in Q^h vanish) and the four ascending edges are replaced by (the edges of) the subpath L_{Q^h} of L_Q from v'_{d-1} to v'_f, in which all edges connect levels $h-1$ and h (using indices as above). In particular, when $d = f$ (i.e. when Q^h has only one (white) tile), the five edges of $\mathcal{E}^h$ shrink into one edge $a'_d = (v'_{d-1}, v'_d)$.

When $\Delta := f - d > 0$, the transformation needs to be examined more carefully. The $\Delta/2+1$ white tiles and the $\Delta/2$ black tiles in Q^h alternate. The horizontal edges $g_d, g_{d+2}, \ldots, g_f$ in $\mathcal{E}^h$ belong to the same nontrivial hammock in Γ^h, say, $\Gamma^h(p)$. More precisely, for $q = d+1, d+3, \ldots, f-1$, the edges $g_{q-1} = (v'_{q-1}, v_{q-2})$ and $g_{q+1} = (v'_{q+1}, v_q)$ are contained in the cycle ζ_{τ_q} related to the black tile τ_q. Also τ_{q-1} is the *leftmost* white tile in $F_T(t(\tau_q))$ and τ_{q+1} is the *rightmost* white tile in $F_T(b(\tau_q))$. Therefore, g_{q-1} is the *first* edge in the lower path β_{τ_q} and g_{q+1} is the *last* edge in the upper path γ_{τ_q} in ζ_{τ_q}. For a similar reason, unless $q = f-1$, g_{q+1} is simultaneously the first edge in $\beta_{\tau_{q+2}}$.

This and Lemma 6.5 imply that if we take the union of cycles ζ_{τ_q} for $q = d+1$, $d+3, \ldots, f-1$ and delete from it the horizontal edges in $\mathcal{E}^h$, then we obtain two directed horizontal paths in level h of Γ: path $\mathcal{P}_1^h$ from v'_d to v'_f which passes the vertices $v'_d, v'_{d+2}, \ldots, v'_f$ in this order, and path $\mathcal{P}_2^h$ from v_{d-1} to v_{f-1} which passes the vertices $v_{d-1}, v_{d+1}, \ldots, v_{f-1}$ in this order (these paths may contain other vertices as well). When $f = d$, each of these paths consists of a single vertex.

In the new graph Γ', the path $\mathcal{P}_1^h$ preserves and continues to be a horizontal path in level h, whereas $\mathcal{P}_2^h$ is shifted by $-\xi_n$ and turns into the directed horizontal path, denoted by $\mathcal{P}_2'^h$, that passes the vertices $v'_{d-1}, v'_{d+1}, \ldots, v'_{f-1}$ in level $h-1$. These paths are connected in Γ' by the (zigzag) path $\mathcal{Z}^h := (v'_{d-1}, v'_d, v'_{d+1}, \ldots, v'_f)$ whose edges connect levels $h-1$ and h. (Under the transformation, the hammock $\Gamma^h(p)$ becomes split into two hammocks, one (in level h) containing the path $\mathcal{P}_1^h$, and the other (in level $h-1$) containing the image $\mathcal{P}_2'^h$ of $\mathcal{P}_2^h$.) In view of (6.5),

(6.6) the last vertex v'_f of $\mathcal{P}_1^h$ (which is simultaneously the first vertex of $\mathcal{P}_2'^{(n+1)}$ when $h < n-1$) and the first vertex v'_{d-1} of $\mathcal{P}_2'^h$ (which is simultaneously the last vertex of $\mathcal{P}_1^{n-1}$ when $h > 1$) are critical for T' in levels h and $h-1$, respectively.

The transformation of $\sigma(\Gamma)$ into $\sigma(\Gamma')$ within the fragment Q^h with $d < f$ is illustrated in the picture; here for brevity we write ρ_q instead of ρ_{τ_q}, and omit σ in notation for vertices, edges and paths.

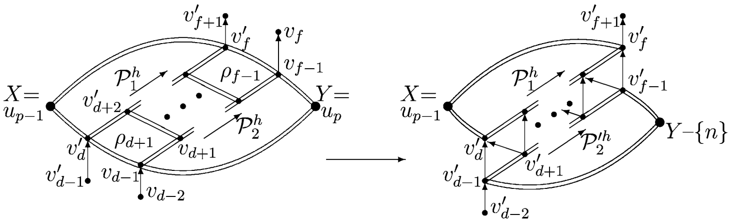

Notice that the paths $\mathcal{P}_1^{h-1}$ and $\mathcal{P}_1^h$ are connected in Γ by one ascending edge (namely, $a'_{d(h)}$) going from the end $v'_{f(h-1)=d(h)-1}$ of the former path to the beginning $v'_{d(h)}$ of the latter one; we call it the *bridge* between these paths and denote by b'_h. Similarly, there is an ascending edge (namely, $a_{d(h)-1}$) going from the end $v_{d(h)-2}$ of $\mathcal{P}_2^{h-1}$ to the beginning $v_{d(h)-1}$ of $\mathcal{P}_2^h$, the bridge between these paths, denoted by b_h. Under the transformation, b_h is shifted and becomes the last edge $(v'_{d(h)-2}, v'_{d(h)})$ of the (zigzag) path $\mathcal{Z}^{h-1}$ and the beginning of $\mathcal{P}_2^{\prime h}$ merges with the end of $\mathcal{P}_1^{h-1}$. (Cf. (6.6).)

Thus, concatenating the paths $\mathcal{P}_1^1, \dots, \mathcal{P}_1^{n-1}$ and the bridges $b'_1, b'_2, \dots, b'_{n-1}$ (where $b'_1 := a'_1$), we obtain a directed path from z_0 to z_{n-1}^ℓ in both Γ and Γ', denoted by $\mathcal{P}_1$. Accordingly, we construct directed paths $\mathcal{P}_2$ (in Γ) and $\mathcal{P}'_2$ (in Γ') by concatenating, in a due way, the paths $\mathcal{P}_2^h$ with the bridges b_h, and the paths $\mathcal{P}_2^{\prime h}$ with the shifts of these bridges, respectively.

Let Γ_1 and Γ_2 be the subgraphs of Γ whose images under σ lie in the regions D^ℓ and D^r of D_T, respectively. Let Γ'_2 be the subgraph of Γ' whose image under σ lies in $D^r - \epsilon_n$. We observe that:

(6.7) (i) the common vertices of $\mathcal{P}_1$ and $\mathcal{P}'_2$ are exactly the critical vertices (indicated in (6.5)) of the path L_Q in $G_{T'}$;
(ii) if a directed path P in Γ' goes from a vertex x of Γ_1 to a vertex y of Γ'_2, then P contains a critical vertex v of L_Q; moreover, there exist a directed path P' from x to v in Γ_1 and a directed path P'' from v to y in Γ'_2.

Here (ii) follows from the facts that both $\mathcal{P}_1, \mathcal{P}'_2$ are directed paths and that all edges not in $\mathcal{P}_1 \cap \mathcal{P}'_2$ that connect Γ_1 and Γ'_2 go from vertices of $\mathcal{P}'_2$ to vertices of $\mathcal{P}_1$ (as they are ascending edges occurring in zigzag paths $\mathcal{Z}^h$).

The 1-contraction operation acts symmetrically, in a sense; below we mainly describe the moments where there are differences from the n-contraction case.

Let $Q = (e_0, \tau_1, e_1, \dots, \tau_r, e_r)$ be the 1-strip in T, $R_Q = (v_0, a_1, v_1, \dots, a_r, v_r)$ the right boundary of Q, and $L_Q = (v'_0, a'_1, v'_1, \dots, a'_r, v'_r)$ the left boundary of Q. So $e_0 = (v_0, v'_0) = z_0 z_1^\ell$ and $e_r = (v_r, v'_r) = z_{n-1}^r z_n$.

Since label 1 is minimal, if a 1-edge e belongs to a tile $\tau \in T$, then either $e = \mathrm{bl}(\tau)$ or $e = \mathrm{rt}(\tau)$. For consecutive tiles τ_q, τ_{q+1} in Q: the height of τ_{q+1} is greater by one than the height of τ_q if both tiles are white, and the heights are equal if one of these tiles is black. Like the previous case, for $h = 1, \dots, n$, define h-th fragment of Q to be the maximal part $Q^h = (e_{d(h)-1}, \tau_{d(h)}, e_{d(h)}, \dots, \tau_{f(h)}, e_{f(h)})$ of Q in which all tiles are of height h.

The fragment Q^h produces $(f-d)/2$ horizontal and four ascending edges in Γ, where $d := d(h)$ and $f := f(h)$. Each tile τ_q with $q-d$ even is white and produces the horizontal edge $g_q := (v'_{q-1}, v_q)$ in Γ^h. Also τ_d contributes the ascending edges $e_{d-1} = \mathrm{bl}(\tau_d) = (v_{d-1}, v'_{d-1})$ and $a_d = \mathrm{br}(\tau_q) = (v_{d-1}, v_d)$, and τ_f contributes the ascending edges $e_f = \mathrm{rt}(\tau_f) = (v_f, v'_f)$ and $a'_f = \mathrm{lt}(\tau_f) = (v'_{f-1}, v'_f)$. Let $\mathcal{E}^h$ be the set of edges in Γ produced by Q^h.

Under the 1-contraction operation, Γ is transformed into $\tilde{\Gamma}' := \Gamma_{T/1}$ as follows: for each h, the horizontal edges of $\mathcal{E}^h$ disappear and the four ascending edges are replaced by the subpath R^h_Q of R_Q from v_{d-1} to v_f (using indices as above) in which all edges connect levels $h-1$ and h. When $d = f$, $\mathcal{E}^h$ shrinks into one edge $a_d = (v_{d-1}, v_d)$.

When $\Delta := f - d > 0$, the horizontal edges $g_d, g_{d+2}, \dots, g_f$ in $\mathcal{E}^h$ belong to the same nontrivial hammock, $\Gamma^h(p)$ say. On the other hand, for $q = d+1, d+3, \dots, f-1$, the edges g_{q-1} and g_{q+1} belong to the cycle ζ_{τ_q} related to the black tile τ_q. Also τ_{q-1} is the *rightmost* white tile in $F_T(t(\tau_q))$ and τ_{q+1} is the *leftmost* white tile in $F_T(b(\tau_q))$. Therefore, g_{q-1} is the *last* edge in the lower path β_{τ_q} and g_{q+1} is the *first* edge in the upper path γ_{τ_q} in ζ_{τ_q}. Unless $q = f-1$, g_{q+1} is simultaneously the last edge in $\beta_{\tau_{q+2}}$. In view of Lemma 6.5, this implies that if we take the union of cycles ζ_q and then delete from it the horizontal edges of $\mathcal{E}^h$, then we obtain two horizontal paths in $\Gamma^h(p)$: path $\tilde{\mathcal{P}}^h_1$ from v'_{f-1} to v'_{d-1} which passes the vertices $v'_{f-1}, v'_{f-3}, \dots, v'_{d-1}$ in this order, and path $\tilde{\mathcal{P}}^h_2$ from v_f to v_d which passes the vertices $v_f, v_{f-2}, \dots, v_d$ in this order. (So, both paths are directed by decreasing the vertex indices, in contrast to the direction of the corresponding paths in the n-contraction case.)

For our further purposes, it will be sufficient to examine the transformation of Γ only within its part related to a single fragment Q^h. In the new graph $\tilde{\Gamma}'$, the path $\tilde{\mathcal{P}}^h_2$ preserves and continues to be a horizontal path in level h, whereas $\tilde{\mathcal{P}}^h_1$ is shifted by $-\xi_1$ and turns into a horizontal path in level $h-1$, denoted by $\tilde{\mathcal{P}}'^h_1$, which passes the vertices $v_{f-1}, v_{f-3}, \dots, v_{d-1}$. These paths are connected in $\tilde{\Gamma}'$ by the zigzag path $\tilde{\mathcal{Z}}^h = (v_f, v_{f-1}, \dots, v_d, v_{d-1})$ whose edges are ascending and connect levels $h-1$ and h. The picture illustrates the transformation $\sigma(\Gamma) \mapsto \sigma(\tilde{\Gamma}')$ within the fragment Q^h.

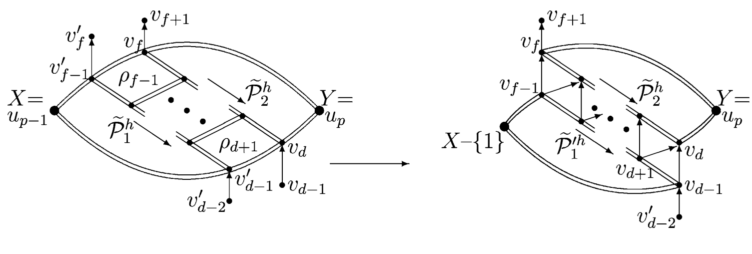

One can see that the first vertex v_f of $\tilde{\mathcal{P}}^h_2$ and the last vertex v_{d-1} of $\tilde{\mathcal{P}}'^h_1$ are critical vertices for $\tilde{T}' := T/1$ in levels h and $h-1$, respectively (and they are the only critical vertices for $\tilde{T}'$ occurring in these paths). Under the transformation, the hammock $\Gamma^h(p)$ (concerning Q^h) becomes split into two hammocks in $\tilde{\Gamma}'$: hammock H_1 in level $h-1$ that contains the path $\tilde{\mathcal{P}}_1'^h$, and hammock H_2 in level h that contains $\tilde{\mathcal{P}}_2^h$. Comparing the part of $\Gamma^h(p)$ between $\tilde{\mathcal{P}}_1^h$ and $\tilde{\mathcal{P}}_2^h$ with the zigzag path $\mathcal{Z}^h$, one can conclude that

(6.8) if a vertex x of H_1 and a vertex y of H_2 are connected in $\tilde{\Gamma}'$ by a directed path from x to y, then there exists a directed path from $x+\xi_1$ to y in Γ.

6.4 Proof of "$\prec^*$ implies $\prec_\Gamma$"

Based on the above explanations, we are now ready to prove Proposition 6.2. We use induction on the number of edge labels of a g-tiling and apply the n- and 1-contraction operations.

Let $A, B \in \mathfrak{S}_T$, $A \neq B$ and $A \prec^* B$. We have to show the existence of a directed path from the vertex A to the vertex B in the graph $\Gamma = \Gamma_T$.

First we consider the case $|A| < |B|$. Let $A' := A - \{n\}$ and $B' := B - \{n\}$. Then $A', B' \in \mathfrak{S}_{T'}$, where T' is the n-contraction T/n of T. Also $|A'| \le |B'|$. Therefore, $A \prec^* B$ implies $A' \prec^* B'$, and by induction the graph $\Gamma' := \Gamma_{T'}$ contains a directed path P from A' to B'. Note that $n \in A$ would imply $n \in B$. So either $n \notin A, B$, or $n \in A, B$, or $n \notin A$ and $n \in B$. Consider these cases, keeping notation from Subsect. 6.3.

Case 1: $n \notin A, B$. Then $A' = A$, $B' = B$, and both A', B' belong to the graph Γ_1. Since the path $\mathcal{P}_1$ in the boundary of Γ_1 is directed, P as above can be chosen so as to be entirely contained in Γ_1. (For if P meets $\mathcal{P}_1$, take the first and last vertices of P that occur in $\mathcal{P}_1$, say, x, y (respectively), and replace in P its subpath from x to y by the subpath of $\mathcal{P}_1$ connecting x and y, which must be directed from x to y since Γ' is acyclic.) Then P is a directed path from A to B in Γ, as required.

Case 2: $n \in A, B$. Then both A', B' belong to the graph Γ_2'. Like the previous case, one may assume that P is entirely contained in Γ_2'. Since $\Gamma_2' + \xi_n$ is a subgraph of Γ, $P + \xi_n$ is the desired path from A to B in Γ.

Case 3: $n \notin A$ and $n \in B$. Then $A' = A$ is in Γ_1 and B' is in Γ_2'. By (6.7), there exist a directed path P' from A to v in Γ_1 and a directed path P'' from v to B' in Γ_2', where v is a critical vertex $v'_{f(h)}$ in $\mathcal{P}_1 \cap \mathcal{P}_2'$. Concatenating P', the (fully white) ascending edge $e_{f(h)} = (v'_{f(h)}, v_{f(h)})$, and the path $P'' + \xi_n$ (going from $v_{f(h)}$ to $B'n = B$), we obtain a directed path from A to B in Γ, as required.

Now consider the case $|A| = |B| =: h$. Note that the reduction by label n as above does not work when the element n is contained in B but not in A (since in this case $|B - \{n\}|$ becomes less than $|A - \{n\}|$ and we cannot apply induction). Nevertheless, we can use the 1-contraction operation (which, in its turn, would not fit to treat the case $|A| < |B|$, since the concatenation of the paths $\tilde{\mathcal{P}}_2^h$ and the corresponding bridges is not a directed path). Let $A' := A - \{1\}$ and $B' := B - \{1\}$; then $|A'| \le |B'|$ (since $1 \in B$ would imply $1 \in A$, in view of $A \lessdot B$). The vertices A, B belong to the horizontal subgraph Γ^h of Γ; suppose they occur in p-th and p'-th hammocks of Γ^h, respectively. The existence of a directed path from A to B is not seen immediately only when $p = p'$ and, moreover, when the 1-strip for T "splits" the hammock $\Gamma^h(p)$. (Note that $p > p'$ is impossible; otherwise there would exist a directed path from B to A, implying $B \lessdot A$.) Let H_1, H_2 be the hammocks in $\tilde{\Gamma}'$ created from $\Gamma^h(p)$ by the 1-contraction operation as described above. If both A', B' belong to the same H_i, then the existence (by induction) of a directed path from A' to B' in $\tilde{\Gamma}'$ (and therefore, in H_i) immediately yields the result. Let $A' \in H_1$ and $B' \in H_2$ (the case $A' \in H_2$ and $B' \in H_1$ is impossible). Then the existence of a directed path from A to B in Γ follows from (6.8).

Thus, $A \prec_\Gamma B$ is valid in all cases, as required. This completes the proof of Proposition 6.2, yielding Theorems 6.1 and 4.1 and completing the proof of Theorem B.

7 Additional results and a generalization

In this concluding section we gather in an additional harvest from results and methods described in previous sections; in particular, we generalize Theorem A to the case of two permutations. Sometimes the description below will be given in a sketched form and we leave the details to the reader.

We start with associating to a permutation ω on $[n]$ the directed path P_ω on the zonogon Z_n in which the vertices are the points $v^i_\omega := \sum(\xi_j : j \in \omega^{-1}[i])$, $i = 0, \dots, n$, and the edges are the directed line segments e^i_ω from v^{i-1}_ω to v^i_ω. So P_ω begins at $v^0_\omega = z_0$, ends at $v^n_\omega = z_n$, and each edge e^i_ω is a parallel translation of the vector $\xi_{\omega^{-1}(i)}$. Also a vertex v^i_ω represents i-th ideal $I^i_\omega = \omega^{-1}[i]$ for ω (cf. Sect. 2). Note that if the spectrum of a g-tiling T on Z_n contains all sets I^i_ω, then the graph G_T contains the path P_ω, in view of Proposition 3.11(i). When ω is the longest permutation ω_0 on $[n]$, P_ω becomes the right boundary $\mathrm{rbd}(Z_n)$ of Z_n. When ω is the identical permutation, denoted as id, P_ω becomes the left boundary $\mathrm{lbd}(Z_n)$.

Consider two permutations ω', ω on $[n]$ and assume that the pair (ω', ω) satisfies the condition:

(7.1) for any $i, j \in [n]$, either $I^i_{\omega'} \lessdot I^j_\omega$ or $I^i_{\omega'} \supseteq I^j_\omega$.

In particular, this implies that P_ω lies to the right of $P_{\omega'}$, i.e. each point v^i_ω lies to the right of $v^i_{\omega'}$ in the corresponding horizontal line, with possibly $v^i_\omega = v^i_{\omega'}$.

For the closed region $Z(\omega', \omega)$ bounded by the paths $P_{\omega'}$ (the left boundary) and P_ω (the right boundary), we can consider a g-tiling T. It is defined by axioms (T2), (T3) as before and slightly modified axioms (T1), (T4), where (cf. Subsect. 3.1): in (T1), the first condition is replaced by the requirement that each edge in $(P_{\omega'} \cup P_\omega) - (P_{\omega'} \cap P_\omega)$ belong to exactly one tile; and in (T4), it is now required that $D_T \cup \sigma(P_{\omega'} \cap P_\omega)$ be simply connected, where, as before, D_T denotes $\bigcup(\sigma(\tau) : \tau \in T)$. Also one should include in the graph $G_T = (V_T, E_T)$ all common vertices and edges of $P_{\omega'}, P_\omega$. Note that such a T possesses the following properties:

(7.2) (i) the union of tiles in T and the edges in $P_{\omega'} \cap P_\omega$ is exactly $Z(\omega', \omega)$; and
(ii) all vertices in $\mathrm{bd}(Z(\omega', \omega)) = P_{\omega'} \cup P_\omega$ are nonterminal.

This is seen as follows. Let an edge e of height h belong to two tiles $\tau, \tau' \in T$ (where the height of an edge is the half-sum of the heights of its ends). Suppose $\tau \cup \tau'$ contains no edge of height h lying to the right of e. Then one of these tiles, say, τ, is black and either $e = \mathrm{br}(\tau)$ or $e = \mathrm{rt}(\tau)$. Assuming $e = \mathrm{br}(\tau)$ (the other case is similar), take the white tile τ'' with $\mathrm{rt}(\tau'') = \mathrm{rt}(\tau)$. Then the edge $\mathrm{br}(\tau'')$ has height h and lies to the right of e. So e cannot belong to the "right boundary" of $\bigcup(\tau \in T)$.

By similar reasonings, e cannot belong to the "left boundary" of $\bigcup(\tau \in T)$. This yields (i). Property (ii) for the vertices z_0, z_n easily follows from (i), and is trivial for the other vertices.

In view of (7.2)(i), we may speak of T as a g-tiling *on* $Z(\omega', \omega)$. We proceed with several observations.

(O1) When ω', ω obey (7.1), at least one g-tiling, even a pure tiling, on $Z(\omega', \omega)$ does exist (assuming $\omega \neq \omega'$). (This generalizes a result in [5] where $\omega' = \mathrm{id}$ and ω is arbitrary; in this case (7.1) is obvious.) It can be constructed by the following procedure, that we call *stripping* $Z(\omega', \omega)$ *along* P_ω *from below*. At the first iteration of this procedure, we take the minimum i such that the edges $e^i_{\omega'}, e^i_\omega$ are different, and take the edge e^k_ω such that $\omega'^{-1}(i) = \omega^{-1}(k) =: c$. Then $k > i$. Let P' be the part of P_ω from $v^{i-1}_{\omega'} = v^{i-1}_\omega$ to v^{k-1}_ω. Using (7.1) for this i and $j = i, \dots, k-1$, one can see that the label $\omega^{-1}(j) =: c_j$ of each edge e^j_ω of P' is greater than c. So we can form the cc_j-tiles τ_j with $\mathrm{br}(\tau_j) = e^j_\omega$. Then $\mathrm{bl}(\tau_i) = e^i_{\omega'}$ and $\mathrm{rt}(\tau_{k-1}) = e^k_\omega$. Therefore, these tiles determine c-strip Q connecting the edge $e^i_{\omega'}$ to the edge e^k_ω. Replace in P_ω the subpath P' followed by the edge e^k_ω by the edge $e^i_{\omega'}$ followed by the left boundary of Q (beginning at $v^i_{\omega'}$ and ending at v^k_ω). The obtained path $\tilde{P}$ determines the permutation ω'' (i.e. $\tilde{P} = P_{\omega''}$) for which the set $I^j_{\omega''}$ is expressed as $I^{j-1}_\omega \cup \{c\}$ for $j = i, \dots, k$, and is equal to I^j_ω otherwise. Using this, one can check that (7.1) continues to hold when ω is replaced by ω''. At the second iteration, we handle the pair (ω', ω'') (satisfying $|P_{\omega'} \cap P_{\omega''}| > |P_{\omega'} \cap P_\omega|$) in a similar way, and so on until the current "right" path turns into $P_{\omega'}$. The tiles constructed during the procedure give a pure tiling T on $Z(\omega', \omega)$, as required.

(O2) Conversely, let ω', ω be two permutations on $[n]$ such that P_ω lies to the right of $P_{\omega'}$. Suppose that there exists a pure tiling T on $Z(\omega', \omega)$. Let $i, j \in [n]$. Then $I^i_{\omega'}$ is the set of edge labels in the subpath P' of $P_{\omega'}$ from its beginning to $v^i_{\omega'}$, and I^j_ω is the set of edge labels in the subpath P of P_ω from its beginning to v^j_ω. Take arbitrary elements $a \in I^i_{\omega'}$ and $b \in I^j_\omega$, and consider in T the a-strip Q_a and the b-strip Q_b, each having the first edge on $P_{\omega'}$ and the last edge on P_ω. So Q_a begins with an edge in P'; if it ends with an edge in P, then a is a common element of $I^i_{\omega'}, I^j_\omega$. In its turn, Q_b ends in P; if it begins in P', then $b \in I^i_{\omega'} \cap I^j_\omega$. On the other hand, if Q_a ends in $P_\omega - P$ and Q_b begins in $P_{\omega'} - P'$, then these strips must cross at some tile $\tau \in T$. Moreover, since T is a pure tiling, such a τ is unique (which is easy to show). It is clear that Q_a contains the edge $\mathrm{bl}(\tau)$, and Q_b contains $\mathrm{br}(\tau)$. Hence, $a < b$, implying (7.1).

(O3) One more useful observation is that, for permutations $\omega' \neq \omega$, the existence of a pure tiling on $Z(\omega', \omega)$ (subject to the requirement that P_ω lie to the right of $P_{\omega'}$) is equivalent to satisfying the *weak Bruhat relation* $\omega' \prec \omega$. The latter means that $\mathrm{Inv}(\omega') \subset \mathrm{Inv}(w)$, where $\mathrm{Inv}(w'')$ denotes the set of inversions for a permutation ω''. This can be seen as follows. Let P_ω lie to the right of $P_{\omega'}$ and let T be a pure tiling on $Z(\omega', \omega)$. It is easy to see that there exists a tile $\tau \in T$ that contains two consecutive edges $e^i_{\omega'}, e^{i+1}_{\omega'}$ in $P_{\omega'}$. Then $e^i_{\omega'} = \mathrm{bl}(\tau)$ and $e^{i+1}_{\omega'} = \mathrm{lt}(\tau)$, and the label $c := \omega'^{-1}(i)$ is less than the label $c' := \omega'^{-1}(i+1)$. Therefore, $(c, c') \notin \mathrm{Inv}(\omega')$. Remove τ from T, obtaining a pure tiling on $Z(\omega'', \omega)$, where ω'' is formed from ω' by swapping c

and c', i.e. $\omega''(c) = i + 1$ and $\omega''(c') = i$. Then $\mathrm{Inv}(\omega'') = \mathrm{Inv}(\omega') \cup \{(c, c')\}$. Repeat the procedure for ω''. Eventually, when the current left path turns into P_ω, we reach ω. This yields $\mathrm{Inv}(\omega') \subset \mathrm{Inv}(\omega)$. Note that the number $|T|$ of steps in the procedure is equal to $|\mathrm{Inv}(\omega) - \mathrm{Inv}(\omega')|$, implying $|\mathfrak{S}_T| = \ell(\omega) - \ell(\omega') + n + 1$.

(O4) Conversely, let $\omega' \prec \omega$. Take $i \in [n]$ such that $e^i_{\omega'} \ne e^i_\omega$ and $e^j_{\omega'} = e^j_\omega$ for $j = 1, \dots, i-1$. Let $c := \omega^{-1}(i)$ and $k := \omega'(c)$ (i.e. the edges e^i_ω and $e^k_{\omega'}$ have the same label c). Clearly $k > i$. Let d be the label of $e^{k-1}_{\omega'}$. Then the inequality $k - 1 \ge i$ and the choice of i imply $\omega(d) > i$. So we have $\omega'(d) < \omega'(c)$ and $\omega(d) > \omega(c)$. This is possible only if $d < c$ and $(d, c) \in \mathrm{Inv}(\omega) - \mathrm{Inv}(\omega')$ (in view of $\mathrm{Inv}(\omega') \subset \mathrm{Inv}(\omega)$). Using these facts, one can form the tile τ with $\mathrm{bl}(\tau) = e^{k-1}_{\omega'}$ and $\mathrm{lt}(\tau) = e^k_{\omega'}$ (taking into account that $d < c$). Replacing in $P_{\omega'}$ the edges $e^{k-1}_{\omega'}, e^k_{\omega'}$ by the other two edges of τ, we obtain the path corresponding to the permutation ω'' satisfying $\mathrm{Inv}(\omega'') = \mathrm{Inv}(\omega') \cup \{(d, c)\}$. Repeating the procedure step by step, we eventually reach ω, and the tiles constructed during the process give the desired pure tiling on $Z(\omega', \omega)$.

(O5) Using flip techniques elaborated in [4], one can show the existence of a pure tiling on $Z(\omega', \omega)$ provided that a g-tiling T on it exists. More precisely, extend T to a g-tiling $\tilde{T}$ on Z_n by adding a pure tiling on $Z(\mathrm{id}, \omega')$ and a pure tiling on $Z(\omega, \omega_0)$. If $\tilde{T}$ has a black tile (contained in T), then, as is shown in [4] (Proposition 5.1), one can choose a black tile τ with the following properties: (a) there are five nonterminal vertices Xi, Xk, Xij, Xik, Xjk (in set notation) such that $i < j < k$, the vertices Xij, Xik, Xjk are connected by edges to $t(\tau)$ (lying in the cone of τ at $t(\tau)$); and (b) replacing Xik by Xj (the *lowering flip* w.r.t. the above quintuple) makes the spectrum of some other g-tiling $\tilde{T}'$ on Z_n. Moreover, the transformation $\tilde{T} \mapsto \tilde{T}'$ is local and involves only tiles having a vertex at Xik (and the new tiles have a vertex at Xj). This implies that the tiles of $\tilde{T}'$ contained in $Z(\omega', \omega)$ form a g-tiling T' on it, whereas the other tiles (lying within $Z(\mathrm{id}, \omega') \cup Z(\omega, \omega_0)$) are exactly the same as those in $\tilde{T}$. Since the flip decreases (by one) the total size of sets in the spectrum, we can conclude that a pure tiling for $Z(\omega', \omega)$ does exist.

(O6) Suppose that the path P_ω lies to the right of $P_{\omega'}$ and that the sets $I^i_{\omega'}, I^j_\omega$ (over all $i, j = 1, \dots, n$) form a ws-collection. By Theorem B, $\mathcal{C}$ is extendable to a largest ws-collection $\mathcal{C}'$, and by Theorem 3.1, there exists a g-tiling T on Z_n with $\mathfrak{S}_T = \mathcal{C}'$. All sets in $\mathcal{C}$ are nonterminal vertices of T, and by Proposition 3.11(i), all edges in $P_{\omega'}$ and P_ω are edges of T. Let T' be the set of tiles $\tau \in T$ such that $\sigma(\tau)$ lies in the simply connected region in D_T bounded by $\sigma(P_{\omega'})$ and $\sigma(P_\omega)$. Then T' is a g-tiling on $Z(\omega', \omega)$.

Summing up the above observations, we obtain the following

Theorem 7.1 *For distinct permutations ω', ω on $[n]$, the following are equivalent*:

(i) *ω', ω satisfy* (7.1);
(ii) *ω', ω satisfy the weak Bruhat relation $\omega' \prec \omega$*;
(iii) *P_ω lies to the right of $P_{\omega'}$ and $Z(\omega', \omega)$ admits a pure tiling*;
(iv) *P_ω lies to the right of $P_{\omega'}$ and $Z(\omega', \omega)$ admits a generalized tiling*;
(v) *P_ω lies to the right of $P_{\omega'}$ and $\{I^i_\omega, I^i_{\omega'} : i = 1, \dots, n\}$ is a ws-collection.*

Now return to the case of one permutation ω. Let us apply the procedure of stripping $Z(\omega,\omega_0)$ along $\mathrm{rbd}(Z_n)$ *from above* (cf. the procedure in (O1)). One can check that the pure tiling T'' on $Z(\omega,\omega_0)$ obtained in this way has the spectrum $\mathfrak{S}_{T''}$ to be exactly the ω-checker $\mathcal{C}^0_\omega$ defined in (2.1). Such a T'' for $n=5$ and $\omega=31524$ is illustrated in the picture.

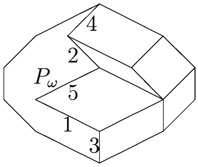

Remark We refer to T'' as above as the *standard tiling* on $Z(\omega,\omega_0)$. This adopts, to the ω case, terminology from [4] where a similar tiling for $\omega=\mathrm{id}$ is called the standard tiling on the zonogon Z_n. The spectrum of the latter consists of all intervals in $[n]$; this is just the collection of $X\cap Y$ over all vertices X in $\mathrm{lbd}(Z_n)$ (i.e. the ideals for id) and all vertices Y in $\mathrm{rbd}(Z_n)$ (i.e. the ideals for ω_0). The spectrum of T'' possesses a similar property: it is the collection $\{I^i_\omega\cap I^j_{\omega_0}: i,j\in[n]\}$ (with repeated sets ignored). (Cf. (2.1) where the term $[j..n]$ is just j-th ideal for ω_0. One can see that withdrawal of the condition $j\le\omega^{-1}(k)$ results in the same collection $\mathcal{C}^0_\omega$.) Also one can check that the same tiling T'' is obtained if we make stripping $Z(\omega,\omega_0)$ *along* P_ω from above. It turns out that a similar phenomenon takes place for any permutations ω',ω obeying (7.1): one can show that stripping $Z(\omega',\omega)$ along P_ω (or along $P_{\omega'}$) from above results in a pure tiling on $Z(\omega',\omega)$ whose spectrum consists of all (different) sets of the form $I^i_{\omega'}\cap I^j_\omega$, $i,j\in[n]$; we may refer to it as the *standard tiling for* (ω',ω).

By reasonings in Sect. 2, for any maximal ω-chamber ws-collection $\mathcal{C}$, the collection $\mathcal{D}:=\mathcal{C}\cup\mathcal{C}^0_\omega$ is a largest ws-collection. So, by Theorem 3.1, $\mathcal{D}=\mathfrak{S}_T$ for some g-tiling T on Z_n. In view of Proposition 3.11, T must include the subtiling T'' as above. Then each edge in $(P_{\mathrm{id}}\cup P_\omega)-(P_{\mathrm{id}}\cap P_\omega)$ belongs to exactly one tile in $T':=T-T''$; this implies that T' is a g-tiling on $Z(\mathrm{id},\omega)$, and we can conclude that $\mathcal{C}=\mathfrak{S}_{T'}$. Conversely (in view of Theorem 2.1), for any g-tiling T' on $Z(\mathrm{id},\omega)$, $\mathfrak{S}_{T'}$ is a maximal ω-chamber collection. One can see that the role of ω-checker can be played, in essence, by the spectrum of *any* pure, or even generalized, tiling T'' on $Z(\omega,\omega_0)$ (i.e. Theorem 2.1 remains valid if we take $\mathfrak{S}_{T''}$ in place of $\mathcal{C}^0_\omega$). Thus, we obtain the following

Corollary 7.2 (a) *Any maximal ω-chamber ws-collection in $2^{[n]}$ is the spectrum of some g-tiling on $Z(\mathrm{id},\omega)$, and vice versa. In particular, any ω-chamber set X lies to the left of the path P_ω (regarding X as a point).*

(b) *For any fixed g-tiling T'' on $Z(\omega,\omega_0)$, $X\subseteq[n]$ is an ω-chamber set if and only if $X\notin\mathfrak{S}_{T''}-\mathcal{I}_\omega$ and X (ws) $\mathfrak{S}_{T''}$, where $\mathcal{I}_\omega:=\{I^0_\omega,\dots,I^n_\omega\}$.*

(c) *$X\subseteq[n]$ is an ω-chamber set if and only if X (ws) $\mathcal{I}_\omega$ and $X\lessdot I^{|X|}_\omega$.*

(Note that X (ws) $\mathcal{I}_\omega$ and $X\lessdot I^{|X|}_\omega$ easily imply that either $X\lessdot I^i_\omega$ or $X\supseteq I^i_\omega$ holds for each i.) In light of (a) in this corollary, it is not confusing to refer to an

ω-chamber set X as a *left set* for ω. A reasonable question is how to characterize, in terms of ω, the corresponding sets X lying to the right of P_ω (i.e. when X belongs to the spectrum of some g-tiling on the right region $Z(\omega,\omega_0)$ for ω). To do this, suppose we turn the zonogon at 180° and reverse the edges. Then the path P_ω reverses and becomes the path $P_{\bar\omega}$ for $\bar\omega := \omega_0\omega$, each set $X \subseteq [n]$ is replaced by $[n]-X$, and $Z(\omega,\omega_0)$ turns into the region $Z(\mathrm{id},\bar\omega)$ lying to the left of $P_{\bar\omega}$. The spectrum of any g-tiling on $Z(\mathrm{id},\bar\omega)$ is formed by left sets $[n]-X$ for $\bar\omega$, and when going back to the original X (thus lying in the region $Z(\omega,\omega_0)$), we observe that X is characterized by the condition:

(7.3) for each $i \in X$, X contains all $j \in [n]$ such that $j > i$ and $\bar\omega(j) > \bar\omega(i)$; equivalently: X contains all j such that $j > i$ but $\omega(j) < \omega(i)$.

We refer to such an X as a *right set* for ω.

Finally, consider again two permutations ω', ω and let $\omega' \prec \omega$. Representing the middle region $Z(\omega',\omega)$ as the intersection of $Z(\mathrm{id},\omega)$ and $Z(\omega',\omega_0)$ and relying on the analysis above, we can conclude with the following generalization of Theorem A.

Theorem A′ *Let ω, ω' be two permutations on $[n]$ satisfying $\omega' \prec \omega$. Then all maximal ws-collections $\mathcal{C} \subseteq 2^{[n]}$ whose members X are simultaneously left sets for ω and right sets for ω' have the same cardinality*; namely, $|\mathcal{C}| = \ell(\omega) - \ell(\omega') + n + 1$. *These collections $\mathcal{C}$ are precisely the spectra of g-tilings on $Z(\omega',\omega)$.*

The sets X figured in this theorem can be alternatively characterized by the condition: for $i, j \in [n]$, if $\omega'(i) \prec \omega'(j)$, $\omega(i) \prec \omega(j)$ and $j \in X$, then $i \in X$, i.e. X is an ideal of the partial order on $[n]$ that is the intersection of two linear orders, one being generated by the path $P_{\omega'}$, and the other by P_ω.

Acknowledgements We thank the referees for useful remarks and suggestions. This work is partially supported by the Russian Foundation of Basic Research and the Ministry of National Education of France (project RFBR 10-01-9311-CNRSL-a). A part of this research was done while the third author was visiting RIMS, Kyoto University and IHES, Bures-sur-Yvette and he thanks these institutes for hospitality.

References

1. Berenstein, A., Fomin, S., Zelevinsky, A.: Parametrizations of canonical bases and totally positive matrices. Adv. Math. **122**, 49–149 (1996)
2. Birkhoff, G.: Lattice Theory, 3rd edn. Am. Math. Soc., Providence (1967)
3. Danilov, V., Karzanov, A., Koshevoy, G.: Tropical Plücker functions and their bases. In: Litvinov, G.L., Sergeev, S.N. (eds.) Tropical and Idempotent Mathematics. Contemp. Math. **495**, 127–158 (2009)
4. Danilov, V., Karzanov, A., Koshevoy, G.: Plücker environments, wiring and tiling diagrams, and weakly separated set-systems. Adv. Math. **224**, 1–44 (2010)
5. Elnitsky, S.: Rhombic tilings of polygons and classes of reduced words in Coxeter groups. J. Combin. Theory, Ser. A **77**, 193–221 (1997)
6. Fan, C.K.: A Hecke algebra quotient and some combinatorial applications. J. Algebraic Combin. **5**, 175–189 (1996)
7. Knuth, D.E.: Axioms and Hulls. Lecture Notes in Computer Science, vol. 606. Springer, Berlin (1992)
8. Leclerc, B., Zelevinsky, A.: Quasicommuting families of quantum Plücker coordinates. Am. Math. Soc. Transl., Ser. 2 **181**, 85–108 (1998)
9. Stembridge, J.: On the fully commutative elements of Coxeter groups. J. Algebr. Comb. **5**, 353–385 (1996)