HAL
open science

Tilings of a Domain on a Hexagon Mesh with Balanced 3-Tiles

Gilles Radenne

To cite this version:

Gilles Radenne. Tilings of a Domain on a Hexagon Mesh with Balanced 3-Tiles. Discrete Models: Combinatorics, Computation, and Geometry, DM-CCG 2001, 2001, Paris, France. pp.287-300, 10.46298/dmtcs.2287 . hal-01182967

HAL Id: hal-01182967

https://inria.hal.science/hal-01182967v1

Submitted on 6 Aug 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Discrete Mathematics and Theoretical Computer Science Proceedings **AA (DM-CCG)**, 2001, 287–300

Tilings of a Domain on a Hexagon Mesh with Balanced 3-Tiles

Gilles Radenne

received February 3, 2001, revised April 20, 2001, accepted May 4, 2001.

In this article, we study the question of tilings on a hexagon mesh with balanced 3-tiles. This problem has been studied by Conway and Lagarias in [CL90], by studying the tiling groups, in fact a group containing the tiling-groups, and their Cayley graphs. We will use two different approaches. The first one is based on matchings in bipartite graphs, which in this case are in correspondance with tilings of domains by lozenges, and thus can be efficiently studied, using Thurston's algorithm. (See [Thu90]). The second one is based on a color and balancing approach of Thurston's algorithm, exposed in [Fou96].

Keywords: Hexagon, Tiling, Balancing, Matching

1 Definitions and basic facts

1.1 Hexagon domains

Consider a covering of the plane with regular hexagons, which also will be called cells, with a 3-coloring. We will call the colors a, b and c. (i, j, k) will always be a permutation of the three colors.

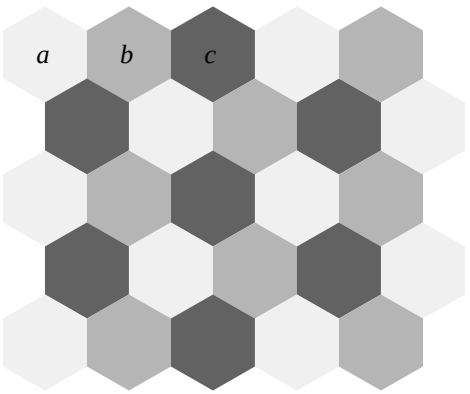

We consider the set of balanced 3-tiles, that is the tiles obtained by gluing together three hexagons of three differents colors. There are five such tiles, two of which we will call triangles:

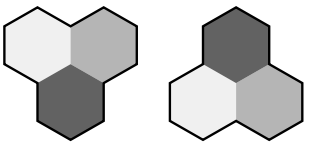

The three others will be called bars, and we will define the color of a bar as the color of its center hexagon.

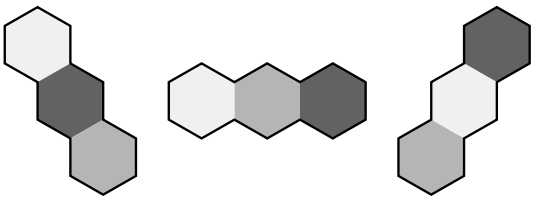

three bars of different colors

1365–8050 © 2001 Maison de l'Informatique et des Mathématiques Discrètes (MIMD), Paris, France

We now consider a domain $\mathcal{D}$ in the plane. What we will call a tiling of $\mathcal{D}$ is a tiling with bars and triangles. A necessary condition for a domain to admit a tiling is evidently to be balanced, that is to have equal number of hexagons of the three colors.

Such tilings have been studied in [CL90] using tiling groups, but in the case of the triangles tiles the tiling group itself has very poor interest. In fact, the border of a domain $\mathcal{D}$ is in the tiling group iff the domain is balanced.

To prove this, it is enough to prove that any balanced set of three hexagons, even a disconnected one, admits a signed tiling by triangles. This is true, as each colored hexagon to tile can by "translated" in the plane with the following sets of signed tiles:

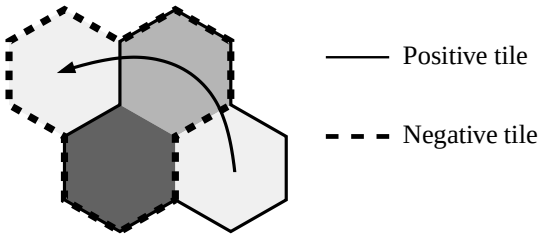

Thus given a balanced set of threee hexagons, even a disconnected one, we can build a signed tiling by "moving" the hexagons to be tiled and placing them together, and covering them with a tile, in the following way (The plus and minus represent the positive and negative tiles):

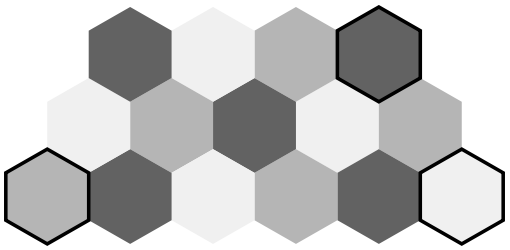

Initial set

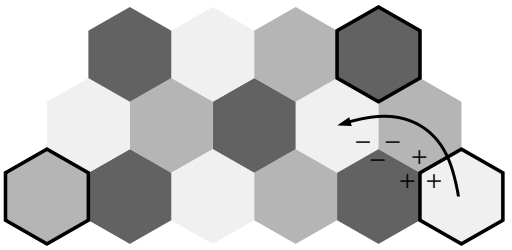

Moving a *a*-hexagon

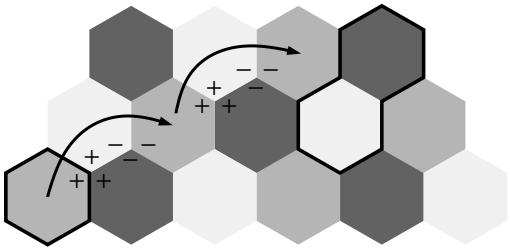

Moving a *b*-hexagon

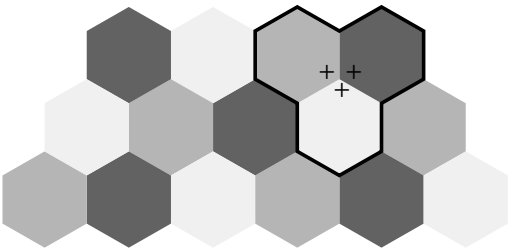

Tiling the new set

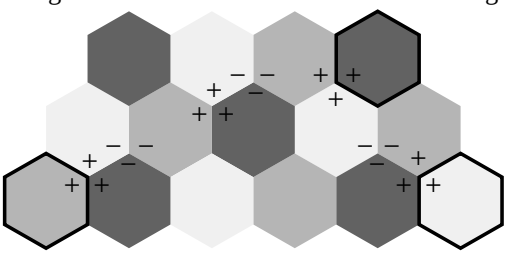

The complete signed tiling

So unless specified otherwise, we will thus study only balanced domains.

1.2 Special tilings

Among these possible tilings of a domain, we will define particular tilings, easier to study.

A perfect tiling will be a tiling using only triangles tiles.

We will call $Per\,(\mathcal{D})$ the set of all perfect tilings of $\mathcal{D}$.

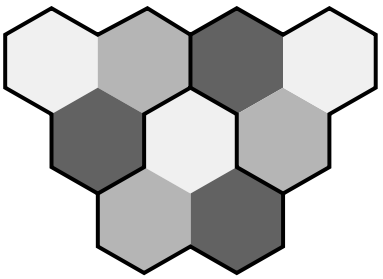

A perfect tiling of a domain

A monochromatic tiling will be a tiling whose all bar tiles are of the same color. If they are of color i, for example, it will be called an i-monochromatic tiling.

Note that a domain admitting a monochromatic tiling admits not necessarily monochromatic tilings for the other colors. We will call $Mon_i(\mathcal{D})$ the set of all i-monochromatic tilings of $\mathcal{D}$.

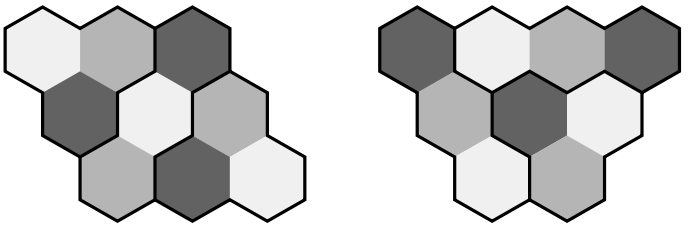

A a- and a b-monochromatic tiling.

For each color i, $Per(\mathcal{D}) \in Mon_i(\mathcal{D})$.
More precisely, $Per(\mathcal{D}) = Mon_i(\mathcal{D}) \bigcap Mon_j(\mathcal{D}) \bigcap Mon_k(\mathcal{D}) = Mon_i(\mathcal{D}) \bigcap Mon_j(\mathcal{D})$ for every permutation i, j, k of a, b, c.

Many balanced domains admit tilings, whereas there are less domains with monochromatic tilings, and domains with a perfect tiling are quite scarce.

It is shown in [CL90] that for a given domain $\mathcal{D}$, the parity of the number of bars in a tiling is an invariant of the domain. We will call it the parity of the domain. [CL90] provides us with an esay way of computing the parity of a domain, using Cayley Graphs. So, if $\mathcal{D}$ admits a tiling with an odd number of bars, that is if $\mathcal{D}$ is odd, there cannot be a perfect tiling of $\mathcal{D}$.

2 Monochromatic tilings and matchings in bipartite graphs

2.1 Correspondance between matchings and tilings

We will begin by reminding what happens in tilings by lozenges: tilings are in one-to-one correspondance with matchings in the bipartite cell-graph of the domain:

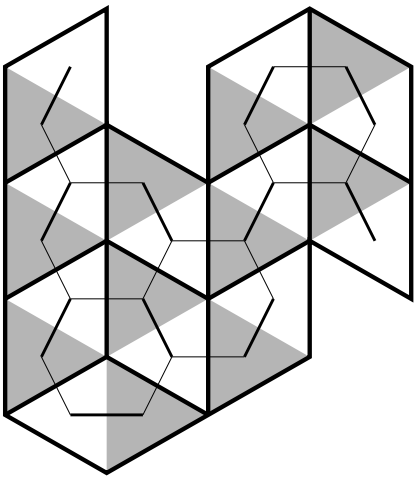

A matching in a bipartite graph can be found in polynomial time. Moreover, in this precise case when the domain is without holes, Thurston's algorithm gives us a tiling, thus a matching, in linear time.

Now let's come back to the hexagon case. The cell graph is now tripartite, inscribed in a triangle mesh:

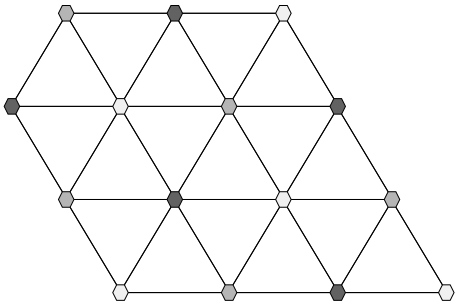

Consider now a domain $\mathcal{D}$, and a i-monochromatic tiling of $\mathcal{D}$. We will call $G_{\mathcal{D}}$ the tripartite cell graph of $\mathcal{D}$. A $i$tile defines then two links, one between a i cell and an adjacent j cell, and one between the same i cell and an adjacent k cell. As each cell is in one unique tile, this defines two matchings in the bipartite graphs given by restricting $G_{\mathcal{D}}$ to i and j cells for one, and i and k cells for the other. For two colors i and j we will call $G_{\mathcal{D}}^{i,j}$ the bipartite graph obtained from $G_{\mathcal{D}}$ by restriction to the cells of color i and j. For a bipartite graph G, we will call $Mat\,(G)$ the set of perfects tilings of G.

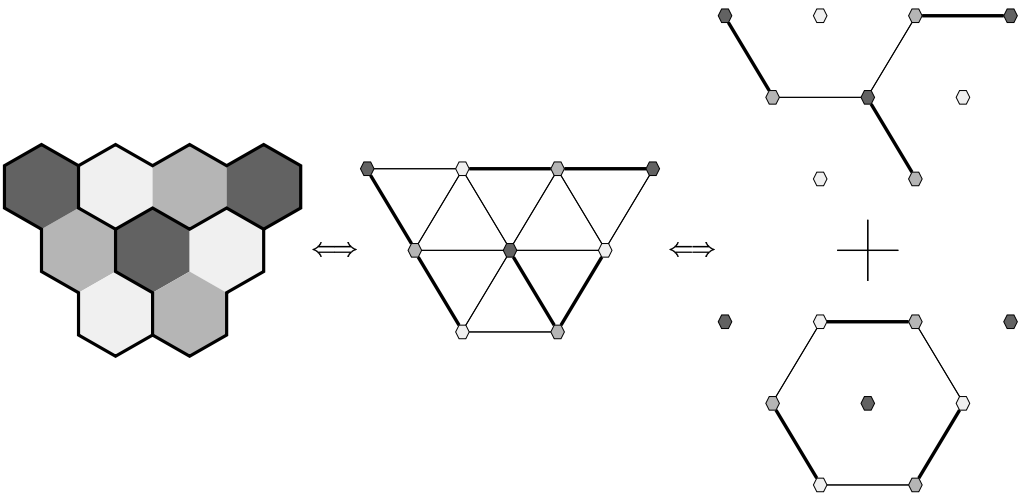

Conversely, given a i cell, if we link it with two adjacent j and k cells, the three cells form either a triangle or a i bar, so from two matchings bertween the i and j cell and between the i and k cell we obtain a i-monochromatic tiling. It is clear that these correspondances are injective, which gives us:

Theorem 1 *For a domain $\mathcal{D}$, the set $Mon_i(\mathcal{D})$ of i-monochromatic tilings of $\mathcal{D}$ is in one-to-one correspondance with $Mat(G_{\mathcal{D}}^{i,j}) \times Mat(G_{\mathcal{D}}^{i,k})$. So $\mathcal{D}$ admits a i-monochromatic tiling iff neither $Mat(G_{\mathcal{D}}^{i,j})$ nor $Mat(G_{\mathcal{D}}^{i,k})$ is the empty set.*

Corollary 1 *If $\mathcal{D}$ admits a i-monochromatic tiling and a j-monochromatic tiling, it admits a k-monochromatic tiling.*

2.2 Computing a monochromatic tiling

So the problem of the existence of a monochromatic graph has been reduced to matchings in bipartite graphs. There are known algorithms to compute such matchings when they exist, in polynomial time. Can we in the present case compute it in linear time? First, let's have a look at the bipartite graph $Mat(G_{\mathcal{D}}^{i,j})$ for $(i,j) = (a,b)$.

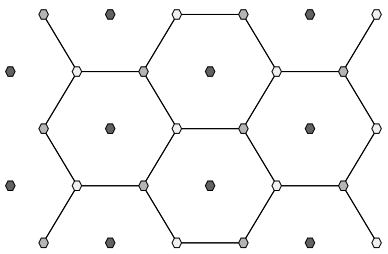

This graph has an hexagonal structure, the same as the cell-graph of an triangle mesh, so we can define a domain $\mathcal{D}^{i,j}$ composed of triangles, whose cell graph is $G^{i,j}_{\mathcal{D}}$. It can be obtained from $\mathcal{D}$ by changing i and j cells in triangle in the following way:

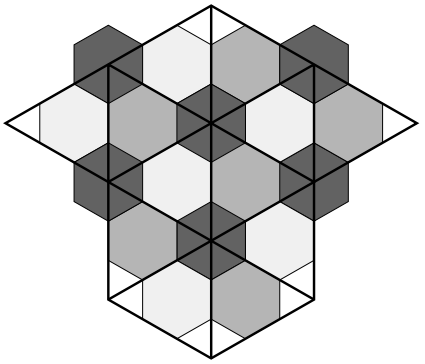

As $\mathcal{D}$ is without hole the only case where $\mathcal{D}^{i,j}$ can have a hole is the following case, in which the hole can be considered as a place where the border meets the same point twice, which we will consider as two different points in the border.

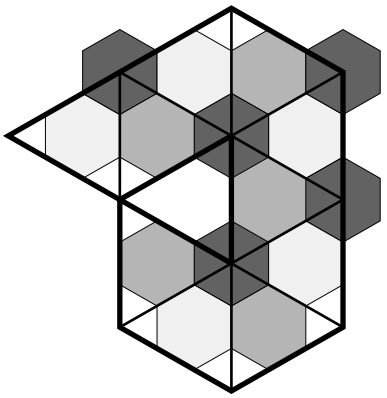

Even with such a hole, Thurston algorithm still works, so we can compute a matching of $Mat(G^{i,j}_{\mathcal{D}})$ in linear time, thus a i-monochromatic tiling can also be computed in linear time.

2.3 The lattice of monochromatic graphs

Given a domain $\mathcal{D}$, we will now study the set $Mon_i(\mathcal{D})$ of i-monochromatic tilings of $\mathcal{D}$, the way we can obtain them from a given tiling by elementary modification in the tiling, and the graph structure it induces on $Mon_i\mathcal{D}$.

As $Mon_i\mathcal{D} \cong Mat(G^{i,j}_{\mathcal{D}}) \times Mat(G^{i,k}_{\mathcal{D}})$, let's see what happen on $Mat(G^{i,j}_{\mathcal{D}})$.

As shown in [Des], $Mat(G^{i,j}_{\mathcal{D}})$ has a natural lattice structure, and we can obtain a tiling from another by a succession of flips of three lozenges forming an hexagon. Furthermore, in its cover graph, two vertices are linked iff the corresponding tilings differ by an elementary flip.

Thus $Mon_i\mathcal{D}$ has also a natural lattice structure, the product lattice of $Mat(G^{i,j}_{\mathcal{D}})$ by $Mat(G^{i,k}_{\mathcal{D}})$, though this lattice structure lacks the natural interpretation as a height function, or equivalently a pile of unit cubes, given in [Des]. The elementary transition between two tilings is induced by the elementary transition in

$Mat(G^{i,j}_{\mathcal{D}})$ or in $Mat(G^{i,k}_{\mathcal{D}})$. An exploration of all the cases gives us three possible transition. With $i = a$ and $j = b$, the transitions induced in $Mon_i\mathcal{D}$ by a flip in $Mat(G^{i,j}_{\mathcal{D}})$ are:

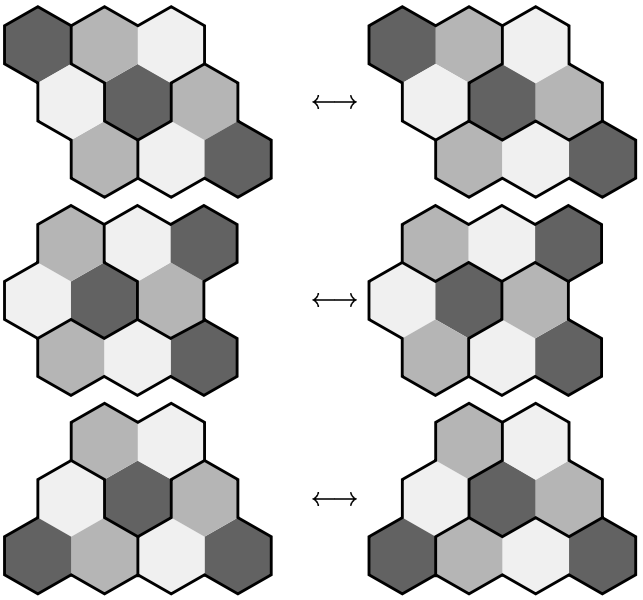

There are two kind of elementary transitions: Transitions 1 and 2 are just reflections in a patch of 3 tiles, whereas transition 3 is between two completely different tilings of a patch. It is the only elementary transitions which changes the numer of bars and triangles of the tiling.

If we consider the graph of i-monochromatic tilings whose edges correspond to elementary transitions, it's the product of the tiling graphs of $G^{i,j}_{\mathcal{D}}$ and $G^{i,k}_{\mathcal{D}}$. As these two graphs are bipartite, the product is bipartite too. This graph is the cover graph of the lattice structure of $Mon_i\mathcal{D}$, so it is the product graph of the cover graphs of $G^{i,j}_{\mathcal{D}}$ and $G^{i,k}_{\mathcal{D}}$. As [Des] gives an algorithm for finding the lattice, thus the cover graph of tilings with lozenges, by applying this algorithm to $G^{i,j}_{\mathcal{D}}$ and $G^{i,k}_{\mathcal{D}}$ we can obtain the graph of $Mon_i\mathcal{D}$.

Let's for example consider the following domain $\mathcal{D}$, and its c-monochromatic tilings:

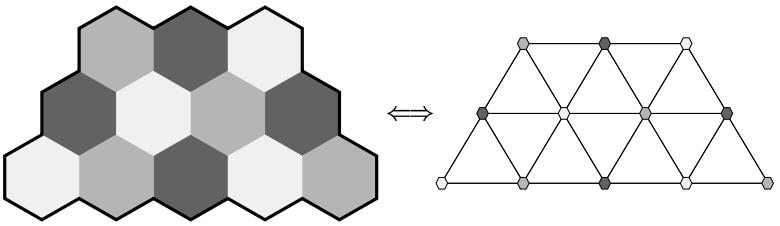

The graphs $Mat(G^{c,a}_{\mathcal{D}})$ and $Mat(G^{c,b}_{\mathcal{D}})$ are:

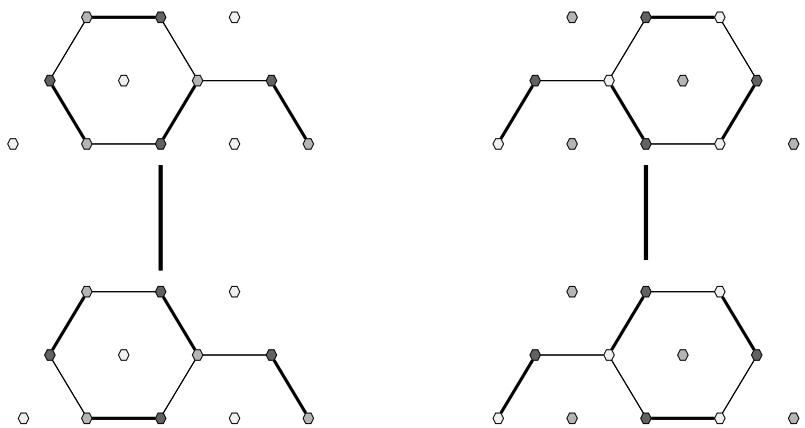

Which yields the following graph for $Mon_c(\mathcal{D})$:

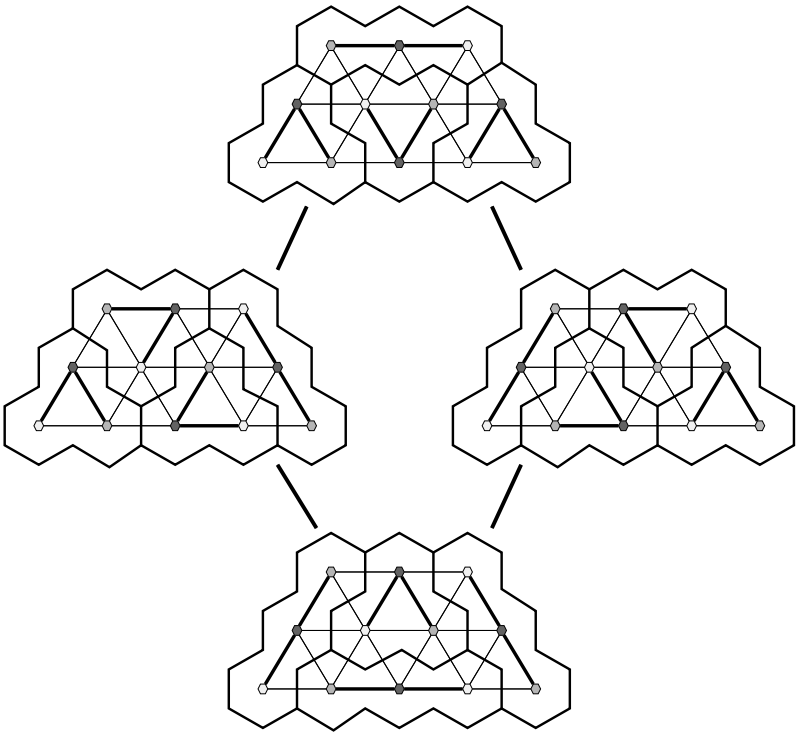

With this approach, we can test the existence of a perfect tiling of $\mathcal{D}$ by computing all the monochromatic tilings of $\mathcal{D}$ and testing all of them to find the perfect ones. But such an algorithm has a high complexity, as the number of lozenge tilings of $G_{\mathcal{D}}^{i,j}$ is at most exponential in regard to the number of cells. (The number of tilings of the union of two domains is at least the product of the number of tilings of each of the domains) So we need a better approach for perfect tilings.

3 Extending Thurston Approach to hexagons

3.1 Balancing interpretation of Thurston's Algorithm

In the lozenges and dominoes cases, Thurston's algorithm ([Thu90]) can be seen as a way to test whether a domain is balanced and locally balanced, which is according, to Hall's theorem, equivalent to the existence of a tiling, as shown in [Fou96]. We will thus extend this approach to hexagons.

A directed path in a regular hexagon mesh will be called direct for color i or i-direct if no edge of this path has a hexagon of color i on its right. We will define the height increase for i of an edge as $+1$ if its left hexagon is of color i, -1 if its right hexagon is of color i, and 0 if neither hexagon is of color i. The height increase $\Delta h_i(P)$ of a directed path P (i-direct or not) will naturally be the sum of the increases of its components. Obviously, a i-direct path will have a nonnegative height increase for i.

Let's now consider a connected domain (which may be not balanced) $\mathcal{D}$, and an anti-clockwise closed path P enclosing this domain. The height increase $\Delta h_i(P)$ for a color is independent of the start point chosen for the path, so we can define the functions Δh_i on domains, which map a domain to the height-increases in the corresponding color. The Δh_i functions on domains are additive in respect to the union of adjacent domains of empty intersection, so for a domain $\mathcal{D}$, $\Delta h_i(\mathcal{D})$ is the sum of the images of its hexagons by h_i. Note that we have for any path P, $\Delta h_i(P) + \Delta h_j(P) + \Delta h_k(P) = 0$, so for a domain $\mathcal{D}$, $\Delta h_i(\mathcal{D}) + \Delta h_j(\mathcal{D}) + \Delta h_k(\mathcal{D}) = 0$

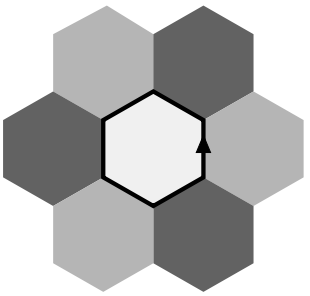

The increase in height of an hexagon is $+6$ for the color of this hexagon, and -3 for the others colors, so if we define a functions N_i which counts the number of hexagons of color i of a domain and N which counts the total number of hexagons, we have $\Delta h_i(\mathcal{D}) = 3(2N_i(\mathcal{D}) - N_j(\mathcal{D}) - N_k(\mathcal{D})) = 3(3N_i(\mathcal{D}) - N(\mathcal{D}))$.

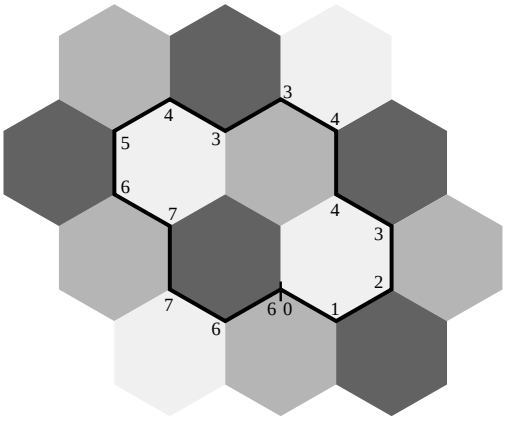

Computing of $\Delta h_a = +6$

We obtain that a domain is balanced iff its three height-increases are 0. As the sum of the three height-increase functions is always 0, the nullity of two of the functions is sufficient. Another consequence of this formula is that the height-increase of a domain is always dividible by 3, so given two points and two paths from one to the other, the difference between height-increase along these paths is dividible by 3.

Now, if a domain $\mathcal{D}$ is balanced, we can define a height h_i function on its border $\partial\mathcal{D}$, by giving the value 0 to an arbitrary origin point of the border and giving to each point of the border the height-increase for i of a path along the border from the origin to this point. As $\mathcal{D}$ is balanced, its height-increases are null, so the h_i are well defined, modulo a constant.

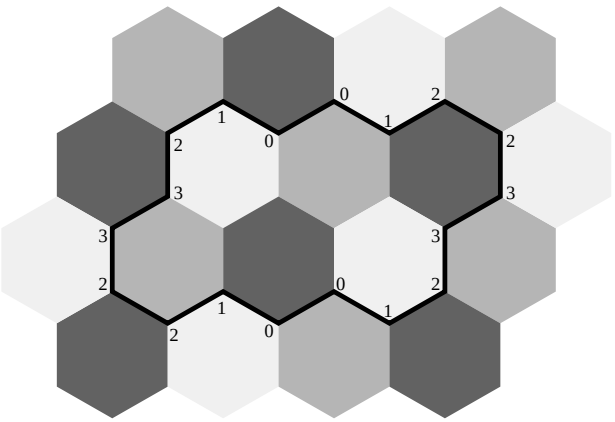

Computing of h_a for a balanced domain

From this function, we now define a function on all the vertices of $\mathcal{D}$, $\widetilde{h}_i$ (the restriction of $\widetilde{h}_i$ on the border $\partial\mathcal{D}$ is not necessarily equal to h_i, as we will see.): $\widetilde{h}_i(x) = min(h_i(y) + d_i(y,x))$, where $d_i(y,x)$ is the minimum height increase for i of a i-direct path from y to x, and the minimum is taken among all the point y on $\partial\mathcal{D}$. We can take $y = x$, which yields that for all x in $\partial\mathcal{D}$, $\widetilde{h}_i(x) \leqslant h_i(x)$.

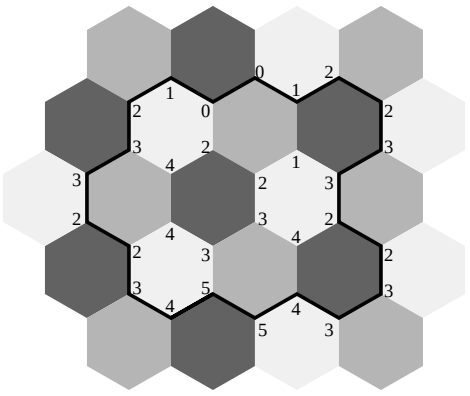

Computing of $\widetilde{h}_a$

We now claim the following theorem:

Theorem 2 *If $\mathcal{D}$ admits a perfect tiling, then for each color i, $\widetilde{h}_i$ and h_i coincide on $\partial\mathcal{D}$.*

(In fact, the existence of monochromatics tilings for each color is a sufficient condition, but as the question of monochromatic tiling has been solved in a much better way, we formulate this in terms of perfect tilings.)

We will say that $\mathcal{D}$ is i-balanced when $\widetilde{h}_i$ restricted to $\partial\mathcal{D}$ coincide with h_i, and well-balanced when it is i-balanced for all color i.

Let's proof this claim: Suppose $\mathcal{D}$ is not well-balanced. There exists then a color i, two points x and y ($x \neq y$) on $\partial\mathcal{D}$, and a i-direct path P from y to x, such that if we call $\Delta h_i(P)$ the height increase for i along P, $\widetilde{h}_i(x) = h_i(y) + \Delta h_i(P) < h_i(x)$. This path P divides $\mathcal{D}$ in two sub-domains, we will call the one on its left $\mathcal{D}_1$ and the other on its right $\mathcal{D}_2$. It is clear that $\Delta h_i(\mathcal{D}_1) = \widetilde{h}_i(x) - h_i(x) < 0$, and $\Delta h_i(\mathcal{D}_1) = 3(2N_i(\mathcal{D}_1) - N_j(\mathcal{D}_1) - N_k(\mathcal{D}_1))$, so $2N_i(\mathcal{D}_1) < N_j(\mathcal{D}_1) + N_k(\mathcal{D}_1)$. So either $N_j(\mathcal{D}_1) > N_i(\mathcal{D}_1)$ or $N_k(\mathcal{D}_1) > N_i(\mathcal{D}_1)$. We will take j and k such as $N_j(\mathcal{D}_1) > N_i(\mathcal{D}_1)$. as P is i-direct, no cell of $\mathcal{D}_1$ meets a hexagon of color i of $\mathcal{D}_2$, so we cannot match all the j-hexagons of $\mathcal{D}_1$ with a i hexagon of $\mathcal{D}$, and so $Mat(G^{i,j}_{\mathcal{D}}) = \varnothing$. Thus $Mon_i(\mathcal{D}) = Mon_j(\mathcal{D}) = \varnothing$. But $Mon_k(\mathcal{D})$ is not necessarily empty. For example, a a bar is a-balanced, but is neither b-balanced nor c-balanced, and it admits a trivial a-monochromatic tiling.

Such a path P will be called an obstruction line.

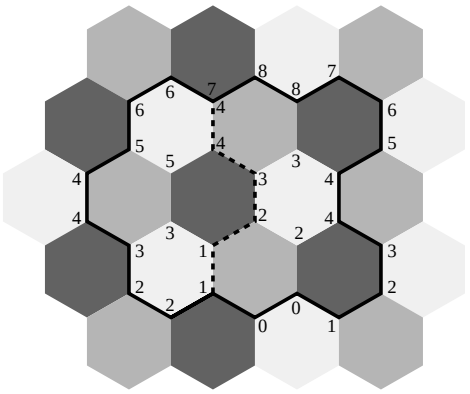

Example of obstruction line for c.

In fact, the proof above shows that $Mon_i(\mathcal{D}) \neq \varnothing$ implies that $\mathcal{D}$ is i-balanced.

Note that a well-balanced domain can in some cases admit no monochromatic tiling at all. Furthermore, to be well-balanced is not a very strong condition, so this theorem is not very useful for the computations of perfect tilings.

Why should we define height functions then ? In fact, they can give us many more informations by the mean of fracture lines, which have been introduced in [Thu90] and used in [Fou96] and [Des].

Given a well-balanced domain $\mathcal{D}$, a path P of $\mathcal{D}$, from one point y of the border to another one x, will be called a fracture line for i if $h_i(x) = h_i(y) + \Delta h_i(P)$. Informally, a fracture line means that $\mathcal{D}$ is

well-balanced, but barely: With $\mathcal{D}_1$ and $\mathcal{D}_2$ defined as above, we have $2N_i(\mathcal{D}_1) = N_j(\mathcal{D}) + N_k(\mathcal{D})$, so if a i-monochromatic tiling exist, as $\mathcal{D}_1$ touches no i-cell of $\mathcal{D}_2$, we have $N_i(\mathcal{D}_1) \geqslant N_j(\mathcal{D}_1)$ and $N_i(\mathcal{D}_1) \geqslant N_k(\mathcal{D}_1)$ (these inequalities are true whenever $\mathcal{D}_1$ is given by a i-direct path and $\mathcal{D}$ admits a i-monochromatic tiling).

Thus $N_i(\mathcal{D}_1) = N_j(\mathcal{D}_1) = N_k(\mathcal{D}_1)$, and in this tilings every i-cell of $\mathcal{D}_1$ is matched with a j-cell and a k-cell of $\mathcal{D}_1$, so this tiling is composed of a i-monochromatic tiling of $\mathcal{D}_1$ and a i-monochromatic tiling of $\mathcal{D}_2$. It is clear that two i-monochromatic tilings, the first one of $\mathcal{D}_1$ and the other one of $\mathcal{D}_2$ form a monochromatic tiling of $\mathcal{D}$.

This decomposition of i-monochromatic tilings is of course true for a perfect ones, which gives us the following theorem:

Theorem 3 *If $\mathcal{D}$ is a well-balanced domain, every fracture line of $\mathcal{D}$ divides $\mathcal{D}$ in two sub-domains $\mathcal{D}_1$ and $\mathcal{D}_2$ such that every perfect tiling of $\mathcal{D}$ is composed of a perfect tiling of $\mathcal{D}_1$ and a perfect tiling of $\mathcal{D}_2$.*

If one of the sub-domain does not admit a perfect tiling, then $\mathcal{D}$ doesn't either.

Note that a fracture line can divide a domain into non-balanced sub domains, which means the domain does not admit perfect tiling. This is the case in the example given for the computing of $\widetilde{h}_a$:

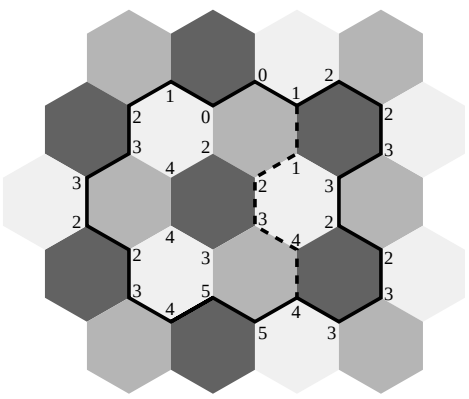

So if we have fracture lines for different colors, their union divides $\mathcal{D}$ in sub-domains with the same decomposition property. Note that fracture lines for different colors can be incompatible, it means that they divide a domain into sub-domain which are not even balanced, which implies that the domain admits no perfect tiling.

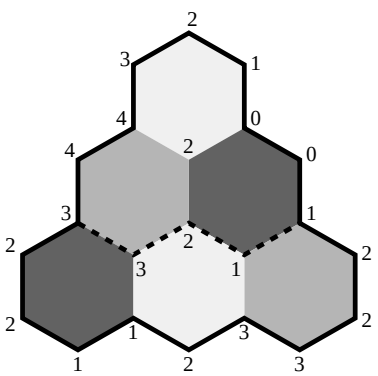

In this domain, the fracture line for the three colors divides the domain into single cells, so it is not perfectly tilable. (The fracture line for the other colors are obtained from the line for a by rotation). We can also conclude by seeing that one of the subdomains obtained from a single fracture line is not perfectly tilable.

On the other hand, fracture lines can directly define a perfect tiling of $\mathcal{D}$:

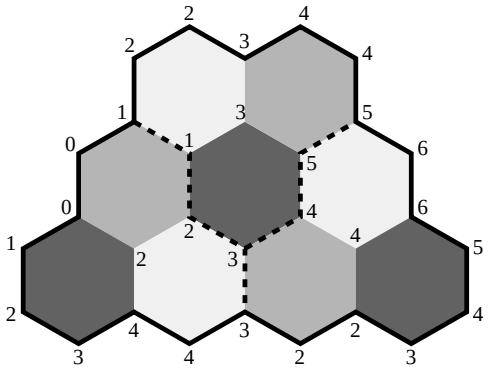

Example of a domain whose fracture line for c provide a perfect tiling.

A well-balanced domain without any fracture line will be called strictly balanced.

A highly important point is that the subdomains given by fracture lines can have their own fracture lines, and so can be divided in smaller subdomains:

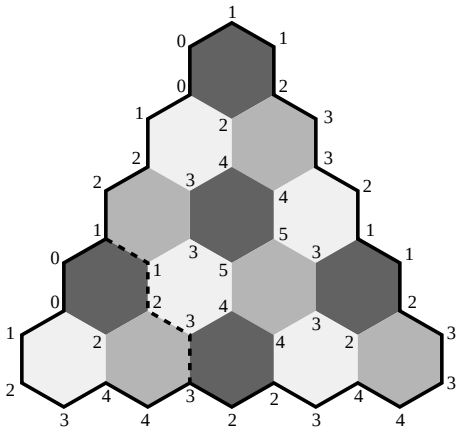

In this domain, the two fracture lines given by the other colors can be obtained from this one by rotation. So the set of fracture line divides this domain in three tiles placed in its edges, and a fourth triangular domain, which has fracture lines showing that it admits no perfect tilings..

So we have a cheap (that is linear) study tool that can be applied to any balanced domain $\mathcal{D}$. It doesn't always give informations on the existence of prefect tilings, but it is nonetheless worthy.

One more thing to add about perfect tilings is that distinct perfect tilings of a domain are not numerous. Contrary to monochromatic tilings, there is no elementary transition between two perfect tilings. In fact, the smaller domain admitting two distinct perfect tilings is the following one:

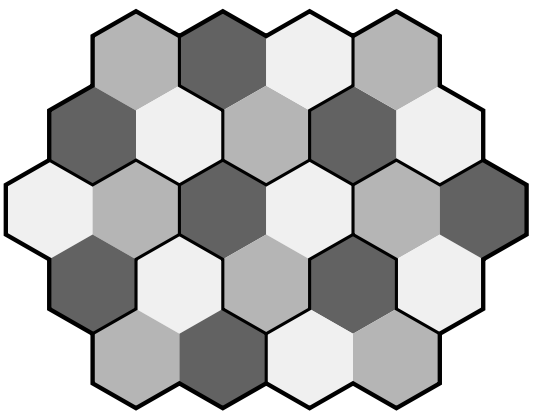

The other perfect tiling of this domain is obtained from this one by a symmetry along the horizontal axis.

3.2 A conjecture on the existence of particular perfect tiling

With this balancing approach, we do not yet have an algorithm which finds a perfect tiling or shows that such a tiling does not exist. Such an algorithm could be obtained from the following conjecture, let it be proved :

Conjecture 1 *Let $\mathcal{D}$ be a stricly balanced domain. If $\mathcal{D}$ admits a perfect tiling, then for each color i, it admits a perfect tiling with a tile placed along a maximum of h_i on the border, in the following way:*

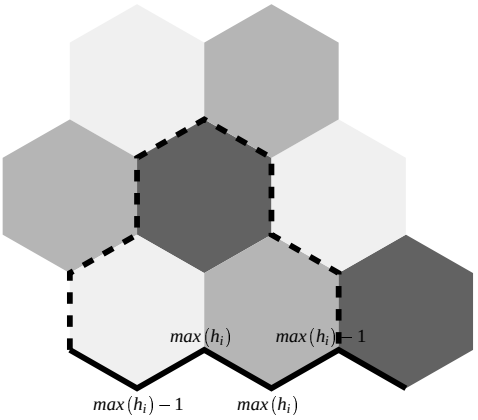

(The position of $max(h_i)$ is imposed by its maximality and the condition that $\mathcal{D}$ does not have any fracture line)

The idea of this conjecture is to extend the iterative approach of Thurston's algorithm.
If this conjecture is proven true, we get the following algorithm for finding a perfect tiling:

- **Goal:** Find a perfect tiling of a balanced domain$\mathcal{D}$.
- **Step 1** Compute all the h_i and $\widetilde{h}_i$ function on $\mathcal{D}$. If $\mathcal{D}$ has an obstruction line, then the domain admits no perfect tiling. Otherwise, if $\mathcal{D}$ admits fracture lines, choose a color i for which there are fracture lines, divide $\mathcal{D}$ is sub-domains $\mathcal{D}_1, \ldots, \mathcal{D}_n$ along these lines and recursively apply the algorithm on all the $\mathcal{D}_l$'s, $1 \leqslant l \leqslant n$.
- **Step 2**: As $\mathcal{D}$ has no fracture line, choose a color i and place a tile along a maximum of h_i.
- **Step 3**: If the remaining domain is not empty, go back to step 1.

  This algorithm could be implemented within at most quadratic complexity.

4 Conclusion and perspectives

We have solved the problem of existence and finding of monochromatic tilings, as well as the graphs $Mon_i(\mathcal{D})$ of all such tilings on $\mathcal{D}$. The problem of perfect tilings is unsolved, but it could be solved by proving the conjecture. However, we could study if it is possible to obtain simply perfect tilings from monochromatic tilings. In particular, it would be interesting to find an answer to the following question: Given a even domain $\mathcal{D}$ which admits monochromatic tilings for all colors, does it always admit a perfect tiling ? (We have not yet found counterexamples. In fact, the parity of a domain is one of the main arguments used by Conway in [CL90] to study perfect tilings.)

Another field of study would be general tilings, in particular the graph of all tilings on $\mathcal{D}$. If we take as edges the transitions of the following types, do we always obtain a connected graph ? (The elementary transitions in a monochromatic graph can of course be obtained from these.)

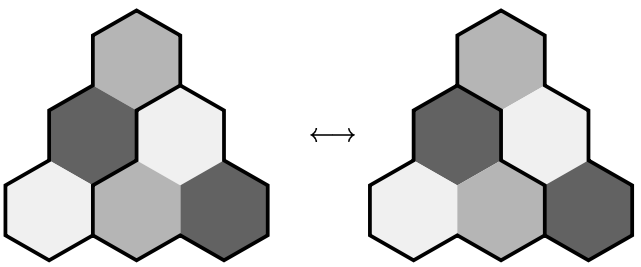

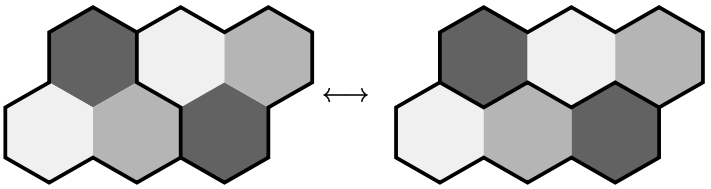

References

[CL90] C.H. Conway and J.C. Lagarias. Tiling with polyominoes and combinatorial group theory. *Journal of combinatorial theory*, Series A:183–207, 1990.

[Des] Sebastien Desreux. Can we enumerate exactly once every tiling with lozenges of a general domain $\mathcal{D}$? *submitted to Theorical Computer Science*. (Sebastien.Desreux@liafa.jussieu.fr).

[Fou96] Jean-Claude Fournier. Pavage des figures planes sans trous par des dominos : Fondement graphique de l'algorithme de Thurston, parallélisation, unicité et décomposition. *Theorical Computer Science*, 159:105–128, 1996.

[Thu90] William P. Thurston. Conway's tiling groups. *American Mathematics Monthly*, pages 753–773, October 1990.