

PERIODIC VERSUS ARBITRARY TESSELLATIONS OF THE PLANE
USING POLYOMINOS OF A SINGLE TYPE

H. A. G. Wyshoff and J. van Leeuwen

RUU - CS - 82 - 11

July 1982



Rijksuniversiteit Utrecht

Vakgroep informatica

Princetonplein 5
Postbus 80.002
3508 TA Utrecht
Telefoon 030-531454
The Netherlands

PERIODIC VERSUS ARBITRARY TESSELLATIONS OF THE PLANE

USING POLYOMINOS OF A SINGLE TYPE

H. A. G. Wijshoff and J. van Leeuwen

Technical Report RUU-CS-82-11

July 1982

Department of Computer Science
University of Utrecht
P.O. Box 80.002, 3508 TA Utrecht
the Netherlands.

PERIODIC VERSUS ARBITRARY TESSELLATIONS OF THE PLANE USING POLYOMINOS OF A SINGLE TYPE

H. A. G. Wijskoff and J. van Leeuwen

Department of Computer Science, University of Utrecht
P.O. Box 80.002, 3508 TA Utrecht, the Netherlands

Abstract. Given N parallel memory modules, nontrivial problems arise if the elements of an (infinite) two-dimensional array are to be distributed in storage such that any set of N elements arranged according to a given data template T can be accessed rapidly in parallel. Array embeddings that allow for this are called skewing schemes and have been studied in connection with vector-processing and SIMD machines. In 1975 H. D. Shapiro proved that there exists a valid skewing scheme for a template T (of size N) if and only if T tessellates the plane. In this paper we settle an important conjecture of Shapiro and prove that for polyominos P (of size N) a valid skewing scheme exists if and only if there exists a valid periodic skewing scheme. (Periodicity implies a rapid technique to actually locate data elements according to the skewing scheme.) The proof shows that when a polyomino P tessellates the plane then it can tessellate the plane "periodically", i.e., with the instances of P arranged in a regular lattice. As a result the existence of a tessellation with polyomino P can be decided in polynomial time.

Keywords and phrases: SIMD machines, parallel memories, two-dimensional arrays, data templates, polyominos, skewing schemes, tessellations, periodicity, lattices, polynomial time.

1. Introduction.

The problem addressed in this paper has a deep motivation from the

Theory of data organisation for large computers such as vector-processing and SIMD-machines (see e.g. Thurber [4]). The characterising feature of these machines is the availability of a multitude of arithmetic units and memory modules that can operate independently in parallel. We assume that processors and memories are connected by a crossbar switch or some equivalent interconnection network that has an equally efficient routing strategy. (There is an abundant literature on the possible interconnection networks that can route all processor-memory assignments rapidly and free of conflict.) Clearly the effectiveness of these machines for numeric computation depends to large extent on being able to store the data elements in the available memories in such a manner that memory conflicts are avoided whenever a batch of data are fetched.

About 1970 Budnik and Kuck [2] pointed out that nontrivial problems arise if the elements of a two-dimensional array have to be stored and distributed over N memory modules such that any set of N elements arranged according to some common pattern or template can be accessed in one cycle, without conflict. A data template T simply consists of a fixed set of N locations or cells relative to a designated base cell or "handle" $(0,0)$. An instance of T is obtained by adding a fixed displacement to all locations of T . An assignment of array elements to memories is called a skewing scheme. A skewing scheme is valid for a template T if it provides for the conflict-free parallel access to the data in any instance of T (located within the array). Clearly there does not always exist a valid skewing scheme for the templates at hand, but in many interesting cases they do.

In 1975 Shapiro (see [3]) added two significant results to this theory. First of all, he proved that there exists a valid skewing scheme for T in all finite cases (square arrays) if and only if there is one for the infinite array with domain $(-\infty:\infty; -\infty:\infty)$. Secondly, he proved that there is a valid skewing scheme for T if and only if T

(as a combinatorial structure) tessellates the two-dimensional plane. As the argument is important, we briefly digress and include our simplified proof of this fact.

Lemma 1.1. Let T_1 and T_2 be instances of T with handles located in h_1 and h_2 respectively. T_1 and T_2 overlap if and only if h_1 and h_2 can be covered by a single instance of T .

Proof.

Suppose T_1 and T_2 overlap in a cell x . It means that $\vec{h_1x}$ and $\vec{h_2x}$ both lead into a cell of the template when used as displacements from the handle. Let h_3 be the "fourth" corner of the parallelogram spanned by h_1, x and h_2 (see figure 1) and imagine an instance T_3 of T is located with its handle in h_3 . It follows that both h_1 and h_2 must be covered by this instance T_3 . The converse is estab-




figure 1.

lished along similar lines. \square

Lemma 1.2. Let T_1 and T_2 be disjoint instances of T . If an instance T_3 overlaps T_1 and T_2 then the cells it covers in T_1 are distinct from the cells it covers in T_2 even when considered as elements of the template.

Proof.

Let the handles of T_1, T_2 and T_3 be located in h_1, h_2 and h_3 respectively. Suppose that T_3 covers a cell x of T_1 and a cell y of T_2 that are identically located with respect to h_1 and h_2 (respectively). It follows that, as displacements, $\vec{h_1x}$ and $\vec{h_2y}$ are identical. Let h_4 be the "fourth" corner of the parallelogram spanned by h_1, x and h_3 (see figure 2) and imagine an instance T_4 of T is located with its handle in h_4 . Observe that h_4, h_3, y and h_2 form a parallelogram as

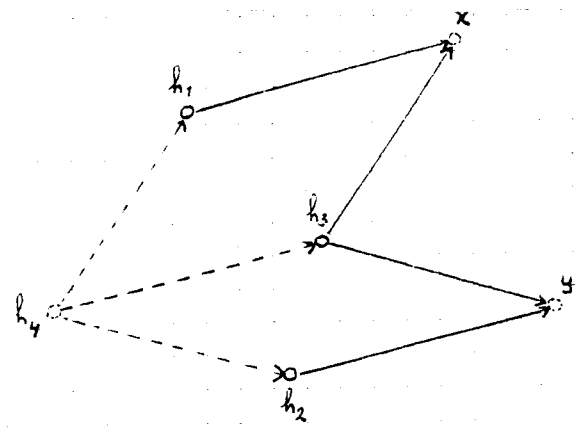


figure 2.

well and that as a result $\overrightarrow{h_4 h_1} = \overrightarrow{h_3 x}$ and $\overrightarrow{h_4 h_2} = \overrightarrow{h_3 y}$. As x and y are both covered by T_3 , it follows that h_1 and h_2 must both be covered by T_4 and thus, using lemma 1.1, that T_1 and T_2 must overlap. This contradicts the disjointness of T_1 and T_2 . \square

Theorem 1.3 There exists a valid skewing scheme s for T if and only if T tessellates the plane.

Proof.

(Note that a skewing scheme is a mapping $s : (-\infty : \infty ; -\infty : \infty) \rightarrow [1 \dots M]$.)

\Rightarrow . Let s be a valid skewing scheme for T . Consider the arrangement A in which an instance of T is located at every cell p with $s(p) = 1$. (Note that every instance of T must have one cell assigned to memory 1 and thus A is not empty.) Any two instances T_1 and T_2 in A must be disjoint. (If not, then lemma 1.1 would ensure the existence of an instance T_3 containing two cells mapped to 1, contradicting the fact that s was valid.) To prove that A is a tessellation we need to show that every cell q is covered. Consider the M possible instances T_i of T that cover q and let their handles be located in cells p_i ($1 \leq i \leq M$). The p_i must all be assigned to different memories (or else another contradiction could be derived with lemma 1.1) and thus there is exactly one p_j such that T_j covers q and $s(p_j) = 1$. This T_j is indeed in A by definition.

\Leftarrow . Let T tessellate the plane. Number the cells of T from 1 to M

and consider the skewing scheme s obtained by repeating this numbering throughout all instances in the tessellation and assigning to each cell the value it got in the numbering. To prove that s is valid for T , consider an arbitrary instance T_i of T . If T_i coincides with an instance from the tessellation, then its cells are trivially assigned to different memories. Otherwise lemma 1.2 shows that T_i takes disjoint bytes out of every instance of T in the tessellation that it intersects when viewed as parts of the original template. Thus all cells in T_i are assigned different numbers even now. \square

General skewing schemes (or: tessellations) are not necessarily of much use in practice. There seems no guarantee that a skewing scheme s is finitely encoded or indeed recursive. This led Shapiro ([3], sect. IV) to consider a number of further constraints to force a skewing scheme to be finitely represented in computer memory. The weakest condition is periodicity which, for the purposes of this paper, is interpreted as the condition that the instances of T that tessellate the plane have their handle located at the points of a finitely generated lattice. (In the literature various other notions of periodicity occur in relation to skewing schemes.) To compute an s -value of a periodic skewing scheme one only needs to have the fixed numbering of T in store and can proceed as in the proof of theorem 1.3 to find the required instance T_j . (The effectiveness follows because we need to test only N times that a point belongs to the lattice.) In practice a tabular method would be used.

In this paper we settle an important conjecture of Shapiro [3] and prove that for templates that have the shape of a polymino there exists a valid skewing scheme if and only if there exists a valid periodic skewing scheme. The proof heavily relies on the geometric interpretation of the problem. We show in fact that when a polymino of size N tessellates the plane, then it can tessellate the plane according to a lattice that

has a basis of cardinality equal to 2. (The precise result is that a periodic tessellation can be designed in which every polyomino is surrounded by 4 polyominoes or every polyomino is surrounded by 6 polyominoes.) As a corollary we show that the existence of a valid skewing scheme for a polyomino of size N can be decided in polynomial time.

The paper is organized as follows. In section 2 we give definitions and some preliminary results dealing with the relative positioning of polyominoes (i.e., of instances of some fixed polyomino). In section 3 we define tessellations and derive a more operational notion of periodicity for tessellations. In section 4 we derive an important condition for the existence of a periodic tessellation (and its consequences of a very regular structure indeed). In section 5 a tedious counting argument involving Euler's formula for planar graphs is given showing that whenever a tessellation with a polyomino exists, then the nodes and numbers must exist required for the condition derived in section 4. In section 6 we prove the polynomial algorithm for the existence of a tessellation with a given polyomino and offer some final comments.

WARNING. We shall often not distinguish between a polyomino and its instances and thus speak just of "polyominoes", without explicitly mentioning that they are all of the same type. We also assume that the orientation of all polyominoes is fixed (as it are instances of one template) and thus we do not consider rotations and reflections when discussing tessellations.

2. Definitions and preliminary results.

All notions introduced in this section are rather straightforward and presented without much commentary. Definitions 2.1 to 2.3 are taken from Shapiro [3].

Definition 2.1 A data template is an N -tuple $T = \{(0,0), (a_1, b_1), \dots, (a_{N-1}, b_{N-1})\}$ with no two components identical and whose first element is $(0,0)$. For consistency we let $(a_0, b_0) \equiv (0,0)$.

Definition 2.2 An instance $T(x,y)$ of a data template T is the N -tuple obtained by the componentwise addition of the displacement (x,y) to T : $T(x,y) = \{(x,y), (a_1+x, b_1+y), \dots, (a_{N-1}+x, b_{N-1}+y)\}$.

Definition 2.3 A polyomino is a data template of which the cells form a rook-wise connected set (when embedded in the plane).

Rook-wise connectedness means that every two cells of the template can be connected by a chain of cells within the template, with every two consecutive cells of the chain sharing a full side.

From now on we fix a polyomino P of size N and introduce some notions pertaining to its set of instances $P(x,y)$. Some of the results would actually be meaningful for arbitrary data templates also.

Definition 2.4 The relative position π of cells (x_1, y_1) and (x_2, y_2) is the "bi-directional" vector $r = \pm \overrightarrow{(x_2 - x_1, y_2 - y_1)}$. The relative position of $P(x_1, y_1)$ and $P(x_2, y_2)$ is the relative position of (x_1, y_1) and (x_2, y_2) .

It is best to think of r as a vector pointing "both ways". Intuitively it is the vector needed to go from one cell to the other. (Observe that the relative position of (x_1, y_1) and (x_2, y_2) is equal to the relative position of (x_2, y_2) and (x_1, y_1) .) In polyominoes $P(x_1, y_1)$ and $P(x_2, y_2)$ all corresponding cells have the same relative position, namely the relative position of $P(x_1, y_1)$ and $P(x_2, y_2)$.

Definition 2.5 $P(x_1, y_1)$ and $P(x_2, y_2)$ overlap if there exist components

(a_i, b_i) and (a_j, b_j) of P such that $(a_i + x_1, b_i + y_1) = (a_j + x_2, b_j + y_2)$.

Lemma 2.6. $P(x_1, y_1)$ and $P(x_2, y_2)$ overlap if and only if P contains two components that are in the same relative position as $P(x_1, y_1)$ and $P(x_2, y_2)$.

Proof.

Clearly $P(x_1, y_1)$ and $P(x_2, y_2)$ overlap if and only if $(a_i, b_i) = (a_j, b_j) + (x_2 - x_1, y_2 - y_1)$ or, equivalently, $(a_j, b_j) = (a_i, b_i) + (x_1 - x_2, y_1 - y_2)$.

□

Also recall lemma 1.1 at this point.

Let $P(x_0, y_0)$ be a fixed instance of P . With every polyomino $P(x, y)$ there actually is a second polyomino (its "buddy") that has the same relative position to $P(x_0, y_0)$.

Definition 2.7. The buddy of $P(x, y)$ with respect to $P(x_0, y_0)$ is the instance $\varphi_{x_0, y_0}(P(x, y)) = P(2x_0 - x, 2y_0 - y)$.

Observing that $(2x_0 - x, 2y_0 - y) = (x_0, y_0) + (x_0 - x, y_0 - y)$ it should be clear (see figure 3) that the buddy of $P(x, y)$ is symmetrically located exactly at the "opposite" side of the polyomino $P(x_0, y_0)$.




figure 3.

It also follows that buddies are paired, i.e., if P_1 is the buddy of P_2 then P_2 is the buddy of P_1 . The mapping φ_{x_0, y_0} implicit in definition 2.7 will be an important one for "filling up" the space surrounding $P(x_0, y_0)$. The following

properties of φ_{x_0, y_0} are worth noting.

Lemma 2.8. φ_{x_0, y_0} preserves relative positions.

Proof.

We have to establish that $\varphi_{x_0, y_0}(P(x_1, y_1))$ and $\varphi_{x_0, y_0}(P(x_2, y_2))$ have the same relative position as $P(x_1, y_1)$ and $P(x_2, y_2)$. A simple calculation




figure 4.

would suffice, but the argument is best seen from a geometric interpretation (see figure 4). In fact it is useful to think of φ_{x_0, y_0} as a reflection of the designated cells around (x_0, y_0) that "carries" the polyomino along in an unreflected manner. This certainly preserves the relative position of corresponding cells throughout the mapping. \square

Using lemma 2.6 it follows in particular that φ_{x_0, y_0} maps disjoint instances of P to disjoint images (buddies).

Lemma 2.9. φ_{x_0, y_0} does not "introduce" overlap, i.e., if $P(x, y)$ and $P(x_0, y_0)$ are disjoint then $\varphi_{x_0, y_0}(P(x, y))$ is disjoint from these instances as well.

Proof.

Let $P(x, y)$ and $P(x_0, y_0)$ be disjoint. Using lemmas 2.6 and 2.8 we easily conclude that $\varphi_{x_0, y_0}(P(x, y))$ must be disjoint from $P(x_0, y_0)$ as well. Suppose $\varphi_{x_0, y_0}(P(x, y))$ is not disjoint from $P(x, y)$ and actually overlaps it in cell $\gamma = (u, v)$. The situation is shown in figure 5 where for sim-




figure 5.

licity we have set $P \equiv P(x, y)$, $P_0 \equiv P(x_0, y_0)$ and $\varphi \equiv \varphi_{x_0, y_0}(P(x, y))$. By observing how γ is located with respect to the handles of $P(x, y)$ and $\varphi_{x_0, y_0}(P(x, y))$ we can identify four more cells $(\alpha, \beta, \delta, \mu)$ that are of interest because of their similar location with respect to the handles of $P(x, y)$, $P(x_0, y_0)$ or $\varphi_{x_0, y_0}(P(x, y))$. Note that α, β, \dots are on a straight line, with a relative position equal to $\pm(x_0 - x, y_0 - y)$ between every consecutive two. Figure 5 shows to which polyomino each of the α, β, \dots must belong.

As $\alpha \in P(x, y)$ and $\gamma \in P(x, y)$ and P is a polyomino, there must be a rook-wise connected chain π of cells leading from α to γ that only uses cells within $P(x, y)$. As β and δ (both $\in P(x_0, y_0)$) are in the same relative position, the chain π' obtained by shifting π over the vector $(x_0 - x, y_0 - y)$ must connect them and run entirely within the polyomino $P(x_0, y_0)$. (This is so because π really is a chain that is fixed for the template.) Because α, γ and β, δ are interlaced π and π' necessarily intersect. Any cell where the chains intersect will belong to both $P(x, y)$ and $P(x_0, y_0)$. This contradicts the fact that they are disjoint. \square

Note that lemma 2.8 holds for templates in general, but that lemma 2.9 makes essential use of the fact that we deal with polyominoes. In later sections "buddying" will be important in analysing disjoint placements of polyominoes around an instance $P(x_0, y_0)$.

3 Tessellations

Tessellations (or "tilings") are a familiar subject in recreational mathematics. Definition 3.1 is taken from Shapiro [3] (although we have added the distinction between partial and total tessellations). For the results of this paper a much deeper understanding of the combinatorial structure of tessellations is required. Thus we shall spend most of our effort analysing (local) conditions that guarantee tessellations.

Definition 3.1 A partial tessellation (using P) is any collection of non-overlapping instances of the polyomino P , i.e., any collection of instances of P with the property that every cell of the plane is in at most one instance. A tessellation is said to be total if every cell of the plane is in exactly one instance of P in the collection.

If no adverb is added we will assume a tessellation to be total. If there exists a total tessellation using P then we say that P "tessellates the plane". Clearly total tessellations are infinite.

Definition 3.2 A (total) tessellation is periodic if it is the collection of instances $P(x, y)$ with (x, y) ranging over the elements of a lattice with finite basis.

Lattices do not normally lead to tessellations. To better grasp the structure of periodic tessellations we need a few more concepts. In the following we assume that polyominoes are drawn as sets of cells on the two-dimensional grid.

Definition 3.3 The boundary B of the (embedded) polyomino P is the set of grid-lines of unit length that bound the interior of P from the exterior. The size of B is denoted as $|B|$. The boundary $B(x, y)$ of $P(x, y)$ is B shifted over by (x, y) .

To be able to refer to them we will number the grid-lines in B going around clockwise as r_0, r_1, \dots , starting from a fixed reference element $r_0 \in B$. Shifting over (x, y) this numbering translates into an "equal" numbering $r_0(x, y), r_1(x, y), \dots$ of $B(x, y)$. We will later see that grid-lines are an important ingredient of our combinatorial reasoning. Note that numbers like $r_i(x, y)$ are merely names of grid-lines with respect to some $P(x, y)$. Different numbers for different instances may refer to the same ac-

Actual grid-line

Now suppose that some partial or total tessellation τ is given. We say $P(x, y) \in \tau$ borders $P(x_0, y_0) \in \tau$ if $B(x, y) \cap B(x_0, y_0) \neq \emptyset$. It is clear that the boundary of $P(x_0, y_0)$ splits up in a number of consecutive segments $[r_{i_0}(x_0, y_0) - r_{i_1}(x_0, y_0)]$, $[r_{i_2}(x_0, y_0) - r_{i_3}(x_0, y_0)]$, ... ($r_{i_0} \leq r_{i_1} \leq r_{i_2} \leq r_{i_3} \leq \dots$) that are the borderline with instances $P(x, y)$. We will assume that each segment is maximal for the particular $P(x, y)$.

Lemma 3.4 If τ is a total tessellation then each $P(x, y) \in \tau$ that borders $P(x_0, y_0) \in \tau$ generates exactly one, contiguous segment on the boundary of $P(x_0, y_0)$.

Proof

If $P(x, y)$ borders $P(x_0, y_0)$ it will lead to at least one segment on $B(x_0, y_0)$. Suppose there are two grid-lines $r_i(x_0, y_0)$ and $r_j(x_0, y_0)$ with $j > i+1$ that border $P(x, y)$ while none of the grid-lines in between do. It follows that there is a "hole" A totally enclosed by the parting and rejoining boundaries of the two polyominoes. The corresponding "segments" on $B(x_0, y_0)$ and $B(x, y)$ are necessarily disjoint and it is easily argued that the circumference of A is less than $|B|$ grid-lines long. We can now argue as follows. Because τ is a tessellation, the entire region A must be tessellated itself with one or more instances of P . Clearly one polyomino will not fit, because its circumference ($|B|$) is larger than A 's. But it is easily shown that any tessellation of A with more instances of P cannot have a circumference of less than $|B|$ grid-lines either. Thus A is "too small" to contain even a single instance of P and must be a real hole, contradicting the fact that τ is a tessellation. \square

The argument in the proof will be of use again later.

Definition 3.5 A partial segmentation of $B(x_0, y_0)$ is any set of (maximal and disjoint) index segments I_0, I_1, \dots along $B(x_0, y_0)$ that each are

the borderline with some $P(x, y) \in \tau$. A segmentation is called total if $\bigcup_i I_i = B(x_0, y_0)$.

A (partial or total) segmentation of a $B(x, y)$ will be denoted as $\text{Seg}(B(x, y))$. Its size, i.e., the number of segments in it, will be denoted by $|\text{Seg}(B(x, y))|$. Its "length" is defined in an obvious manner.

We say that $P(x_1, y_1), \dots, P(x_k, y_k)$ partially surround the instance $P(x_0, y_0)$ if the polyominoes (including $P(x_0, y_0)$) are all disjoint but each $P(x_i, y_i)$ ($i > 0$) borders $P(x_0, y_0)$. We say that $P(x_1, y_1), \dots, P(x_k, y_k)$ completely surround $P(x_0, y_0)$ if, in addition, each grid-line of $B(x_0, y_0)$ is contained in some $B(x_i, y_i)$ ($i > 0$). The size of a (partial or complete) surrounding will be the number (k) of distinct polyominoes in it. Clearly partial surroundings lead to partial segmentations and complete surroundings lead to total segmentations of $B(x_0, y_0)$. Of course lemma 3.4 applies to the elements of a complete surrounding of $P(x_0, y_0)$ just the same. Surroundings and the segmentations they induce are the key to a further understanding of periodic tessellations.

Definition 3.6. A tessellation τ is regular if the same segmentation is induced in every $B(x, y)$ with $P(x, y) \in \tau$, i.e., $\text{Seg}(P(x_1, y_1)) = \text{Seg}(P(x_2, y_2))$ for every $P(x_1, y_1)$ and $P(x_2, y_2)$ in the tessellation.

In a periodic tessellation the relative positions of the surrounding polyominoes must be the same for every $P(x, y) \in \tau$. It easily follows that periodic tessellations must be regular. But the same is true of regular tessellations. Thus, "regularity" exactly characterizes periodic tessellations.

Lemma 3.7. There exists a regular tessellation using P if and only if there exist an instance $P(x_0, y_0)$ and a complete surrounding $P(x_1, y_1), \dots, P(x_k, y_k)$ of it such that $\text{Seg}(P(x_i, y_i)) \subseteq \text{Seg}(P(x_0, y_0))$ ($i > 0$).

Proof.

(Note that the segmentations of the $P(x_i, y_i)$ referred to in the second part of the lemma will be partial for $i > 0$.)

\Rightarrow . Let τ be regular. Consider any $P(x_0, y_0) \in \tau$ and the polyominoes (of τ) completely surrounding it. The desired property now follows immediately from definition 3.6.

\Leftarrow . Suppose there exists a complete surrounding $P(x_1, y_1), \dots, P(x_k, y_k)$ of $P(x_0, y_0)$ such that $\text{Seg}(P(x_i, y_i)) \subseteq \text{Seg}(P(x_0, y_0))$ ($i > 0$). Because we can shift the entire configuration anywhere, one can surround $P(x_0, y_0)$ where-ever (x_0, y_0) is located. Observe that $|\text{Seg}(P(x_0, y_0))| = k$, by virtue of lemma 3.4. Consider any $P(x_i, y_i)$ ($1 \leq i \leq k$) and surround it by polyominoes just like $P(x_0, y_0)$. Because of the assumed property of the original partial segmentations the new polyominoes "grip" with the existing ones without conflict. Repeating this, every polyomino can be surrounded and the tessellation that results must be regular. \square

Our further efforts will focus on a better understanding of periodic tessellations.

4. Conditions for periodic tessellations.

To further analyze tessellations using a polyomino P we again inspect their embedding on the infinite grid. Given a (partial or total) tessellation τ , let G_τ be the graph of boundaries of the instances $P(x, y) \in \tau$. The nodes of G_τ will be the (grid-) points where at least three boundaries meet, its edges are the segments of unit-length grid-lines between any two nodes. (The precise definitions in case of partial τ are easy and left to the reader.) Clearly G_τ is a planar graph with nodes of degree 3 or 4. The length of an edge e will be the number of unit-length grid-lines it is composed of, denoted as $|e|$.

Definition 4.1 A three-node (four-node) is any grid-point g where three (four) non-overlapping instances of P meet. The branches of g are


the three (four) edges that meet in g (taken in consecutive order).

We will normally refer to the three- and four-nodes of some G_r with r total, but the definition applies to any local configuration of some $P(x_0, y_0)$ and a (partial or complete) surrounding just the same. In the latter case we speak of three-nodes (four-nodes) admitted by P . An edge will simply extend to either a node or a grid-point where two boundaries part.

Lemma 4.2 Suppose P admits a three-node g with branches T_1 , T_2 and T_3 . Then there exists with every $P(x_0, y_0)$ a partial surrounding $P(x_1, y_1), \dots, P(x_6, y_6)$ such that $P(x_i, y_i)$ also borders $P(x_{i+1}, y_{i+1})$ for $1 \leq i \leq 6$ (and $x_7 \equiv x_1, y_7 \equiv y_1$). The length of the partial segmentation induced by the partial surrounding is $2 \cdot |T_1| + 2 \cdot |T_2| + 2 \cdot |T_3|$.

Proof

Suppose P admits a three-node g as described. It means that we can choose any $P(x_0, y_0) \approx P_0$ we like and find two additional non-overlapping instances $P(x_1, y_1) \approx P_1$ and $P(x_2, y_2) \approx P_2$ that border it,



with the three of them meeting in a three-point g . The typical situation is shown in figure 6, where the handles of the various instances are indicated. Note that each instance borders the other two along precisely two consecutive branches of g . Let \vec{a} , \vec{b} and \vec{c} be the vectors pointing from one handle to the next, starting from and returning to (x_0, y_0) .

figure 6.

Consider an instance $P(x_3, y_3) \approx P_3$ located at the cell $(x_3, y_3) = (x_0, y_0) + \vec{b}$ ($= (x_2, y_2) + \vec{a}$), see figure 7. As P_3 has the same relative position to P_0 as P_2 has to P_1 , it does not overlap but does border P_0 . For a similar reason it does not overlap




figure 7

P_2 , but it does border it. The connectedness and isomorphism of the polyominoes are easily used to prove that P_3 cannot reach around P_0 and P_2 to intersect P_1 and thus P_3 does not overlap or even border this one. It follows that P_1, P_2 and P_3 are a correct "beginning" of the partial surrounding claimed in the lemma.

We use the mapping φ_{x_0, y_0} introduced in section 2 to further extend the partial surrounding we now have. Thus let $P(x_4, y_4) = \varphi_{x_0, y_0}(P(x_1, y_1)) \cong P_4$ and, going around in counter-clockwise order, let $P(x_5, y_5) = \varphi_{x_0, y_0}(P(x_2, y_2)) \cong P_5$ and also $P(x_6, y_6) = \varphi_{x_0, y_0}(P(x_3, y_3)) \cong P_6$. The situation so obtained is indicated in figure 8.




figure 8.

Because P_4 has the same relative position to P_3 as P_0 has to P_1 , it does not overlap but borders P_3 . Likewise P_6 borders P_1 but does not overlap it. Using lemma 2.8 and 2.9 it follows that P_4, P_5 and P_6 are all disjoint from P_0 as well (but they border it just like P_1, P_2 and P_3 do). We conclude that P_1 to P_6 form a partial surrounding of P_0 as claimed in the lemma.

Finally consider the partial segmentation induced on the boundary of $P(x_0, y_0)$. The argument of lemma 3.4 is easily used to argue that no other instances of P but P_1 through P_6 border $P(x_0, y_0)$. Each P_i ($1 \leq i \leq 6$) gives rise to exactly one segment along $B(x_0, y_0)$. Note that the junction of P_0, P_1 and P_2 at g (leading to segments T_1 and a "mirrored" T_3) is isomorphically repeated as a junction of P_4, P_5 and P_3 ("mirrored" segment T_1 , and T_2) and as a junction of P_5, P_6 and P_0 ("mirrored" segment T_2 , and T_3), if we simply observe relative positions and the

impact with the boundary of $P(x_0, y_0)$. This identifies the 6 segments along $B(x_0, y_0)$, with a total length equal to $|T_1| + |T_3| + |T_1| + |T_2| + |T_2| + |T_3| = 2 \cdot (|T_1| + |T_2| + |T_3|)$. \square

Lemma 4.3 Suppose P admits a four-node g with branches T_1, T_2, T_3 and T_4 (in, say, counter-clockwise order). Then there exists with every $P(x_0, y_0)$ a partial surrounding $P(x_1, y_1), \dots, P(x_4, y_4)$ such that $P(x_i, y_i)$ is "near" to $P(x_{i+1}, y_{i+1})$ for $1 \leq i \leq 4$ (and $x_5 \equiv x_1, y_5 \equiv y_1$). The length of the partial segmentation induced on $B(x_0, y_0)$ is equal to $2 \cdot |T_1| + 2 \cdot |T_4|$.

Proof

The argument is very similar to that of lemma 4.2, but several additional observations are required. Suppose P admits a four-point as described. It means that we can "fit" 4 instances $P(x_0, y_0) \cong P_0, P_1(x_1, y_1) \cong P_1, P(x', y') \cong P'$ and $P(x_2, y_2) \cong P_2$ and position this four some anywhere we like, with the boundaries meeting (tightly) at g .

It is assumed that P_0 is in the " T_1/T_4 corner" of g and that the other polyominoes are listed in counter-clockwise order.

Now drop P' (we will justify this later) and position an instance $P(x_{12}, y_{12}) \cong P_{12}$ in the T_2/T_3 corner of g instead, with $(x_{12}, y_{12}) = (x_2, y_2) + \vec{a}$ (see figure 9). Observing relative positions it will be clear that P_{12} borders both P_1 and P_2 , but it does not necessarily fit like P' in this corner. Thus it may actually remove the branches T_2 and T_3 as common boundaries,

although it need not "miss" both simultaneously.

Let $P(x_3, y_3) = \varphi_{x_0, y_0}(P(x_1, y_1)) \cong P_3$ and $P(x_4, y_4) = \varphi_{x_0, y_0}(P(x_2, y_2)) \cong P_4$




figure 9.

and consider the situation of the polyominoes now located around P_0 (see figure 9). It should be clear that P_1 to P_4 are a correct partial surrounding of P_0 . By reasoning from the topology of the polyominoes one can see that no instance of P can be fitted to P_0 in between any two consecutive P_i and P_{i+1} ($1 \leq i \leq 4$, with $5 \equiv 1$), simply because this is not possible in the "hole" left between P_0, P_1, P_2 and P_2 . Thus we have a partial surrounding as claimed in the lemma. (It is also clear now why we were interested in P_{12} instead of P' . If we look at the repeated instances of g obtained from equally positioned instances of P , then in three out of four cases there is a "fixed" polyomino located in the P' -corner in the same relative position as P_{12} , leaving no room to maintain P' .)

In case P_{12} does not "fill" the south-east corner at g , it will be clear from figure 9 that only 4 segments are induced along the boundary of $P(x_0, y_0)$. There are a T_1 and a "mirrored" T_4 , a "mirrored" T_1 and a T_4 respectively. The total length of the partial segmentation is equal to $|T_1| + |T_4| + |T_1| + |T_4| = 2 \cdot (|T_1| + |T_4|)$. If P_{12} does fill the south-east corner at g , then it gives rise to (possibly new) branches T_2' and T_3' . But these branches necessarily extend such that they coincide with the (mirrored version of) T_4 and T_1 branches at the next instance of g , while going around. This means that we actually get a complete surrounding of $P(x_0, y_0)$ but the segmentation will still have size 4 and length $2 \cdot (|T_1| + |T_4|)$. \square

(Note that the proof explains the lemma's notion of the surrounding polyominoes being "near" to one another.)

Corollary 4.4. Let τ be a periodic tessellation using P . Then either (a) every node of G_τ is a three-node and every $P(x, y) \in \tau$ is completely surrounded by 6 other instances of P , or (b) every node of G_τ is a four-node and every $P(x, y) \in \tau$ is completely surrounded by 4 other

instances of P

Proof.

The proof follows by more closely examining the arguments in lemma 4.2 and lemma 4.3. Let τ be a periodic tessellation of the plane using P . Consider an arbitrary node g of G_τ . Clearly g is either a three-node or a four-node.

In case g is a three-node there are three instances P_0, P_1 and $P_2 \in \tau$ that meet at g as specified in the beginning of the proof of lemma 4.2. This identifies three relative positions (the "vectors" $\pm \vec{a}$, $\pm \vec{b}$ and $\pm \vec{c}$ in figure 6) which, because of the periodicity of τ , must always lead from one instance of P in τ to another one that necessarily also belongs to τ . It is easily verified that for this reason each of the polyominoes P_3 to P_6 constructed in the proof of lemma 4.2 can be justified as a polyomino actually belonging to τ . As any gap between two consecutive P_i 's ($i > 0$) and the boundary of P_0 would be too small to fit in another instance of P and yet τ is total, it follows that the polyominoes P_1 to P_6 must be a complete surrounding of P_0 . Again arguing from the assumed periodicity of τ , this means that every polyomino in τ is surrounded likewise in exactly the same manner. In particular, each node of G_τ necessarily appears as a three-node.

In case g is a four-node there are four instances P_0, P_1, P' and $P_2 \in \tau$ that meet at g as specified in the beginning of the proof of lemma 4.3. Because of the assumed periodicity of τ it is easily seen that the "third" polyomino P' must actually be the polyomino P_2 identified in the proof of lemma 4.3 for other reasons, which means in particular that P_2 nicely fills the "south-east" corner at g (see figure 9). One can now argue as before that the polyominoes P_1 to P_4 constructed in the proof of lemma 4.3 must all belong to τ and form a complete surrounding of P_0 . Since τ is periodic, it follows that every polyomino in τ is surrounded in exactly the same manner and (hence) that every node of G_τ is a four-node. \square

Corollary 4.5. Let τ be a periodic tessellation of the plane using P . Then the underlying lattice of points (x, y) such that $P(x, y) \in \tau$ has a basis of cardinality equal to 2.

Proof

It is clear that a basis must contain more than one vector. We shall argue that always two vectors suffice for a basis. Let τ be periodic. By corollary 4.4 we know that τ must either consist of (a) polyominoes that are all surrounded in exactly the same manner by 6 other instances, or of (b) polyominoes that are all surrounded likewise by 4 other instances of P . The lattice we are after is generated by the vectors from which all relative positions within τ can be obtained by "iteration". It follows from the proof of lemma 4.2 (viz. figure 8) that two vectors will do in case (a). (Note in figure 8 that e.g. \vec{c} is integrally dependent on \vec{a} and \vec{b} and that \vec{a} and \vec{b} "generate" the entire tessellation.) By the same token it follows from the proof of lemma 4.3 (viz. figure 9) that two vectors suffice in case (b) as well. \square

Corollary 4.6. Let τ be an arbitrary (partial or total) tessellation of the plane using P . For all three-nodes of G_τ we have: $|T_1| + |T_2| + |T_3| \leq \frac{1}{2}|B|$ and for all four-nodes of G_τ we have: $|T_1| + |T_2| + |T_3| + |T_4| \leq |B|$, where T_1, T_2 and T_3 (and T_4) are the branches of the three-node (four-node) in question and B is the boundary of P . (Note that always $|B|$ is even.)

Proof

Consider any three-node g of G_τ . By lemma 4.2 every $P(x_0, y_0)$ can be partially surrounded by a set of polyominoes (not necessarily from τ) that induce a partial segmentation of $B(x_0, y_0)$ of length 2. ($|T_1| + |T_2| + |T_3|$) Hence $|T_1| + |T_2| + |T_3| \leq \frac{1}{2}|B|$. Likewise it follows from lemma 4.3 that for every four-node g of G_τ : $|T_1| + |T_4| \leq \frac{1}{2}|B|$, and by a symmetric argument that $|T_2| + |T_3| \leq \frac{1}{2}|B|$. Hence $|T_1| + |T_2| + |T_3| + |T_4| \leq |B|$ for every four-node. \square

The final result of this section is important because it establishes a local condition that is necessary and sufficient for the existence of a periodic tessellation.

Theorem 4.7. There exists a periodic tessellation of the plane using the polyomino P with boundary B if and only if

(*) P admits a three-node g with branches T_1, T_2 and T_3 such that $|T_1| + |T_2| + |T_3| = \frac{1}{2}|B|$, or

(**) P admits a four-node g with branches T_1, T_2, T_3 and T_4 such that $|T_1| + |T_2| + |T_3| + |T_4| = |B|$.

Proof.

\Rightarrow . Let τ be a periodic tessellation of the plane using P . By corollary 4.4 we know that G_τ consists of either three-nodes or four-nodes. If G_τ consists of three-nodes (and, hence, P admits a three-node) then the argument in corollary 4.4 shows that the surrounding of any $P(x_0, y_0)$ as constructed in the proof of lemma 4.2 must be complete. It follows that $2 \cdot (|T_1| + |T_2| + |T_3|) = |B|$, or: $|T_1| + |T_2| + |T_3| = \frac{1}{2}|B|$, for any three-node in G_τ . If G_τ consists of four-nodes (and, hence, P admits a four-node) then the argument in corollary 4.4 shows likewise that the surroundings constructed in the proof of lemma 4.3 must be complete. Thus $|T_1| + |T_4| = \frac{1}{2}|B|$ and, by symmetry, $|T_2| + |T_3| = \frac{1}{2}|B|$. It follows that $|T_1| + |T_2| + |T_3| + |T_4| = |B|$ in this case.

\Leftarrow . The converse follows once again by closely examining the proofs of lemma 4.2 and lemma 4.3. Suppose P satisfies (*). Observing the length of the induced segmentation, it follows that the surrounding of $P(x_0, y_0)$ constructed in the proof of lemma 4.2 (starting from the three instances of P as they meet at the given three-node g) necessarily is a complete surrounding. Drawing in the 6 isomorphically positioned instances of g around the boundary of $P_0 = P(x_0, y_0)$ in figure 8, it is easily observed that this forces the partial segmentation of each of the $P(x_i, y_i)$ ($1 \leq i \leq 6$) to satisfy the condition of lemma 3.7. Hence, the sur-

rounding can be repeated in all instances of P surrounding $P(x_0, y_0)$ and be extended to a regular (hence: periodic) tessellation of the entire plane.

If P satisfies $(**)$ rather than $(*)$, then a similar argument carries through based on the construction of a surrounding in the proof of lemma 4.3 and shows with lemma 3.7 that again a periodic tessellation can be obtained using P . \square

5 Obtaining periodic tessellations from arbitrary tessellations: a proof of Shapiro's conjecture.

The detailed analyses of the preceding sections will now be used to settle Shapiro's conjecture (cf. section 1) and prove that whenever there is a tessellation of the plane using the polyomino P , there must exist a periodic tessellation using P . Let τ be an arbitrary tessellation of the plane using P . The key idea is a detailed analysis of the "grid"-graph G_τ . Imagine that each edge of G_τ is cut into two equal halves and that the length of each half is charged to the appropriate end-point.

Definition 5.1 The support of a node $g \in G_\tau$, denoted as: $\text{Sup}_\tau(g)$ or just as: $\text{Sup}(g)$ when τ is understood, is equal to the total charge thus accumulated at g , i.e., $\text{Sup}_\tau(g) = \frac{1}{2} \sum |e|$, with the summation extending over all (3 or 4) edges incident to g .

(The reason for looking at the edge-lengths in G_τ should be clear, for the edges are the "branches" of the three-nodes and four-nodes in the graph. The halving is only introduced to simplify later accounting procedures and to avoid that entire edges are counted twice: once at every end-point.) The proof of Shapiro's conjecture heavily relies on the criteria for periodic tessellations in theorem 4.7 and uses the following surprising fact.

Lemma 5.2. In every tessellation of the plane using P there exists a three-node as in (*) or a four-node as in (**).

Proof.

Let N be sufficiently large and consider an arbitrary $N \times N$ window on G_P . Let G'_P be the (planar) graph of nodes and edges obtained by only considering the polyominoes of τ that are strictly located within the window. Clearly G'_P is a connected and finite section of G_P , with a contour C bounding the graph from its "exterior". Among the nodes along C there are likely to be many that are remnants of three-nodes or four-nodes that lost at least one branch (because it was sticking out of the window). Let K be the number of polyominoes of τ strictly contained in the window and (hence) spanning G'_P . Define factors ε (depending on τ , K and N) such that

$\varepsilon_1 \cdot K$ = the number of three-nodes along C that have degree 2 in G'_P ,

$\varepsilon'_1 \cdot K$ = the number of three-nodes along C that (still) have degree 3 in G'_P ,

$\varepsilon_2 \cdot K$ = the number of four-nodes along C that have degree 2 in G'_P ,

$\varepsilon'_2 \cdot K$ = the number of four-nodes along C that have degree 3 in G'_P ,

$\varepsilon''_2 \cdot K$ = the number of four-nodes along C that (still) have degree 4 in G'_P .

Claim 5.2.1. For N sufficiently large each factor ε is less than $\frac{1}{2 \cdot |B|}$, where $|B|$ is the size of the boundary of P .

Proof.

Note that the size of the polyomino is fixed. Thus K increases quadratically in N for $N \rightarrow \infty$. On the other hand, it is easily seen that $|C|$ increases at most linearly in N . Thus the number of nodes along

C can be made less than any factor times K , by choosing N sufficiently large.

For a further analysis of G'_r we define the following values. In each case an expression is obtained either by direct reasoning or by carefully accounting the "contributions" to three-nodes ($\frac{1}{3}$ from each incident polyomino), four-nodes ($\frac{1}{4}$ from each incident polyomino) and edges ($\frac{1}{2}$ from the "initial" node in the clockwise ordering of B):

α_{ij} = the number of polyominoes within the window (hence in G'_r) that have i three-nodes and j four-nodes on their boundary,

α_{i*} = the number of polyominoes etc. that have i three-nodes on their boundary
 $= \sum_j \alpha_{ij}$,

α_{*j} = the number of polyominoes etc. that have j four-nodes on their boundary
 $= \sum_i \alpha_{ij}$,

t = the number of three-nodes within the window (hence in G'_r)
 $= \sum_i \frac{i}{3} \cdot \alpha_{i*} + \frac{2}{3} \cdot \varepsilon_1 K + \frac{1}{3} \cdot \varepsilon'_1 K$,

f = the number of four-nodes within the window (hence in G'_r)
 $= \sum_j \frac{j}{4} \cdot \alpha_{*j} + \frac{3}{4} \cdot \varepsilon_2 K + \frac{2}{4} \cdot \varepsilon'_2 K + \frac{1}{4} \cdot \varepsilon''_2 K$,

n = the total number of nodes within the window (hence in G'_r)

$$= t + f,$$

$$\begin{aligned} e &= \text{the total number of edges (branches) within the window} \\ &\quad \text{(hence in } G_r') \\ &= \sum_{i,j} \frac{i+j}{2} \cdot \alpha_{ij} + \frac{1}{2} \cdot \varepsilon_1 K + \frac{1}{2} \cdot \varepsilon_1' K + \frac{1}{2} \cdot \varepsilon_2 K + \frac{1}{2} \cdot \varepsilon_2' K + \frac{1}{2} \cdot \varepsilon_2'' K, \end{aligned}$$

$$\begin{aligned} p &= \text{the total number of parts (faces) in which the plane is} \\ &\quad \text{divided by } G_r' \\ &= K + 1. \end{aligned}$$

$$\text{Note that } \sum_i \alpha_{i*} = \sum_j \alpha_{*j} = K.$$

$$\text{Claim 5.2.2 } f = -\frac{1}{2}t + \frac{1}{2}\varepsilon_1 K + \varepsilon_2 K + \frac{1}{2}\varepsilon_2' K + K - 1$$

Proof.

Since G_r' is planar we can apply Euler's well-known polyeder formula: $n + p = e + 2$. Substituting the expressions for n , p and e (etc.) we obtain

$$t + f + K + 1 = e + 2 \Rightarrow$$

$$\Rightarrow \sum_i \frac{i}{3} \cdot \alpha_{i*} + \frac{2}{3} \varepsilon_1 K + \frac{1}{3} \varepsilon_1' K + \sum_j \frac{j}{4} \cdot \alpha_{*j} + \frac{3}{4} \varepsilon_2 K + \frac{1}{2} \varepsilon_2' K + \frac{1}{4} \varepsilon_2'' K +$$

$$+ K + 1 = \sum_{i,j} \frac{i+j}{2} \cdot \alpha_{ij} + \frac{1}{2} (\varepsilon_1 + \varepsilon_1' + \varepsilon_2 + \varepsilon_2' + \varepsilon_2'') K + 2 \Rightarrow$$

$$\Rightarrow \sum_i \frac{i}{6} \cdot \alpha_{i*} + \sum_j \frac{j}{4} \cdot \alpha_{*j} = \frac{1}{6} \varepsilon_1 K - \frac{1}{6} \varepsilon_1' K + \frac{1}{4} \varepsilon_2 K - \frac{1}{4} \varepsilon_2'' K + K - 1.$$

Multiplying the latter equation by 2, the left-hand side contains terms that remind of $t + 2f$. Straightforward manipulation shows:

$$t + 2f = \left(\sum_i \frac{i}{3} \cdot \alpha_{i*} + \sum_j \frac{j}{2} \cdot \alpha_{*j} \right) + \frac{2}{3} \varepsilon_1 K + \frac{1}{3} \varepsilon_1' K + \frac{3}{2} \varepsilon_2 K +$$

$$\varepsilon'_2 K + \frac{1}{2} \varepsilon_2'' K = \varepsilon_1 K + 2 \varepsilon_2 K + \varepsilon_2' K + 2K - 2,$$

and the expression claimed for f easily follows.

Suppose by way of contradiction that γ (hence G'_γ) does not contain any three-nodes satisfying (*) nor any four-nodes satisfying (**). By corollary 4.6 this means that for every three-node g : $|T_1| + |T_2| + |T_3| \leq \frac{1}{2}|B| - 1$ (note that $|B|$ is always even) and for every four-node g : $|T_1| + |T_2| + |T_3| + |T_4| \leq |B| - 1$, where T_i etc. are the branches of the node. It means that for every three-node g : $\text{Sup}(g) \leq \frac{1}{4}|B| - \frac{1}{2}$ and for every four-node g : $\text{Sup}(g) \leq \frac{1}{2}|B| - \frac{1}{2}$. Let L be the total edge length of G'_γ . Note that $L < \sum_{g \in G'_\gamma} \text{Sup}(g)$. (The $<$ -sign holds because there is at least one node along the contour of G'_γ that "lost" a branch which is still accounted for in its support.) Using the expression for f from claim 5.2.2 we can bound L as follows:

$$L < \sum_{\substack{g \in G'_\gamma \\ g \text{ three-node}}} \text{Sup}(g) + \sum_{\substack{g \in G'_\gamma \\ g \text{ four-node}}} \text{Sup}(g) \leq$$

$$\leq t \cdot \left(\frac{1}{4}|B| - \frac{1}{2}\right) + f \cdot \left(\frac{1}{2}|B| - \frac{1}{2}\right) <$$

$$< t \cdot \left(\frac{1}{4}|B| - \frac{1}{2}\right) + \left(-\frac{1}{2}t + \frac{1}{2}\varepsilon_1 K + \varepsilon_2 K + \frac{1}{2}\varepsilon_2' K + K\right) \left(\frac{1}{2}|B| - \frac{1}{2}\right) =$$

$$= -\frac{1}{4}t + \left(\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon_2'\right) K \cdot |B| + \frac{1}{2}K \cdot |B| -$$

$$- \left(\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon_2'\right) K - \frac{1}{2}K \leq$$

$$\leq \frac{1}{2}K \cdot |B| + \left\{ \left(\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon_2'\right) \cdot |B| - \frac{1}{2} \right\} K.$$

As N was chosen sufficiently large, it easily follows from claim 5.2.1 that $\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon_2' < \frac{1}{2|B|}$. Thus our estimate on L re-

duces to

$$L < \frac{1}{2} K \cdot |B|$$

On the other hand, if we let each of the K polyominoes in G_r' contribute one half of every bounding edge (which indeed properly divides the length of every edge over its two bordering polyominoes) then it easily follows that

$$L \geq K \cdot \frac{1}{2} |B| = \frac{1}{2} K \cdot |B|,$$

a contradiction. We conclude that G_r' (hence r) must contain a three-node satisfying (*) or a four-node satisfying (**). \square

Corollary 5.3 In every tessellation of the plane using P there exist infinitely many three-nodes as in (*) or infinitely many four-nodes as in (**).

Proof

The proof of lemma 5.2 shows that for N sufficiently large there is a three-node satisfying (*) or a four-node satisfying (**) in every $N \times N$ window on G_r . The argument is easily completed from here. \square

Theorem 5.4 Let P be a polyomino. If it is possible at all to tessellate the plane using P , then there exists a periodic tessellation of the plane using P .

Proof

The result follows at once from lemma 5.2 and theorem 4.7. (Note the additional observations for periodic tessellations in section 4.) \square

6. Final comments

The study of plane tessellations was motivated from the theory of

data organisation for SIMD-machines. It was argued in section 1 (see also Shapiro [3]) that only periodic tessellations are likely to be of practical interest, because of the simple scheme it implies for storing and retrieving data items. Thus the proof of Shapiro's conjecture has a clear interpretation in practice. It is important to note that the result of theorem 5.4 is entirely "effective". First of all, whenever a tessellation using a polyomino P is given in some computable manner, then the proof of lemma 5.2 shows that one can actually compute (by inspecting an arbitrary $N \times N$ window over the tessellation) a three-node satisfying (*) or a four-node satisfying (**). Secondly, the results underlying theorem 4.7 show that there is an effective way to determine the two generating vectors (i.e. the basis) of the lattice of points where the polyominoes P in a periodic tessellation must be placed. Clearly, given theorem 5.4 only the second observation is important, for one can always determine by "trying" whether P admits a three-node or a four-node with the desired property.

Theorem 6.1 Given a polyomino P , there exists an algorithm that is polynomial in the size of P to decide whether P can tessellate the plane or not.

Proof

By theorem 5.4 we only need to test the conditions for a periodic tessellation using P as expressed in theorem 4.7. Take an instance of P and test at every (grid-) point along the boundary whether 3 or 4 instances can be fitted without overlap and satisfying the length condition for the branches at the node so created. There are only polynomially many cases to consider, and each test for overlap and the condition for the branches takes at most $O(|B|^2)$ (hence: polynomial) time as well. (It is not hard to show that the algorithm can, in fact, be made polynomial in $|B|$ altogether.) \square

(The impact of polynomial time algorithms in the theory of computation is explicated in e.g. Aho, Hopcroft and Ullman [1].)

Severe problems arise if we attempt to generalise theorem 5.4 and e.g. relax the condition that P be a polyomino. The template T shown in figure 10 provides an example that Shapiro's conjecture




figure 10

does not remain valid if we do so. It is easily verified that T tessellates the plane. But the following argument shows that it cannot tessellate the plane periodically. Name the two components f ("first") and s ("second"). Whenever we try to place a second instance of T to fill the narrow gorge between f and s , we get

either an f on an f or an s on an s , and it is easily seen that this cannot be repeated without conflict. Yet there may be a way to relax the condition of periodicity in such a manner that a suitable modification of Shapiro's conjecture remains valid.

Other problems arise if we no longer insist that all instances of the polyomino in a tessellation must have a fixed orientation. One can imagine that a finite set of instances of the polyomino would have to exist such that the plane can be tessellated periodically with this set. At present no results are known in this direction.

7. References.

[1] Aho, A.V., J.E. Hopcroft and J.D. Ullman, The design and analysis of computer algorithms, Addison-Wesley Publ. Comp., Reading, Mass., 1974

[2] Budnik, P. and D.J. Kuck, The organisation and use of parallel memories, IEEE Trans. Comput. C-20 (1971) 1566-1569

[3] Shapiro, H. D., Theoretical limitations on the efficient use of parallel memories, IEEE Trans. Comput. C-27 (1978) 421-428.

[4] Thurber, K. J., Large scale computer architecture: parallel and associative processors, Hayden Book Comp, Rochelle Park NJ, 1976.