

Arbitrary versus Periodic Storage Schemes and Tessellations of the Plane Using One Type of Polyomino

H. A. G. WIJSHOFF AND J. VAN LEEUWEN

*Department of Computer Science, University of Utrecht,
P.O. Box 80.012, 3508 TA Utrecht, The Netherlands*

Given N distinct memory modules, the elements of an (infinite) array in storage are distributed 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 Shapiro (*IEEE Trans. Comput.* **C-27** (1978), 421-428) proved that there exists a valid skewing scheme for a template T if and only if T tessellates the plane. A conjecture of Shapiro is settled and it is proved that for polyominoes P a valid skewing scheme exists if and only if there exists a valid periodic skewing scheme. (Periodicity implies a rapid technique to locate data elements.) The proof shows that when a polyomino P tessellates the plane without rotations or reflections, then it can tessellate the plane periodically, i.e., with the instances of P arranged in a lattice. It is also proved that there is a polynomial time algorithm to decide whether a polyomino tessellates the plane, assuming the polyominoes in the tessellation should all have an equal orientation. © 1984 Academic Press, Inc.

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, 1976). The characterising feature of these machines is the availability of a multitude of arithmetic units and memory modules that can operate independently in parallel. Clearly the effectiveness of these machines depends to a large extent on being able to store the data elements in the available memories such that memory conflicts are avoided whenever data are fetched.

About 1970 Budnik and Kuck (1971) pointed out that nontrivial problems arise if any set of N elements from a 2-dimensional array, arranged according to some common pattern or template, must be accessed in one cycle without conflict. A data template T consists of a fixed set of N locations relative to a designated base 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 T if it provides for the conflict-free parallel access to the data in any instance of T . Clearly there does not always exist a valid skewing scheme for the templates at hand, but in many interesting cases there does.

In 1975 Shapiro (1978) 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 $(-\infty: \infty; -\infty: \infty)$. Second, he proved that there is a valid skewing scheme for T if and only if T (as a combinatorial structure) tessellates the 2-dimensional plane. As the argument is important, we briefly digress and include our simplified proof of this fact.

WARNING. When we speak of tessellations of the plane using a template of some sort we shall require throughout this paper that the templates in a tessellation all have equal orientation, i.e., we do not consider rotations and reflections of the objects when discussing tessellations unless explicitly stated otherwise.

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 h_1x and 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 Fig. 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 established along similar lines. ■

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 same template.*




FIGURE 1

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, $\mathbf{h}_1\mathbf{x}$ and $\mathbf{h}_2\mathbf{y}$ are identical. Let h_4 be the “fourth” corner of the parallelogram spanned by h_1 , x , and h_3 (see Fig. 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 well and that as a result $\mathbf{h}_4\mathbf{h}_1 = \mathbf{h}_3\mathbf{x}$ and $\mathbf{h}_4\mathbf{h}_2 = \mathbf{h}_3\mathbf{y}$. 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 . ■

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 $s: (-\infty: \infty; -\infty: \infty) \rightarrow [1 \cdots N]$.)

(\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 N possible instances T_i of T that cover q and let their handles be located in cells p_i ($1 \leq i \leq N$). 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.




FIGURE 2

(\Leftarrow) Let T tessellate the plane. Number the cells of T from 1 to N 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_1 of T . If T_1 coincides with an instance from the tessellation, then its cells are trivially assigned to different memories. Otherwise Lemma 1.2 shows that T_1 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_1 are assigned different numbers even now. ■

General skewing schemes are not necessarily of use in practice. There is no guarantee that a skewing scheme s is finitely encoded or indeed recursive. This led Shapiro (1978, Sect. IV) to consider a number of constraints to force a skewing scheme to be finitely represented in computer memory. The weakest condition which he introduced was termed periodicity. We shall use periodicity in a somewhat stronger sense, namely, to mean that the instances of T that tessellate the plane have their handle in the points of a (2-dimensional) lattice. Periodicity implies a simple method to store and retrieve data elements quickly (Wijshoff and van Leeuwen (1983)).

In this paper we settle an important conjecture of Shapiro (1978) and prove that for templates that have the shape of a polyomino there exists a valid skewing scheme if and only if there exists a valid periodic skewing scheme. The proof relies on the geometric interpretation of the problem. We show that when a polyomino of size N tessellates the plane, then it can tessellate the plane periodically, i.e., with its instances arranged according to an effectively computable lattice. 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 some preliminary results. In Section 3 we define tessellations and derive some properties of these tessellations. In Section 4 we derive an operational notion of periodicity and 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 is developed, 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 time algorithm for the existence of a tessellation with a given polyomino and offer some final comments to identify the significance of the results within the theory of geometric packing and covering.

2. DEFINITIONS AND PRELIMINARY RESULTS

All notions introduced in this section are rather straightforward and are presented without much commentary. Definitions 2.1 and 2.3 are from Shapiro (1978).

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 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 with no “holes” (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 $P(x, y)$.

DEFINITION 2.4. The relative position π of cells (x_1, y_1) and (x_2, y_2) is the “bi-directional” vector $\mathbf{r} = \pm(x_2 - x_1, y_2 - y_1)$. 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 \mathbf{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 for some i and j : $(a_i, b_i) = (a_j, b_j) + (x_2 - x_1, y_2 - y_1)$ or, equivalently, $(a_j, b_j) = (a_i, b_i) + (x_1 - x_2, y_1 - y_2)$. ■




FIGURE 3

Let $P(x_0, y_0)$ be a fixed instance of P . With every polyomino $P(x, y)$ there 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 Fig. 3) that the buddy of $P(x, y)$ is symmetrically located at the “opposite” side of the polyomino $P(x_0, y_0)$. 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 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 would suffice, but the argument is best seen from a geometric interpretation (see Fig. 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. ■

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

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.




FIGURE 4




FIGURE 5

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 Fig. 5, where for simplicity 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. ■

Lemma 2.9 is a special case of some results of Levi (1934). 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 analyzing disjoint placements of polyominoes around an instance $P(x_0, y_0)$.

3. TESSELLATIONS

Tessellations (or “tilings”) are a familiar subject in mathematics. Definition 3.1 is from Shapiro (1978) (although we have added the dis-

inction between partial and total tessellations). We require the objects in a tessellation to have an equal orientation.

DEFINITION 3.1. A partial tessellation (using P) is any collection of nonoverlapping 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 adjective is added we assume a tessellation to be total. If there exists a total tessellation using P then we say that P tessellates the plane.

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 2-dimensional lattice.

A 2-dimensional lattice is the set of points obtained as the integer linear combinations $\lambda x + \mu y$ of two integral vectors x and y (the basis of the lattice). In the following we assume that polyominoes are sets of cells on the 2-dimensional grid. (Recall that polyominoes have no holes.) We say that two polyominoes border each other if they have at least one gridpoint in common.

DEFINITION 3.3. The boundary B of the (embedded) polyomino P is the set of gridlines 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 by (x, y) .

We number the gridlines 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 a numbering $r_0(x, y), r_1(x, y), \dots$ of $B(x, y)$. Note that numbers like $r_i(x, y)$ are merely names of gridlines with respect to some $P(x, y)$.

Notice that a partial tessellation may contain holes, e.g., cells which are not covered by any instance of this tessellation. The “boundary” of any such hole is formed by the parting and rejoining boundaries of the enclosing instances. The entire boundary will be called the *interior boundary* I of the instances with respect to the hole. Whenever a collection of instances forms no hole then the *exterior boundary* E of that collection is defined as the collection of gridlines belonging to exactly one instance. The length of I (or E , resp.) will be denoted as $|I|$ ($|E|$). In the remaining part of this section we shall prove that when a hole is formed by two or three instances then this hole cannot be covered by further instances of P .

LEMMA 3.4. *Given three nonoverlapping instances $P(x_0, y_0), P(x_1, y_1),$*

and $P(x_2, y_2)$, with $P(x_i, y_i)$ bordering $P(x_{i \pm 1 \pmod 3}, y_{i \pm 1 \pmod 3})$ for $0 \leq i \leq 2$. Let $\mathbf{a} = (x_1 - x_0, y_1 - y_0)$, $\mathbf{b} = (x_2 - x_0, y_2 - y_0)$, and $\mathbf{c} = (x_2 - x_1, y_2 - y_1)$. Then for all (x, y) the instances $P(x, y)$, $P((x, y) \pm \mathbf{a})$, $P((x, y) \pm \mathbf{b})$, and $P((x, y) \pm \mathbf{c})$ do not overlap each other (see Fig. 6).

Proof. Let $P(x, y) \approx P_0$, $P((x, y) + \mathbf{a}) \approx P_1$, $P((x, y) + \mathbf{b}) \approx P_2$, $P((x, y) + \mathbf{c}) \approx P_3$, $P((x, y) - \mathbf{a}) \approx P_4$, $P((x, y) - \mathbf{b}) \approx P_5$, and $P((x, y) - \mathbf{c}) \approx P_6$.

Note that P_0 does not overlap with any P_i . For, P_1 e.g., we can make the following observations:

- (1) P_1 and P_0 do not overlap because $P(x_1, y_1)$ and $P(x_0, y_0)$ do not overlap.
- (2) P_1 and P_2 do not overlap because $P(x_1, y_1)$ and $P(x_2, y_2)$ do not overlap.
- (3) P_1 and P_6 do not overlap because $P(x_2, y_2)$ and $P(x_0, y_0)$ do not overlap.
- (4) P_1 and P_4 do not overlap because of 1 and Lemma 2.9.

Because the polyominoes are (rook-wise) connected and have the same form and orientation, we also observe:

- (5) P_1 and P_3 do not overlap, and
- (6) P_1 and P_5 do not overlap.

The reason for this is that P_1 cannot reach around P_2 and P_0 (or P_0 and




FIGURE 6

P_6 , respectively) to intersect P_3 or P_5 . Note that P_2 and P_0 (and P_0 and P_6 , respectively) border each other.

By symmetry similar observations hold for P_2, P_3, \dots, P_6 . Hence P_0 through P_6 do not overlap each other. (Note that P_1 borders P_2 borders P_3 borders \dots borders P_6 borders P_1 . So P_0 is totally enclosed by $P_1 \dots P_6$). ■

The next two lemmas will place restrictions on the sizes of the interior and exterior boundary of a collection of instances of P .

LEMMA 3.5. *Given a polyomino P with boundary B ,*

(i) *Let two instances of P form a hole h and have a gridline in common. Then the size of the interior boundary I of these instances with respect to h is strictly less than $|B|$.*

(ii) *Let three instances of P form a hole h and at least two of them have a gridline in common. Then the size of the interior boundary of these instances with respect to h is strictly less than $|B|$.*

Proof. We shall first give a proof of the second part of the lemma, because it is the more difficult one.

(ii) Let $P(x_0, y_0) \approx P_0$, $P(x_1, y_1) \approx P_1$, and $P(x_2, y_2) \approx P_2$ be three nonoverlapping instances, which form a hole h . Furthermore two of them have a gridline in common. Consider the set S of gridlines belonging to the interior boundary of $\{P_0, P_1, P_2\}$ with respect to hole h or to at least 2 instances (see Fig. 7). Clearly $|S| \geq |I| + 1$. We will prove that S consists of




FIGURE 7

gridlines belonging to disjoint parts of the boundary B . By this we mean that if $r_k(x_i, y_i), r_l(x_j, y_j) \in S$ and $r_k(x_i, y_i) \neq r_l(x_j, y_j)$ then $k \neq l$.

Whenever $r_k(x_i, y_i), r_l(x_i, y_i) \in S$ and $r_k(x_i, y_i) \neq r_l(x_i, y_i)$ then it trivially follows that k cannot be equal to l . Suppose now $r_k(x_i, y_i), r_l(x_j, y_j) \in S$, $i \neq j$, $r_k(x_i, y_i) \neq r_l(x_j, y_j)$, and $k = l$. Let $i = 0$ and $j = l$. Because of symmetry the other cases can be handled analogously.

From Lemma 3.4 follows that $P_0, P_1, P_2, P(x_0 + x_2 - x_1, y_0 + y_2 - y_1) \approx P_3$, and $P(2x_0 - x_1, 2y_0 - y_1) \approx P_4$ do not overlap each other. Note that the relative positions of P_2, P_0 , and P_1 to each other are the same as those of P_3, P_4 , and P_0 to each other. So we can extend both P_2 and P_3 to P'_2 and P'_3 in such a way that P'_2 exactly covers P_2 and the hole h and P'_3 exactly covers P_3 and the corresponding hole (see Fig. 8). Furthermore P_0, P_1, P'_2, P'_3 , and P_4 do not overlap each other. We assumed that $k = l$. Thus both P'_2 and P'_3 border P_0 along the same gridline $r_k(x_0, y_0)$ (if $r_k(x_1, y_1) \notin B(x_0, y_0)$) or both P_1 and P_4 border P_0 along a gridline $r_{k'}(x_0, y_0)$, with $r_{k'}(x_0, y_0) = r_k(x_1, y_1)$ (if $r_k(x_1, y_1) \in B(x_0, y_0)$). This however contradicts the fact that P_0, P_1, P'_2, P'_3 , and P_4 do not overlap each other.

It follows now that the set S consists of gridlines belonging to disjoint parts of the boundary B . So $|S| \leq |B|$. Together with $|S| \geq |I| + 1$ this gives $|I| < |B|$.

(i) Let $P(x_0, y_0) \approx P_0$ and $P(x_1, y_1) \approx P_1$ be two nonoverlapping instances which form a hole h and have a gridline in common. Consider again the set S of gridlines belonging to the interior boundary of P_0 and P_1 with respect to h or to both $B(x_0, y_0)$ and $B(x_1, y_1)$. Then because of Lemma 2.9, P_0, P_1 , and $P_2 \approx \varphi_{x_1, y_1}(P(x_0, y_0))$ do not overlap each other.




FIGURE 8

Extend both P_0 and P_1 to P'_0 and P'_1 in the same way as in (ii). Then again P'_0 , P'_1 , and P_2 do not overlap each other. Similar to the proof of (ii) a contradiction occurs if we assume that two gridlines of S correspond to the same gridline of B . ■

LEMMA 3.6. *Given a polyomino P with boundary B . Let E be the exterior boundary of any nonempty finite collection of instances of P , which do not overlap each other and form no hole. Then $|E| \geq |B|$.*

*Proof*¹. Let \mathcal{C} be a nonempty finite collection of instances of P , which do not overlap each other and form no hole. Consider an arbitrary instance $P(x_0, y_0) \approx P_0 \in \mathcal{C}$. We will prove that every element of the boundary B_0 of P_0 induces a corresponding element of E .

Let $r_k(x_0, y_0) \in B_0$. We have to show that there exists a $P(x, y) \in \mathcal{C}$ such that $r_k(x, y) \in E$. We will prove this by inspecting the boundaries of the instances of P .

For every $P(x_i, y_i) \approx P_i \in \mathcal{C}$ let C_i be its convex hull. Let the straight line segments on the boundary of each C_i be $\alpha_i^1, \alpha_i^2, \dots, \alpha_i^n$. With every α_i^m we can identify the set $R_i^m = \{r_l(x_i, y_i), r_{l+1}(x_i, y_i), \dots, r_{l+s}(x_i, y_i)\}$, such that $r_l(x_i, y_i)$ and $r_{l+s}(x_i, y_i)$ meet the two endpoints of α_i^m (see Fig. 9). Notice that $\sum_{j=1}^n |R_{ij}^j| = |B|$, with i_j arbitrary. The following claim actually ends the proof.

Claim 3.6.1. $\forall j \exists i R_i^j \subseteq E$.

Proof. Let j be given. Consider an arbitrary instance P_i with its convex hull C_i . Take a line l perpendicular to the line segment α_i^j and project all the line segments of the C_k 's on this line l . Because each C_k is convex it follows




FIGURE 9

¹ The argument is due to N. van Diepen.

that the line segments $\alpha_k^1, \alpha_k^2, \dots, \alpha_k^n$ have to be in one of the two halfplanes obtained by cutting the plane along α_k^j . From this argument follows that there has to be an outermost point p on the line l such that p is an element of the projection and, in particular, p is the projection of an $\alpha_{k_0}^j$.

Suppose that $R_{k_0}^j \not\subseteq E$. Then there has to be a P_k ($k \neq k_0$) such that $B_k \cap R_{k_0}^j \neq \emptyset$. Consider C_k . C_k cannot totally cover P_{k_0} , otherwise $k = k_0$. So the area of C_k is strictly greater than the area of $C_{k_0} - P_{k_0}$. Thus C_k has to stick out of $C_{k_0} - P_{k_0}$. This however contradicts the fact that p is the outermost point of the projection. ■

THEOREM 3.7. (i) *If $P(x_1, y_1)$ and $P(x_2, y_2)$ form a hole and have a gridline in common, then there exists no total tessellation using P which contains both $P(x_1, y_1)$ and $P(x_2, y_2)$.*

(ii) *If $P(x_1, y_1)$, $P(x_2, y_2)$, and $P(x_3, y_3)$ form a hole and have a gridline in common, then there exists no total tessellation using P which contains both $P(x_1, y_1)$, $P(x_2, y_2)$, and $P(x_3, y_3)$.*

Proof. (i) From Lemmas 3.5(i) and 3.6.

(ii) From Lemmas 3.5(ii) and 3.6. ■

4. CONDITIONS FOR PERIODIC TESSELLATIONS

Recall the definition of the boundary $B(x, y)$ of an embedded polyomino $P(x, y)$ and its numbering. Now suppose that some partial or total tessellation τ is given. 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 for all $i > 0$: $B(x_i, y_i) \cap B(x_0, y_0) \neq \emptyset$. 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 gridline 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. It is clear that the boundary of $B(x_0, y_0)$ splits up in a number of consecutive segments $[r_{i_0}(x_0, y_0) \cdots r_{i_1}(x_0, y_0)]$, $[r_{i_2}(x_0, y_0) \cdots r_{i_3}(x_0, y_0)]$, ... ($r_0 \leq r_{i_0} \leq r_{i_1} < r_{i_2} \leq r_{i_3} < \dots$) that are the borderline with instances $P(x_i, y_i)$. We will assume that each segment is maximal for the particular $P(x_i, y_i)$.

LEMMA 4.1. *If τ is a total tessellation then each $P(x, y)$ that surrounds $P(x_0, y_0) \in \tau$ generates exactly one, contiguous segment on the boundary.*

Proof. Directly from Theorem 3.7. ■

DEFINITION 4.2. A partial segmentation of $B(x_0, y_0)$ is any set of (maximal and disjoint) indexed segments I_0, I_1, \dots along $B(x_0, y_0)$ that 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 segmentation of $B(x, y)$ will be denoted as $\text{Seg}(B(x, y))$. The number of segments in it will be denoted by $|\text{Seg}(B(x, y))|$. Its “length” is defined in an obvious manner. Clearly partial surroundings lead to partial segmentations and complete surroundings lead to complete segmentations of $B(x_0, y_0)$. Surroundings and the segmentations they induce are the key to a further understanding of periodic tessellations.

DEFINITION 4.3. 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.

LEMMA 4.4. *If a tessellation τ is periodic then τ is regular.*

Proof. Follows from the observation that in a periodic tessellation the relative positions of the surrounding polyominoes must be the same for every $P(x, y) \in \tau$. ■

The reverse, regularity implies periodicity, is valid too, and will be proved later in this section.

LEMMA 4.5. *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 4.3.

(\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 $P(x_0, y_0)$ wherever (x_0, y_0) is located. Observe that $|\text{Seg}(P(x_0, y_0))| = k$, by virtue of Lemma 4.1. 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 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. ■

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. The length of an edge e will be the number of unit-length gridlines of which it is composed, denoted as $|e|$. Clearly G_τ is a planar graph with nodes of degree 3 or 4.

DEFINITION 4.6. A three-node (four-node) is any gridpoint g where three (four) nonoverlapping instances of P meet. The branches of g are the three (four) edges that meet in g (taken in consecutive order).

We normally refer to the three- and four-nodes of some G_τ with τ total but the definition applies to any local configuration of some $P(x_0, y_0)$ and a (partial or complete) surrounding. 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 gridpoint where two boundaries part.

LEMMA 4.7. 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 $P(x_i, y_i) \cap P(x_{i+1}, y_{i+1}) \neq \emptyset$ for $1 \leq i \leq 6$ (and $x_7 \equiv x_1, y_7 \equiv y_1$). The length of the partial segmentation induced 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 for any $P(x_0, y_0) \approx P_0$ we can find two additional nonoverlapping instances $P(x_1, y_1) \approx P_1$ and $P(x_2, y_2) \approx P_2$ that surround it, with the three of them meeting in g . Furthermore $B(x_0, y_0) \cap B(x_1, y_1) \neq \emptyset$, $B(x_1, y_1) \cap B(x_2, y_2) \neq \emptyset$, and $B(x_2, y_2) \cap B(x_0, y_0) \neq \emptyset$. We apply Lemma 3.4 to P_0, P_1 , and P_2 and obtain the situation as shown in Fig. 10.




FIGURE 10

Observe that P_0, P_1, \dots, P_6 do not overlap each other and that the relative position of P_0 and P_1 is the same as that of P_4 and P_0 and of P_3 and P_2 and of P_5 and P_6 and so on. So $P(x_i, y_i) \cap P(x_{i+1}, y_{i+1}) \neq \emptyset$ for $1 \leq i \leq 6$ (and $x_7 \equiv x_1, y_7 \equiv y_1$).

Finally consider the partial segmentation induced on the boundary of $P(x_0, y_0)$. From Theorem 3.7 now follows that no other instances of P but P_1 through P_6 can border $P(x_0, y_0)$. Each P_i ($1 \leq i \leq 6$) gives rise to exactly one segment along $B(x_0, y_0)$. This identifies the 6 segments along $B(x_0, y_0)$, with a total length equal to $|T_1| + |T_3| + |T_1| + |T_2| + |T_2| + |T_3| = 2 \cdot (|T_1| + |T_2| + |T_3|)$. ■

LEMMA 4.8. *Suppose P admits a four-node g with branches T_1, T_2, T_3 , and T_4 . 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 the length of the partial segmentation induced on $B(x_0, y_0)$ is equal to $2 \cdot |T_1| + 2 \cdot |T_4|$ (resp. $2 \cdot |T_4| + 2 \cdot |T_3|$, $2 \cdot |T_3| + 2 \cdot |T_2|$, and $2 \cdot |T_2| + 2 \cdot |T_1|$).*

Proof. Suppose P admits a four-node g as described. So for every $P(x_0, y_0) \approx P_0$ we can find three additional nonoverlapping instances $P(x_1, y_1) \approx P_1, P(x_2, y_2) \approx P_2$, and $P(x_3, y_3) \approx P_3$ that surround it, with the four of them meeting in a four-node g . Let P_0 border P_1 along T_1 , P_1 border P_2 along T_2 , P_2 border P_3 along T_3 , and P_3 border P_0 along T_4 . Consider P_3, P_0, P_1 , then these three instances meet the conditions as stated in Lemma 3.4. So we obtain again a situation as shown in Fig. 6. By means of the same arguments as in the previous proof this gives the desired result. Note that $B(x_1, y_1) \cap B(x_3, y_3) = \emptyset$. The other partial surroundings are obtained by considering P_2, P_3, P_0 (resp. P_1, P_2, P_3 and P_0, P_1, P_2). ■

LEMMA 4.9. *Let τ be a regular 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. Let τ be a regular 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 4.7. This identifies three relative positions (the “vectors” $\pm \mathbf{a}$, $\pm \mathbf{b}$, and $\pm \mathbf{c}$ in Fig. 10) which, because of the regularity 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 4.7 can be justified as a polyomino actually belonging to τ . As any hole 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 (because of Theorem 3.7) 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 regularity 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 one can argue as before that the polyominoes $P(x_1, y_1)$, $P(x_2, y_2)$, $P(x_3, y_3)$, and $P(x_4, y_4)$ of Lemma 4.8 must all belong to τ and form a complete surrounding of P_0 . Since τ is regular, 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. ■

LEMMA 4.10. *Let τ be a regular tessellation of the plane using P . Then the underlying set of points (x, y) such that $P(x, y) \in \tau$ forms a lattice.*

Proof. Let τ be regular. By Corollary 4.9 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 4.7 (viz. Fig. 10) that two vectors will do in case (a).

(Note in Fig. 10 that, e.g., **c** is integrally dependent on **a** and **b** and that **a** and **b** "generate" the entire tessellation.) By the same token it follows from the proof of Lemma 4.8 that two vectors suffice in case (b) as well. ■

THEOREM 4.11. *A (total) tessellation τ is periodic if and only if τ is regular.*

Proof. Directly from Lemma 4.4 and Definition 3.2 together with Lemma 4.10. ■

Thus, "regularity" exactly characterizes periodic tessellations, and all results we obtained for regular tessellations are valid for periodic tessellations as well.

COROLLARY 4.12. *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: $|T_1| + |T_4| \leq \frac{1}{2}|B|$ (resp. $|T_4| + |T_3| \leq \frac{1}{2}|B|$, $|T_3| + |T_2| \leq \frac{1}{2}|B|$, and $|T_2| + |T_1| \leq \frac{1}{2}|B|$). Here 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 .*

Proof. Consider any three-node g of G_τ . By Lemma 4.7 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 $2 \cdot (|T_1| + |T_2| + |T_3|)$. Hence $|T_1| + |T_2| + |T_3| \leq \frac{1}{2}|B|$. Likewise it follows from Lemma 4.8 that for every four-node g of G_τ : $|T_1| + |T_4| \leq \frac{1}{2}|B|$, respectively, $|T_4| + |T_3| \leq \frac{1}{2}|B|$, $|T_3| + |T_2| \leq \frac{1}{2}|B|$, and $|T_2| + |T_1| \leq \frac{1}{2}|B|$. ■

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 4.13. *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 4.9 and Theorem 4.11 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 Lemma 4.9 shows that the surrounding of any $P(x_0, y_0)$ as constructed in the proof of Lemma 4.7 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 Lemma 4.9 shows likewise that the surroundings constructed in the proof of Lemma 4.8 must be complete. Thus $|T_1| + |T_4| = \frac{1}{2}|B|$ and $|T_2| + |T_3| = \frac{1}{2}|B|$. It follows that $|T_1| + |T_2| + |T_3| + |T_4| = |B|$ in this case.

(\Leftarrow) 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 4.7 necessarily is a complete surrounding. Observing the relative positions of $P(x_0, y_0)$ and its surrounding polyominoes, it follows that for each of the $P(x_i, y_i)$: $\text{Seg}(B(x_i, y_i)) \subseteq \text{Seg}(B(x_0, y_0))$. So the conditions of Lemma 4.5 are satisfied and the surrounding can be extended to a regular (hence: periodic) tessellation of the entire plane.

If P satisfies (**) rather than (*), then from Corollary 4.12 follows that $|T_1| + |T_4| = \frac{1}{2}|B|$ and a similar argument carries through, based on the construction of a surrounding in the proof of Lemma 4.8, and shows with Lemma 4.5 that again a periodic tessellation can be obtained using P . ■

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. Sect. 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 endpoint.

DEFINITION 5.1. The support of a node $g \in G_\tau$, denoted as: $\text{Sup}_\tau(g)$ or just as: $\text{Sup}(g)$, 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 endpoint.) The proof of Shapiro's conjecture heavily relies on the criteria for periodic tessellations in Theorem 4.13 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_τ . Let G'_τ 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'_τ is a connected and finite section of G_τ , 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'_τ . Define factors ε (depending on τ , K , and N) such that

$$\begin{aligned} \varepsilon_1 \cdot K &= \text{the number of three-nodes along } C \text{ that have degree 2 in } G'_\tau, \\ \varepsilon'_1 \cdot K &= \text{the number of three-nodes along } C \text{ that (still) have degree 3 in } G'_\tau, \\ \varepsilon_2 \cdot K &= \text{the number of four-nodes along } C \text{ that have degree 2 in } G'_\tau, \\ \varepsilon'_2 \cdot K &= \text{the number of four-nodes along } C \text{ that have degree 3 in } G'_\tau, \\ \varepsilon''_2 \cdot K &= \text{the number of four-nodes along } C \text{ that (still) have degree 4 in } G'_\tau. \end{aligned}$$

Claim 5.2.1. For N sufficiently large each factor ε is less than $1/(2 \cdot |B|)$, 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'_τ 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'_τ) 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'_τ)

$$= \sum_i \frac{i}{3} \cdot \alpha_{i*} + \frac{2}{3} \cdot \varepsilon_1 K + \frac{1}{3} \cdot \varepsilon'_1 K,$$

f = the number of four-nodes within the window (hence in G'_τ)

$$= \sum_j \frac{j}{4} \cdot \alpha_{*j} + \frac{3}{4} \cdot \varepsilon_2 K + \frac{2}{4} \cdot \varepsilon'_2 K + \frac{1}{4} \cdot \varepsilon''_2 K,$$

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

$$= t + f,$$

e = the total number of edges (branches) within the window (hence in G'_τ)

$$= \sum_{i,j} \frac{i+j}{2} \cdot \alpha_{ij} + \frac{1}{2} \cdot \varepsilon_1 K + \frac{1}{2} \cdot \varepsilon'_1 K + \frac{1}{2} \cdot \varepsilon_2 K \\ + \frac{1}{2} \cdot \varepsilon'_2 K + \frac{1}{2} \cdot \varepsilon''_2 K,$$

p = the total number of parts (faces) into which the plane is divided by G'_τ

$$= K + 1.$$

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

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

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

$$\begin{aligned} &\Rightarrow \sum_i \frac{i}{3} \cdot \alpha_{i*} + \frac{2}{3} \varepsilon_1 K + \frac{1}{3} \varepsilon'_1 K + \sum_j \frac{j}{4} \cdot \alpha_{*j} + \frac{3}{4} \varepsilon_2 K + \frac{1}{2} \varepsilon'_2 K + \frac{1}{4} \varepsilon''_2 K \\ &\quad + K + 1 = \sum_{i,j} \frac{i+j}{2} \cdot \alpha_{ij} + \frac{1}{2} (\varepsilon_1 + \varepsilon'_1 + \varepsilon_2 + \varepsilon'_2 + \varepsilon''_2) K + 2 \\ &\Rightarrow \sum_i \frac{i}{6} \cdot \alpha_{i*} + \sum_j \frac{j}{4} \cdot \alpha_{*j} = \frac{1}{6} \varepsilon_1 K - \frac{1}{6} \varepsilon'_1 K + \frac{1}{4} \varepsilon_2 K - \frac{1}{4} \varepsilon''_2 K + K - 1. \end{aligned}$$

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

$$\begin{aligned} t + 2f &= \left(\sum_i \frac{i}{3} \cdot \alpha_{i*} + \sum_j \frac{j}{2} \cdot \alpha_{*j} \right) + \frac{2}{3} \varepsilon_1 K + \frac{1}{3} \varepsilon'_1 K + \frac{3}{2} \varepsilon_2 K + \varepsilon'_2 K + \frac{1}{2} \varepsilon''_2 K \\ &= \varepsilon_1 K + 2\varepsilon_2 K + \varepsilon'_2 K + 2K - 2, \end{aligned}$$

and the expression claimed for f easily follows.

Suppose by way of contradiction that τ (hence G'_τ) does not contain any three-nodes satisfying (*) nor any four-nodes satisfying (**). By Corollary 4.12 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_1 , etc., are the branches of the node. It means that for every three-node g : $\text{Sup}(g) \leq \frac{1}{4}|B| - \frac{1}{2}$. Let L be the total edge length of G'_τ . Note that $L < \sum_{g \in G'_\tau} \text{Sup}(g)$. (The $<$ sign holds because there is at least one node along the contour of G'_τ 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:

$$\begin{aligned} L &< \sum_{\substack{g \in G'_\tau \\ (g_{\text{three-node}})}} \text{Sup}(g) + \sum_{\substack{g \in G'_\tau \\ (g_{\text{four-node}})}} \text{Sup}(g) \\ &\leq t \cdot \left(\frac{1}{4}|B| - \frac{1}{2} \right) + f \cdot \left(\frac{1}{2}|B| - \frac{1}{2} \right) \\ &< t \cdot \left(\frac{1}{4}|B| - \frac{1}{2} \right) + \left(-\frac{1}{2}t + \frac{1}{2}\varepsilon_1 K + \varepsilon_2 K + \frac{1}{2}\varepsilon'_2 K + K \right) \left(\frac{1}{2}|B| - \frac{1}{2} \right) \\ &= -\frac{1}{4}t + \left(\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon'_2 \right) K \cdot |B| + \frac{1}{2}K \cdot |B| \\ &\quad - \left(\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon'_2 \right) K - \frac{1}{2}K \\ &\leq \frac{1}{2}K \cdot |B| + \left\{ \left(\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon'_2 \right) \cdot |B| - \frac{1}{2} \right\} K. \end{aligned}$$

As N was chosen sufficiently large, it easily follows from Claim 5.2.1 that $\frac{1}{4}\varepsilon_1 + \frac{1}{2}\varepsilon_2 + \frac{1}{4}\varepsilon'_2 < 1/(2 \cdot |B|)$. Thus our estimate on L reduces to

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

On the other hand, if we let each of the K polyominoes in G'_τ 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'_τ (hence τ) must contain a three-node satisfying (*) or a four-node satisfying (**). ■

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_τ . The argument is easily completed from here. ■

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 4.13. (Note the additional observations for periodic tessellations in Sect. 4.) ■

6. FINAL COMMENTS

Our study of plane tessellations was motivated from the theory of data organisation for SIMD machines. We argued in Section 1 (see also Shapiro, 1978) that only periodic tessellations are of practical interest. Thus the proof of Shapiro's conjecture has significance within this context. 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 compute (by inspecting any $N \times N$ window) a three-node satisfying (*) or a four-node satisfying (**). Second, the results underlying Theorem 4.13 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 5.3 we only need to test the conditions for a periodic tessellation using P as expressed in Theorem 4.13. 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 takes at most $O(|B|^2)$ (hence: polynomial) time. ■

The study of plane tessellations (tilings, pavings) with regular objects has a long history in mathematics. It has repeatedly been the subject of M. Gardner's column in the *Scientific American* (1975 a, b, c; 1977). A systematic study of tessellations with sets of polyominoes was made by Golomb (1966). In the late sixties Golomb (1970) proved that the question whether an arbitrary finite set of polyominoes tiles the plane (rotational symmetries, etc., allowed) is equivalent to Wang's domino problem (1965) and hence algorithmically undecidable. If the set contains only one polyomino, the decidability question is reportedly still open (Göbel, 1979). Thus the results we proved in this paper, and Theorem 6.1 in particular, may be viewed as a partial answer to this question for a restrictive class of tessellations (requiring polyominoes to have a fixed orientation).

Severe problems arise if we attempt to generalize Theorem 5.4 and, e.g., relax the condition that P is a polyomino. The template T shown in Fig. 11 provides an example that Shapiro's conjecture 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




FIGURE 11




FIGURE 12

periodicity in such a manner that a suitable modification of Shapiro's conjecture remains valid. Other problems arise if we no longer insist that tessellations have to be total. For instance, it is not true that the existence of a partial tessellation with a certain density using a polyomino P implies the existence of a periodic partial tessellation with the same density using P . To illustrate this see Fig. 12, in which a partial tessellation is shown, which covers $\frac{9}{11}$ of the plane. If we want to tessellate the plane in a periodic way with this polyomino then at most $\frac{8}{11}$ of the plane can be covered. (The example is due to H. L. Bodlaender.)

The existence of periodic tessellations in general, using sets of objects and allowing symmetries, is a hard problem for which only a few results have been proved. It is known (see Gardner, 1977) that there exists a set of 2 polygons which tile the plane nonperiodically only. Thus there are many inspiring problems left in the study of tessellations.

REFERENCES

- BUDNIK, P., AND KUCK, D. J. (1971), The organisation and use of parallel memories, *IEEE Trans. Comput.* **C-20**, 1566–1569.
- GARDNER, M. (1975 a, b, c; 1977), Mathematical games, articles in *Sci. Amer.* July, 112–117; August, 112–115; December, 116–119; January, 110–121.
- GÖBEL, F. (1979), Geometrical packing and covering problems, in “Packing and Covering in Combinatorics” (A. Schrijver, Ed.), Math. Centre Tracts 106, pp. 179–199, Mathematisch Centrum, Amsterdam.
- GOLOMB, S. W. (1966), Tiling with polyominoes, *J. Combin. Theory* **1**, 280–296.
- GOLOMB, S. W. (1970), Tiling with sets of polyominoes, *J. Combin. Theory* **9**, 60–71.
- LEVI, P. (1934), Sur une généralisation du théorème de Rolle, *C. R. Acad. Sci. Paris* **198**, 424–425.
- SHAPIRO, H. D. (1978), Theoretical limitations on the efficient use of parallel memories, *IEEE Trans. Comput.* **C-27**, 421–428.
- THURBER, K. J. (1976), “Large Scale Computer Architecture: Parallel and Associative Processors,” Hayden, Rochelle Park, N.J.
- WANG, H. (1965), Games, logic and computers, *Sci. Amer.* **213**, 98–106.
- WUISHOFF, H. A. G., AND VAN LEEUWEN, J. (1983), “Periodic Storage Schemes with a Minimum Number of Memory Banks,” Technical report RUU-CS-83-4, Dept. of Computer Science, University of Utrecht.