The One Dimensional Skewing Problem

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

RUU-CS-89-23
October 1989

Rijksuniversiteit Utrecht

Vakgroep informatica

Padualaan 14 3584 CH Utrecht
Corr. adres: Postbus 80.089, 3508 TB Utrecht
Telefoon 030-531454
The Netherlands

The One Dimensional Skewing Problem

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

RUU-CS-89-23
October 1989

Rijksuniversiteit Utrecht

Vakgroep informatica

Padualaan 14 3584 CH Utrecht
Corr. adres: Postbus 80.089, 3508 TB Utrecht
Telefoon 030-531454
The Netherlands

The One Dimensional Skewing Problem

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

Technical Report RUU-CS-89-23
October 1989
(revised version of Technical Report RUU-CS-86-8)

Department of Computer Science
University of Utrecht
P.O.Box 80.089, 3508 TB Utrecht
The Netherlands

The One Dimensional Skewing Problem

G. Tel* and J. van Leeuwen
Department of Computer Science, University of Utrecht,
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands

H.A.G. Wijshoff
Center for Supercomputing Research and Development,
305 Talbot Laboratory, 104 South Wright Street, Urbana, IL 61801-2932

July 1986 / Revised: October 1989

Abstract

Parallel computers (such as vector machines and array-processors) feature the availability of many, highly pipelined processing units and many memory banks which can be accessed independently in parallel at great speed. Aside from needing adequately parallellized ("vectorized") algorithms, their application requires general storage mappings for distributing and retrieving vector data from memory at low cost. Data mappings of this kind, also known as "skewing schemes", were first considered during the design of the ILLIAC IV in the late nineteen sixties. The design of good skewing schemes is non-trivial. In this paper we analyse the simplest case of skewing an (infinite) one-dimensional array, in such a way that cells that are connected according to a given data template T are mapped in different banks. We derive upper and lower bounds for the minimum number of memory banks in which this can be done in the general case, and an extra upper bound for the special case of (k, l, m) templates. Several additional results are proved for natural classes of templates, which are indicative for the combinatorial complexity of the problem.

1 Introduction

Since the late sixties, parallel computers are designed and built to reach a higher speed of computation than would be possible with conventional architectures. The parallel computers (or "super computers") consist of a (large) number of *arithmetic units* (also called *processing elements* or PE's), which perform operations in parallel. Of course it is necessary that memory is able to store and fetch data at high rate. That is why memory is split up into a number of independent blocks called *memory banks*. The banks and the arithmetical units are connected by an *interconnection network* (see figure 1). In one cycle of the supercomputer each of the arithmetical units performs one fetch, store or arithmetical operation. The set of data elements that is accessed in one cycle is called a *vector of interest.* Because a memory bank can handle only one store or fetch operation in each cycle, it is desired to store the elements of a vector of interest in different banks.

*The work of this author was partially supported by the ESPRIT II Basic Research Actions Program of the EC under contract no. 3075 (project ALCOM).

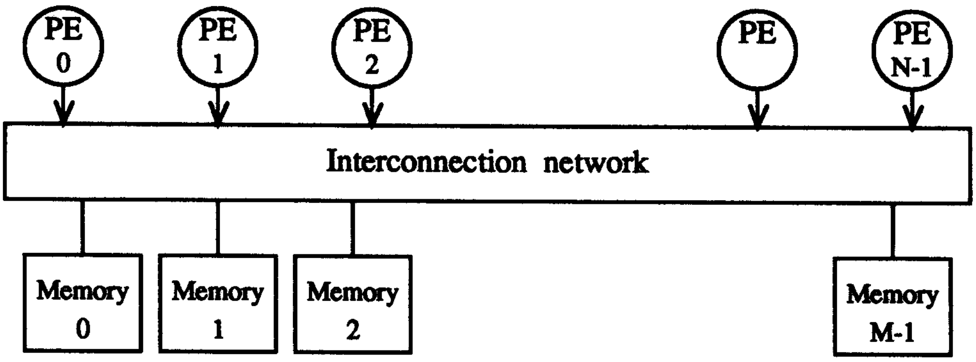

Figure 1: Architecture of a supercomputer.

Data mappings that achieve this are called *skewing schemes*. If two units request a data item from the same memory bank, a *memory conflict* occurs: if it happens it will take two or more cycles to do the work of one.

In this paper we consider the simplest form of data distribution that arises in parallel computers. The set of data elements is an infinite, one-dimensional array, indexed by $\mathbb{Z}$. The vectors of interest will be the images under translations of one (fixed) finite subset of $\mathbb{Z}$ called a *template*. We are interested in finding a skewing scheme that uses the smallest possible number of memory banks to achieve conflict-freeness.

The research on data distributions and skewing schemes was initiated by Budnik and Kuck [1] and Shapiro [3]. For an overview of the theory of skewing schemes we refer to Wijshoff [5]. A more recent result, concerning hierarchical memories and special classes of skewing schemes, is given by Tel and Wijshoff [4].

This paper is organized as follows. Section 2 contains definitions, some general basic results on one dimensional templates and quotes an important result by Shapiro (theorem 2.7). Section 3 deals with a class of templates called (k, l, m) templates. We investigate when this template is linefilling and prove an upper bound on the number of memory banks necessary to skew it validly. Section 4 deals with a class of templates called (k, l, m, n, o) templates. We investigate when this template is linefilling. In section 5 we consider general templates and sketch a "tessellation test" for templates.

2 Preliminaries

In the problem we consider the data is organized in an infinite array, which we want to allocate in some way to M memory banks numbered $0 \ldots M-1$.

Definition 2.1 *A* skewing scheme *is a mapping* $s : \mathbb{Z} \longrightarrow \{0, \ldots, M-1\}$.

This definition is to be interpreted as follows. If $s(i) = p$, then the array element indexed by i is stored in bank number p.

Definition 2.2 *A* template *is a finite subset* $T = \{x_0, \ldots, x_{c-1}\}$ *of* $\mathbb{Z}$. *An* instance $T(y)$ *of a template* T *is the set* $\{x + y | x \in T\}$.

As a convention we shall always assume that $x_0 = 0$ and all other x_i are positive. This is because we are interested in instances, and as can be seen from the definition the instances of $\{x_0, x_1, \ldots, x_{c-1}\}$ are the same as the instances of $\{0, x_1 - x_0, \ldots, x_{c-1} - x_0\}$. The instances of a template T (or, in some cases, all instances of all of some *set* of templates $T_1, \ldots, T_k$) are the vectors of interest. The scheme s "works" for the template T if all elements of a vector of interest are mapped to different banks.

Definition 2.3 *A skewing scheme s is* valid *for a template T iff for all $y \in \mathbf{Z}$, $s \restriction T(y)$ is an injection $T(y) \hookrightarrow \{0, \ldots, M-1\}$. A skewing scheme s is* valid *for a collection of templates $\mathcal{C} = \{T_1, \ldots, T_k\}$ iff s is valid for all $T_i \in \mathcal{C}$.*

($s \restriction T(y)$ means the function s with the domain restricted to $T(y)$. The condition on validness is equivalent to saying that, for all $y \in \mathbf{Z}$ and for all $x_1, x_2 \in T$, $s(y + x_1) \neq s(y + x_2)$.) In terms of memory access, the validity of a skewing scheme means that any instance of a template can be accessed in one cycle without memory conflicts. Given a template T, it is not hard to find a skewing scheme that is valid for T when M is large enough. But, we are interested in using an M as small as possible. First we introduce some handy notations:

$$\begin{aligned} c &: \text{ the number of cells in } T \\ Dif(T) &: \text{ the set } \{|x_i - x_j| : x_i, x_j \in T\} \\ d &: \text{ the size of } Dif(T) \\ len(T) &: \max_{x \in T} x + 1 \end{aligned}$$

If $y \in Dif(T)$ then $y = |x_i - x_j|$ for some $0 \leq x_i, x_j < len(T)$, hence $y < len(T)$. Thus $Dif(T) \subseteq \{0, 1, \ldots, len(T) - 1\}$. We say T is *complete* if $Dif(T) = \{0, 1, \ldots, len(T) - 1\}$. This is equivalent to requiring that $d = len(T)$.

Let $M(T)$ be the minimum number of banks for which there is a valid skewing scheme for T. Obviously $M(T) \geq c$ because $s \restriction T$ must be an injection. A somewhat different characterization of validity is given by the following result.

Proposition 2.4 *s is valid for T iff for all $x_1, x_2 \in \mathbf{Z}$, if $x_1 - x_2 \in Dif(T)$ then $s(x_1) \neq s(x_2)$.*

Proof. Suppose s is valid for T and $x_1 - x_2 \in Dif(T)$. Then there is an instance $T(y)$ of T such that $T(y)$ contains both x_1 and x_2 and because $s \restriction T(y)$ is an injection, $s(x_1) \neq s(x_2)$ follows.

Suppose for all $x_1, x_2 \in \mathbf{Z}$ with $x_1 - x_2 \in Dif(T)$, $s(x_1) \neq s(x_2)$. For an instance $T(y)$ and $x_1, x_2 \in T(y)$ we have $|x_1 - x_2| \in Dif(T)$ and so $s(x_1) \neq s(x_2)$. Hence $s \restriction T(y)$ is an injection. □

An obvious upper bound on $M(T)$ is $len(T)$: set $M = len(T)$ and define $s(x) = x \bmod M$. This scheme is valid because if $x_1 - x_2 \in Dif(T)$ then $x_1 - x_2 < len(T) = M$ so $s(x_1) \neq s(x_2)$. So now we have $c \leq M(T) \leq len(T)$. The following result is less immediate.

Lemma 2.5 *$M(T) \leq d$.*

Proof. Set $M = d$, and start with an arbitrary skewing scheme s, not necessarily valid for T, such as $s(i) = 0$ for all i. We construct a skewing scheme s, using at most M banks, which is valid for T. First construct $s(0)$, then $s(1)$, and so forth ad infinitum. To construct $s(i)$, consider the values of s in points $i - f$, for $f > 0$ and $f \in Dif(T)$. There are at most $d-1$ different values to consider here so (by the pidgin-hole principle) there is at least one $p \in \{0, \ldots, M-1\}$ which is not among them. Set $s(i) = p$. After defining $s(i)$ for all positive values, construct $s(-1), s(-2)$ and so on in the same way, now considering values $s(i + f)$. By the construction of s, its range is a subset of $\{0, \ldots, M-1\}$ and the condition of proposition 2.4 is fulfilled. Hence $M(T) \leq d$. □

We now want to know when the two proven bounds are equalities. For the bound $M(T) \leq len(T)$ we have the following result.

Lemma 2.6 *$M(T) = len(T)$ iff T is complete.*

Proof. Suppose $M(T) = len(T)$. Then $d = |Dif(T)| \geq len(T)$ (see 2.5). As $Dif(T) \subseteq \{0, \ldots, len(T)-1\}$, it follows that $Dif(T) = \{0, \ldots, len(T)-1\}$, hence T is complete.

Suppose T is complete, i.e., all y, $0 \leq y \leq len(T) - 1$ are in $Dif(T)$. This implies that, if s is valid for T, then $s(0), s(1), \ldots, s(len(T)-1)$ are $len(T)$ different numbers (use proposition 2.4). So the number of banks used by s is at least $len(T)$, use $M(T) \leq len(T)$ (see above) to prove that $M(T) = len(T)$. □

For the bound $M(T) \geq c$ we have the following classical result due to Shapiro [3], which takes the following form for the one-dimensional case.

Theorem 2.7 (Shapiro) *$M(T) = c$ iff T tessellates the line.*

We say that T *tessellates* the line (or T is *line filling*) iff there is an (infinite) set $E \subseteq \mathbb{Z}$ such that $\{T(e) : e \in E\}$ is a partition of $\mathbb{Z}$. Informally, tiles of shape T can be used to tile $\mathbb{Z}$. Because of Shapiro's theorem, tiling is an important notion in the literature on skewing (see e.g. Wijshoff [5]). We assume here that whenever we say that a set A tessellates a set B, no rotations, symmetries, etc., are allowed: All tiles in the tessellation must have the same, fixed orientation. An obvious but important lemma about tessellations is:

Theorem 2.8 *If a set A tiles B and B tiles C, then A tiles C.*

Proof. Use tiles of shape A to construct tiles of shape B, and next tile C with these tiles. See figure 2. Formally, if $\{A(e) : e \in E\}$ is a partition of B, and $\{B(f) : f \in F\}$ is a partition of C, then let $G = \{e + f : e \in E, f \in F\}$ and see that $\{A(g) : g \in G\}$ is a partition of C. □

Testing on completeness of a template is easily done in time polynomial in c. Testing whether a template tessellates the line is not so easy: we address this problem in section 5.

Obviously, $d \leq \frac{1}{2}c(c-1) + 1$: there are $\binom{c}{2} = \frac{1}{2}c(c-1)$ possible pairs of different cells in T, and $0 \in Dif(T)$. So by lemma 2.5 $M(T) \leq \frac{1}{2}c(c-1) + 1$. For $c = 3$ and $c = 4$ there are templates for which this bound is sharp. $T = \{0, 1, 3\}$ is complete and hence $M(T) = 4$. $T = \{0, 1, 4, 6\}$ is complete and hence $M(T) = 7$. We do not know whether there exist templates consisting of 5, 6, or more cells that require $\frac{1}{2}c(c-1) + 1$ banks, but we believe that there are none for large c.

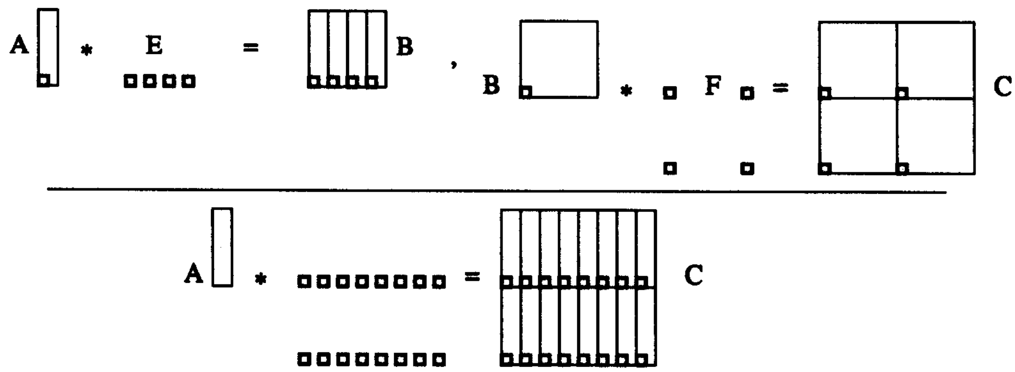

Figure 2: Transitivity of tessellation.

We call T_1 a *sub template* of T_2 if there is a y such that $T_1(y) \subseteq T_2$, notation: $T_1 \subseteq_T T_2$. Informally, one can say that T_1 can be embedded in T_2. Being a subtemplate does not imply set inclusion, note that for $T_1 = \{0, 2\}$ and $T_2 = \{0, 1, 3\}$ we have $T_1 \subseteq_T T_2$ but not $T_1 \subseteq T_2$. Note however that $T_1 \subseteq_T T_2$ implies $Dif(T_1) \subseteq Dif(T_2)$. Finally $Dif(T_1) \subseteq Dif(T_2)$ implies $M(T_1) \leq M(T_2)$, because every skewing scheme that is valid for T_2 is also valid for T_1 (use proposition 2.4).

3 (k, l, m)-Templates

We will now look at a special class of templates, that consist of two rows of consecutive cells (of size k and size m, respectively), and some space (of size l) in between.

Definition 3.1 *The template* (k, l, m) *is the template* $\{0, 1, \ldots, k-1, k+l, k+l+1, \ldots, k+l+m-1\}$.

By symmetry we have $M(k, l, m) = M(m, l, k)$. Thus without loss of generality we can assume $k \geq m$. We see that

$$\begin{aligned} c &= k + m \\ len(k, l, m) &= k + l + m \\ Dif(k, l, m) &= \{0, 1, \ldots, k-1, l+1, l+2, \ldots, k+l+m-1\} \end{aligned}$$

The following results state when (k, l, m) is complete and when it tessellates $\mathbf{Z}$.

Proposition 3.2 (k, l, m) *is complete iff* $l < k$.

Proof. See the representation of Dif above. □

Theorem 3.3 (k, l, m) *tessellates the line iff* $k + m | l$ *or* $(k = m \wedge k | l)$.

Proof. Suppose first that (k, l, m) tiles $\mathbf{Z}$. If $k = m$, the gap of length l between the two blocks of one instance is filled up with blocks of length k, hence $k | l$. Now let $k > m$.

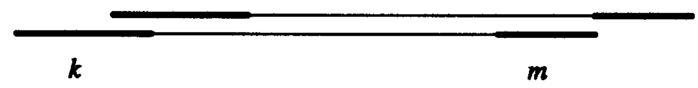

Figure 3: No m parts abut.

We claim that no two instances in the tessellation are placed such that their two m-parts are abut (see figure 3), because their k-parts would otherwise overlap. Next we claim that no two instances in the tessellation are placed such that their two k-parts are abut (see figure 4), because between their m-parts would be a gap of size $k - m$ and this gap cannot be filled up without violating the first claim. So in the tessellation k and m-blocks

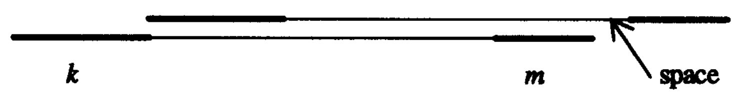

Figure 4: No k parts abut.

alternate. Now fix one instance (see figure 5) and observe that the gap of size l is filled

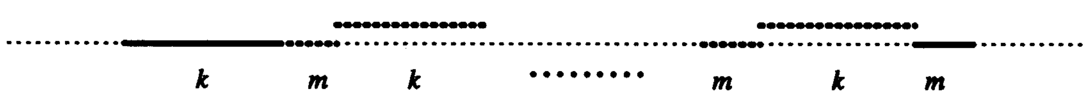

Figure 5: Tessellation with (k, l, m).

with an alternation of k and m-blocks, thus showing $k + m | l$.

To prove the theorem in the other direction, suppose $k + m | l$. With $E = \{\ldots, -2c, -c, 0, c, 2c, \ldots\}$, $\{T(e) : e \in E\}$ is a partition of $\mathbb{Z}$ (see figure 6). Finally suppose $k = m$ and $l = R \times k$. Now $R + 1$ copies of the template tile $\{0, 1, \ldots, (R+1)c - 1\}$ (see figure 7) when based at $0, k, \ldots, R \times k$. A sequence of thes "supertiles" tessellates $\mathbb{Z}$, so by theorem 2.8 our original tile does. □

In the remainder of this section we try to establish a bound on $M(k, l, m)$ for the non-linefilling cases. To achieve a bound on $M(k, l, m)$ we first look at the template $U = (k, 0, m)$, which is in fact a consecutive sequence of $k + m = c$ cells.

We say a skewing scheme is *periodic with period* p iff for all i, $s(i + p) = s(i)$. The sequence $s(0), s(1), \ldots, s(p-1)$ is called the *period sequence*. Note that if this is the case, and s is valid for T, then s is also valid for any template U with $U \equiv_p T$. Here by $U \equiv_p T$ we mean that for each $u \in U$ there is *one* $x \in T$ for which $u \equiv x \pmod{p}$ and vice versa. Obviously, $(k, l, m) \equiv_l (k, 0, m)$ so we will try to find a skewing scheme which is valid for U and periodic with period l. For an integer r, let the skewing scheme s_r be defined by $s_r(x) = x \bmod r$. s_r uses r banks and is periodic with period r. It is valid for U if $c \leq r$.

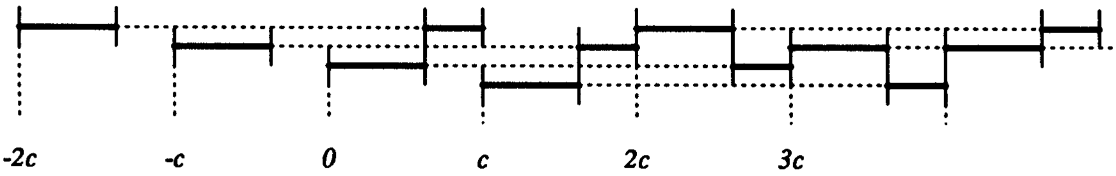

Figure 6: Tessellation where $k + m|l$.

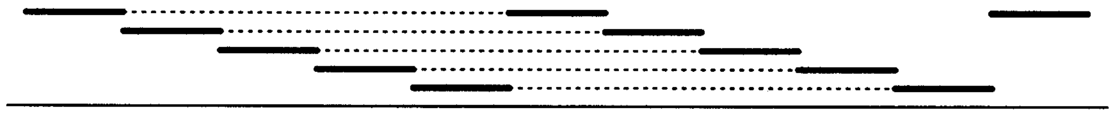

Figure 7: Tessellation where $k = m \wedge k|l$.

We shall now describe a way to "glue" periods of different skewing schemes together. The period sequence of the s_r (for $r \geq c$) all start out with the sequence $0, 1, \ldots, c-1$. Every scheme that consists of a succession of period sequences of an s_r (with $r \geq c$) is also valid for U: the same bank number does not appear twice within distance c. As we want a period length of l for the scheme we are going to construct, we write l as the sum of such r:

$$l = r_1 + \cdots + r_n. \tag{1}$$

The scheme we construct has a period sequence of length l, which is the concatenation of n smaller period sequences of length r_1 through r_n.

Because the number of banks used in the new scheme is equal to the largest of the r_i used, we want to have this maximum as small as possible. As for each i $r_i \geq c$, $n \leq l\mathrm{div}c$ and this implies that the maximum of the summands in (1) is at least $\lceil l/(l\mathrm{div}c)\rceil$. Hence this is the best we can achieve with this method, and indeed it is possible to write

$$l = a \times \lceil l/(l\mathrm{div}c)\rceil + b \times \lfloor l/(l\mathrm{div}c)\rfloor \tag{2}$$

for suitably chosen a and b. (The reader is invited to work out formulae for a and b himself.)

Of course l can often be written as (1) in different ways and also the order of the summands influences the final scheme. One of the possible resulting skewing schemes is represented by an elegant formula:

$$s(x) = (x \times (l\mathrm{div}c) \bmod l)\,\mathrm{div}(l\mathrm{div}c) \tag{3}$$

We shall now prove that this scheme is indeed valid for (k, l, m) by showing that it is periodic with period l and that it is valid for $(k, 0, m)$.

Lemma 3.4 *The scheme s defined above is valid for (k, l, m).*

Proof. First we show that s is periodic with period l. As $x \times (l\text{div}c)$ is reduced modulo l, we have indeed that $s(x) = s(x + l)$. Next we show that s is valid for $U = (k, 0, m)$. As $Dif(U) = \{0, \ldots, c-1\}$ it suffices to show that $x \neq y \wedge s(x) = s(y)$ implies $|x - y| \geq c$. Suppose $s(x) = s(y)$. Then

$$|(x \times (l\text{div}c) \bmod l) - (y \times (l\text{div}c) \bmod l)| < l\text{div}c.$$

Now if

$$|(x \times (l\text{div}c)) - (y \times (l\text{div}c))| < l\text{div}c$$

then $|x - y| < 1$ and $x = y$. Otherwise

$$|x \times (l\text{div}c) - y \times (l\text{div}c) - A \times l| < l\text{div}c$$

for a certain $A \neq 0$. But then

$$|x \times (l\text{div}c) - y \times (l\text{div}c)| > l - l\text{div}c$$

which implies $|x - y| > c - 1$. □

Note that the highest possible bank number for s is $(l-1)\text{div}(l\text{div}c)$ so the number of banks is $(l-1)\text{div}(l\text{div}c) + 1 = \lceil \frac{l}{l\text{div}c} \rceil$. The following has now been proved:

Proposition 3.5 $M(k, l, m) \leq \lceil \frac{l}{l\text{div}c} \rceil$.

Note further that $\lceil \frac{l}{l\text{div}c} \rceil = c + \lceil \frac{l\text{mod}c}{l\text{div}c} \rceil$. If $l \bmod c = 0$ this formula yields c and we know this bound is equal to $M(k, l, m)$ in this case. For all $l \geq c(c-1)$ we have $l \bmod c \leq l\text{div}c$ so the formula yields either c or $c+1$ and also in this case it is equal to $M(k, l, m)$. For $k = m$ however we can do even better than this.

Proposition 3.6 $M(k, l, k) \leq c + \lceil \frac{l\text{mod}k}{l\text{div}c} \rceil$.

Proof. If $l \bmod k = l \bmod c$ then this number of banks is the same as in the previous proposition and we use the scheme derived there. As $c = 2k$, $l \bmod k < l \bmod c$ only if $l\text{div}k$ is an odd number:

$$l = (2q + 1)k + r,$$

$r < k$. We define a periodic skewing scheme with period $2(k + l)$. The period sequence is a succession of blocks of numbers as in figure 8 (for $q = 3$). There are $2q + 2$ A-blocks, which

| A | A | C | A | A | C | A | A | C | A | A | B | C | B | B | C | B | B | C | B | B | C | B |

Figure 8: Period sequence of a scheme for (k, l, k).

are of length k and consist of the string $0, 1, \ldots, k-1$. There are also $2q + 2$ B-blocks, which are also of length k but consist of the string $k, k+1, \ldots, 2k-1$. The lengths of the $2q$ C-blocks are $\lceil r/q \rceil$ or $\lfloor r/c \rfloor$, chosen such that the length of the entire period is $2(k + l)$ and the length of the i^{th} C-block is equal to the length of the $(i + q)^{\text{th}}$ one. For $i \leq q$, the

i^{th} C-block is located after the $2i^{\text{th}}$ A-block and the $(q+i)^{\text{th}}$ C-block is located after the $(2i-1)^{\text{th}}$ B-block.

We leave it to the reader to check the validity of this scheme for (k,l,k). It uses $2k + \lceil r/q \rceil$ memory banks, hence $M(k,l,k) \leq \lceil \frac{l \mathrm{mod} k}{l \mathrm{div} c} \rceil$. □

So with the definition

$$F(k,l,m) = \begin{cases} \lceil (l \bmod k)/(l \mathrm{div} c) \rceil + c & \text{if } k = m \\ \lceil (l \bmod c)/(l \mathrm{div} c) \rceil + c & \text{otherwise} \end{cases}$$

we can summarize the two previous propositions in:

Lemma 3.7 $M(k,l,m) \leq F(k,l,m)$

Because $T \subseteq_T U$ implies $M(T) \leq M(U)$ we also have $M(k,l,m) \leq F(k',l',m')$ if $(k,l,m) \subseteq_T (k',l',m')$. In some cases this gives an improved bound. For example, take $T = (3,7,2)$. We have $F(3,7,2) = 7$, but $(3,7,2) \subseteq_T (3,6,3)$ and $M(3,6,3) = 6$ so that $M(3,7,2) \leq 6$ follows. (We know this bound is sharp as T is not linefilling (theorem 3.3) so it cannot be skewed in 5 banks (theorem 2.7).) Also $F(1,8,5) = 8$ but $(1,8,5) \subseteq_T (2,7,5)$ and $F(2,7,5) = 7$. We see that the bound given by the following theorem is an improvement on the prepious lemma:

Theorem 3.8 $M(k,l,m) \leq \min\{F(k',l',m') : (k,l,m) \subseteq_T (k',l',m')\}$.

Theorem 3.8 is constructive: if a search procedure is used to find the minimum and the k',l',m' for which it is assumed, a corresponding skewing scheme can be found using the constructions in this section. The search will in general not be very expensive: As $F(k',l',m') \geq k' + m''$, the search can be pruned where $k' + m'$ exceeds an earlier found value of F. The number of tested cases will in general be small.

4 (k,l,m,n,o)-Templates

Now we turn our attention to the class of templates that do not consist of two, but of three blocks of consequtive cells.

Definition 4.1 *The template* (k,l,m,n,o) *is the template* $\{0,1,\ldots,k-1,k+l,k+l+1,\ldots,k+l+m-1,k+l+m+n,k+l+m+n+1,\ldots,k+l+m+n+o-1\}$

The analysis of this template is much more complicated than that of (k,l,m). We will derive a condition to test whether (k,l,m,n,o) tessellates $\mathbb{Z}$. To do so, we use a different representation of the template, see figure 9. The three blocks are called P, Q, and R, and have length p, q, and r, respectively, where the blocks are named such that $p \geq q \geq r$. a and b represent the relative distance of Q, respectively R, from P. Note that a and b may be negative (as in figure 9). The two representations can easily be converted into each other. Again let c denote $p+q+r$. The main goal of this section is to determine whether (k,l,m,n,o) tessellates $\mathbb{Z}$ (cf. theorem 3.3). For the analysis we distinguish four cases.

1 $p > q > r$. In case 1 we consider two subcases:

1a $p > q > r$ and $p \neq q + r$.

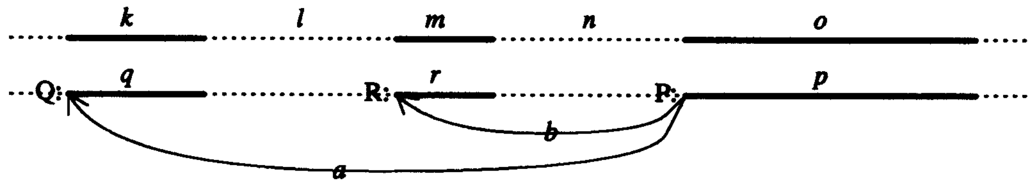

Figure 9: The (k, l, m, n, o) template.

1b $p > q > r$ and $p = q + r$.

2 $p > q = r$. In case 2 we consider two subcases:

2a $p > q = r$ and $p \neq q + r$.

2b $p > q = r$ and $p = q + r$.

3 $p = q > r$.

4 $p = q = r$.

Case 1 Suppose we want to tessellate the line with a three part tile as described. First we show that some configurations of blocks never appear in a tessellation.

QQ By this we mean that two instances are located such that their Q parts abut each other. This does not occur because the P parts of such instances would overlap.

RR For the same reason two R blocks never abut each other.

PRP In a PRP configuration (see figure 10), two instances have relative position $p + r$,

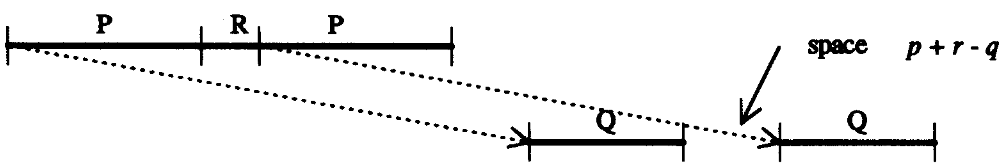

Figure 10: PRP is impossible.

leaving a space of $p + r - q$ between their Q parts. This gap is impossible to fill in, because it is smaller than a P block but larger than an R block. If more than one R block is laid in it, the corresponding P blocks overlap. If a Q block is laid in it, the corresponding P block overlaps with the original PRP configuration.

PQP In a PQP configuration (see figure 11) there is a space of size $(p + q - r)$ between the R blocks corresponding with the P blocks, which cannot be filled in for similar reasons.

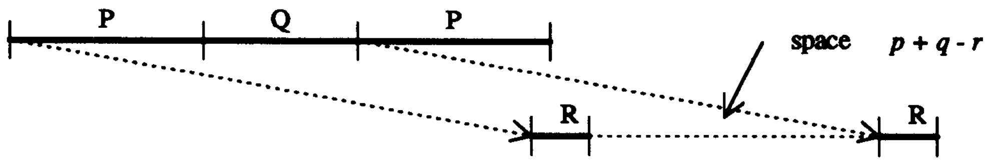

Figure 11: PQP is impossible.

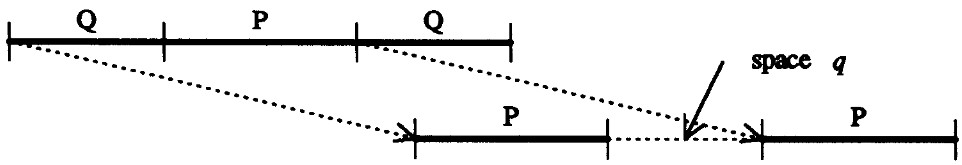

Figure 12: QPQ is impossible.

QPQ If a QPQ configuration occurs (see figure 12), there are two P blocks with q space in between. This gap can be filled only with a Q block, thus forcing the impossible configuration PQP.

RPR For the same reasons an RPR configuration forces a PRP configuration.

Case 1a If we are in case 1a ($p \neq q + r$) there are some more impossible configurations:

PP If a PP configuration occurs (see figure 13), there is a $(p - q)$ gap between the corresponding Q blocks, which is impossible to fill.

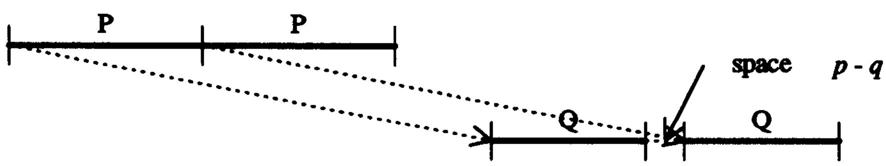

Figure 13: Case 1a: PP is impossible.

RQR In an RQR configuration (see figure 14) two instances have relative position $q + r$ so that either between the two P blocks there would be a gap smaller than r, or these two blocks would overlap.

QRQ In a QRQ configuration also two instances have relative position $q + r$.

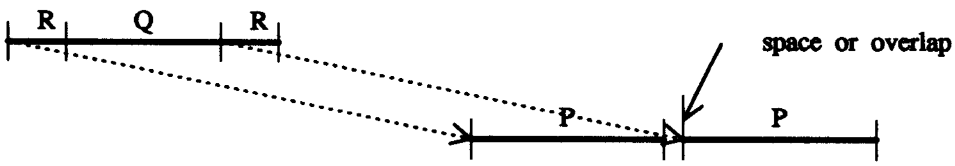

Figure 14: RQR is impossible.

But with these 9 configurations being impossible, there are only two possible infinite configurations of blocks (necessary to fill the line):

(1) ... P Q R P Q R P Q R P Q R ...
(2) ... P R Q P R Q P R Q P R Q ...

Obviously this can be done if and only if $a \equiv p \pmod{c} \wedge b \equiv p + q \pmod{c}$ (for tessellation (1)) or $b \equiv p \pmod{c} \wedge a \equiv p + r \pmod{c}$ (for tessellation (2)). Thus case 1a is now completely solved.

Case 1b For case 1b ($p > q > r$, $p = q + r$) we derive some more impossible configurations. It is possible to have a consecutive sequence of P blocks (see figure 15), but not

Figure 15: Sequence of P blocks.

that this sequence is abut to a Q block on both sides, see figure 16. The instances that

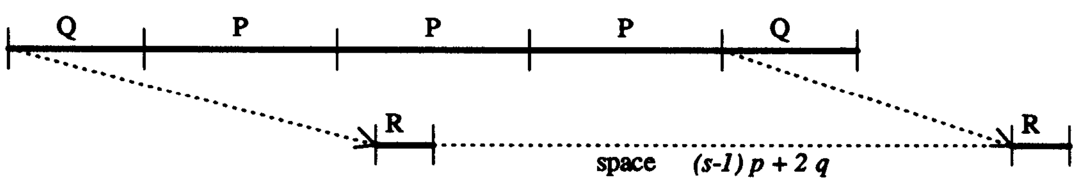

Figure 16: QPP $\cdots$ PPQ is impossible.

the Q blocks are part of have relative position $q + sp$ for some s so that between their R blocks there would be an $sp + q - r$ space. To fill it, R blocks cannot be used (their Q blocks would overlap with the original configuration) and QQ or PQP configurations do not occur. Hence the only possibility to fill this gap is by another QPPPPQ configuration, this time with $s - 1$ P blocks. Repeating the argument forces a QPQ configuration.

A similar argument shows that RPPPPPR (any number of P's) is impossible. Finally a sequence of alternating R and Q blocks has even length (i.e., it does not start *and* end with a Q block or with an R block). If so (see figure 17) the surrounding P blocks have

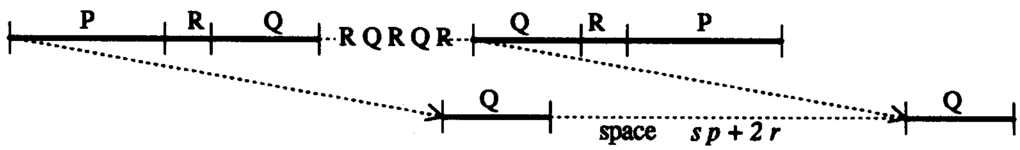

Figure 17: PRQR $\cdots$ QRP is impossible.

relative position $(s+1)p+r$, so between their Q blocks there is a space of $sp+2r$ which is impossible to fill.

We may now conclude, that in a tessellation the line is divided into "slots", each of length p, in which lies either a P block, an RQ pair, or a QR pair, but RQ pairs and QR pairs do not occur both. Let us consider tessellations in which QR pairs occur. (The analysis of the other possibility is similar.)

See figure 18. In some slots we have a P, in others a QR configuration. Of course we

Figure 18: The tessellation in case 1b.

must have

$$a \equiv 0 \pmod{p}, b \equiv q \pmod{p}$$

so let us say

$$a = a'p, b = b'p + q.$$

Let i be a slot with a P block. Then (see figure 19) there is a Q block in slot $i+a'$, hence

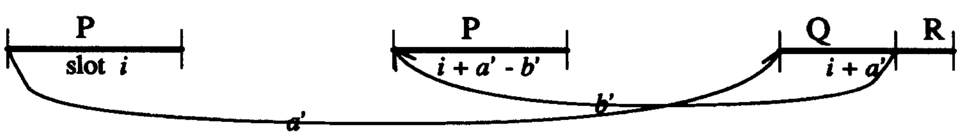

Figure 19: Periodicity of the tessellation.

there is also an R block in slot $i+a'$, which implies that there is a P block in slot $i+a'-b'$. It turns out that the pattern shown in figure 18 is periodic with a period of $a'-b'$ slots. Hence we can tessellate $\mathbf{Z}$ with T if and only if we can tessellate $\mathbf{Z}_{(a'-b')}$ with $\{\bar{0}, \bar{a'}\}$.

Lemma 4.2 *$\{\bar{0},\bar{x}\}$ tessellates $\mathbf{Z}_g$ iff $2|ord(\bar{x})$.*

Proof. Suppose $\{\bar{0},\bar{x}\}$ tessellates $\mathbf{Z}_g$, consider a tessellation in which one instance is based in $\bar{0}$. Then no instance can be based at $\bar{x}$, so the point $2\bar{x}$ is covered with an instance based at $s\bar{x}$. We find that an instance is based at $k\bar{x}$ if and only if $2|k$. We know one is based at $\bar{0} = ord(\bar{x})\bar{x}$, so $2|ord(\bar{x})$.

Suppose $2|ord(\bar{x})$. With instances based at $\bar{0}, 2\bar{x}, \ldots, (ord(\bar{x})-2)\bar{x}$, $x \times \mathbf{Z}_g$ is tiled. $x \times \mathbf{Z}_g$ tiles $\mathbf{Z}_g$, so $\{\bar{0},\bar{x}\}$ does so by theorem 2.8. □

So we find that a tessellation (of the P/QR type) is possible if and only if $2|ord(\bar{a'})$ (in $\mathbf{Z}_{(a'-b')}$). As $ord(\bar{a'}) = (a'-b')/\gcd(a', a'-b')$, this just says that $a'-b'$ contains more factors 2 than a' does. And that is equivalent to saying that a' and b' contain as many factors 2.

Thus case 1b is summarized as follows. The template tessellates the line if the following two conditions are both true:

- Either $a \equiv 0 \pmod p \wedge b \equiv q \pmod p$ (this gives the P/QR type tessellation) or $a \equiv r \pmod p \wedge b \equiv 0 \pmod p$ (this gives the P/RQ type tessellation).
- ($a \bmod p$) and ($b \bmod p$) contain the same number of 2's in their prime decomposition.

Case 2 We now turn our attention to the analysis of case 2: $p > q = r$. This analysis is done using the same approach as for case 1. First we identify some impossible configurations:

QQ This configuration is impossible for the same reason as it was in case 1, as its impossibility resulted only from the fact that $p > q$.

RR The same applies here.

PRP If a PRP configuration starts at position x (see figure 20), there is a gap of size p at position $x + a + q$ between the corresponding Q blocks. It cannot be filled

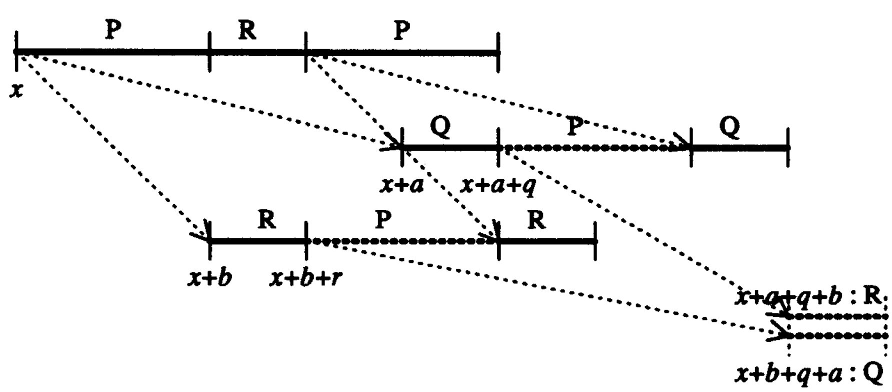

Figure 20: Case 2: PRP is impossible.

using Q blocks or using more than one R block (the corresponding P blocks would overlap) so it is filled with a P block, which has its corresponding R block at position $x + a + q + b$. However between the R blocks (corresponding to the two P blocks in the original configuration) there is also a space of size p, at position $x + b + r$, which can also be filled only with a P block for similar reasons. The corresponding Q block is at position $x + b + r + a$, which is equal to $x + a + q + b$, hence it overlaps with the forementioned R block.

PQP The configuration PQP is impossible for similar reasons.

RPR If an RPR configuration occurs there is a space of size r between the corresponding P blocks and hence this configuration forces a PRP or PQP configuration.

QPQ This configuration is impossible for the same reason.

Case 2a In case 2a ($p > q = r$ and $p \neq q + r$) the configurations PP, RQR, QRQ are also impossible for the same reason as in case 1a. Reasoning as in case 1a, we find that the template tessellates the line if and only if $a \equiv p \pmod{c} \wedge b \equiv p + q \pmod{c}$ or $a \equiv p + r \pmod{c} \wedge b \equiv p \pmod{c}$.

Case 2b For case 2b ($p > q = r$ and $p = q + r$) the analysis is exactly the same as for case 1b and the result is the same. The template tessellates the line if the following two conditions are both true:

- Either $a \equiv 0 \pmod{p} \wedge b \equiv q \pmod{p}$ (this gives the P/QR type tessellation) or $a \equiv r \pmod{p} \wedge b \equiv 0 \pmod{p}$ (this gives the P/RQ type tessellation).
- ($a \bmod p$) and ($b \bmod p$) contain the same number of 2's in their prime decomposition.

Case 3 For case 3 ($p = q > r$) we can again derive some impossible configurations.

RR The corresponding P blocks would overlap.

PP There would be a $(p - r)$-gap between the corresponding R blocks, which is impossible to fill.

QQ There would be a $(p - r)$-gap between the corresponding R blocks, which is impossible to fill.

RPR In the configuration RPR there are two instances at relative position $r + p$. Say the configuration starts at position x (see figure 21). There is a gap of size r between the corresponding P blocks (at position $x - b + p$) and between the corresponding Q blocks (at position $x - b + p + a$). Both must be filled with an R block. But then both the Q block of the first one and the R block of the second one are located at position $x - 2b + a + p$ and hence they overlap.

RQR Again two instances have relative position $r + p$.

PRP Again two instances have relative position $r + p$.

QRQ Again two instances have relative position $r + p$.

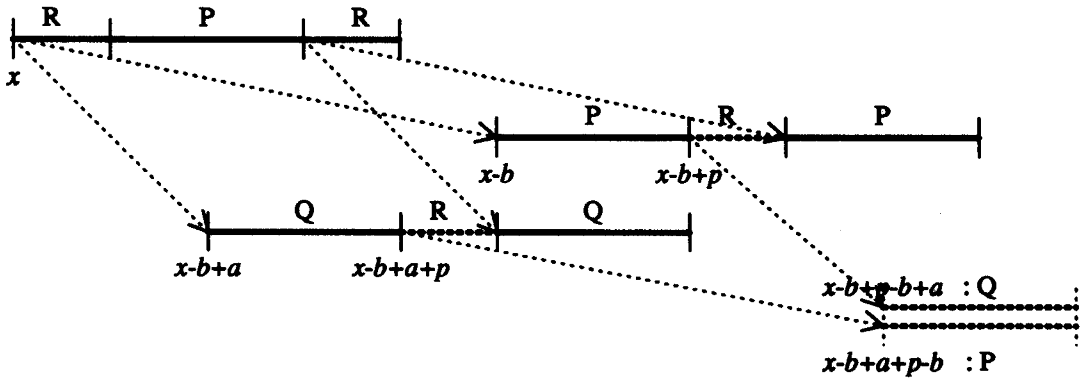

Figure 21: RPR is impossible.

QPQ In a QPQ configuration there is a $(2p-r)$-gap between the corresponding R blocks, which is impossible to fill.

PQP Impossible for the same reason.

We find that the same configurations are impossible as in case 1a, and as a result we find the same conditions again. The template tessellates the line if and only if $a \equiv p \pmod c \wedge b \equiv p+q \pmod c$ or $a \equiv p+r \pmod c \wedge b \equiv p \pmod c$.

Case 4 In case 4 ($p = q = r$) all blocks have the same length p and hence the distances between the blocks must be multiples of p to allow a tessellation. Suppose $a = a'p$ and $b = b'p$. The line can again be divided into slots, now containing a P, a Q, or an R block. If slot i contains a P, then slot $i + a'$ contains a Q and slot $i + b'$ contains an R block and vice versa. It now turns out that T tessellates the line if and only if $\{0, a', b'\}$ does.

When we take $a'' = a'/\gcd(a', b')$ and $b'' = b'/\gcd(a', b')$ then $\{0, a', b'\}$ tessellates the line if and only if $\{0, a'', b''\}$ does, this will be proved later (see corollary 5.2).

Lemma 4.3 *$T' := \{0, a'', b''\}$ tessellates $\mathbb{Z}$ if and only if $3|a'' + b''$.*

Proof. Suppose T' tiles $\mathbb{Z}$ and $\{T'(e) : e \in E\}$ is a partition of $\mathbb{Z}$. Each y is covered by exactly one instance in this set so *one* of y, $y - a''$, $y - b''$ is in E. So if $x \in E$, then $x + a'' \notin E$ and $x + b'' \notin E$. But then it follows that $x + a'' + b'' \in E$! We see that the tessellation is periodic with a period of $a'' + b''$. As there is an integer number of tiles in one period, the period length is a multiple of the size of the tile and hence $3|a'' + b''$.

Now suppose $3|a'' + b''$. As $\gcd(a'', b'') = 1$, this implies that 0, a'', and b'' are three different numbers modulo 3, and hence $\{T'(e) : e \in 3\mathbb{Z}\}$ is a partition of $\mathbb{Z}$. □

As a result we have for case 4 that T tessellates the line if and only if $p|a$ and $p|b$ and $3|a'' + b''$, where $a'' = a/\gcd(a, b)$ and $b'' = b/\gcd(a, b)$.

The results for all the cases are summarized in the following theorem.

Theorem 4.4 *$T = (k, l, m, n, o)$ tessellates the line if and only if one of the following three conditions is true.*

1. $a \equiv p \pmod{c} \wedge b \equiv p + q \pmod{c}$ *or* $a \equiv p + r \pmod{c} \wedge b \equiv p \pmod{c}$
2. $p = q + r$ *and either* $(a \equiv 0 \pmod{p} \wedge b \equiv q \pmod{p})$ *or* $(a \equiv r \pmod{p} \wedge b \equiv 0 \pmod{p})$ *and* $(a \operatorname{div} p)$ *and* $(b \operatorname{div} p)$ *contain the same number of 2's in their prime decomposition.*
3. $p = q = r$ *and* $p|a$ *and* $p|b$ *and* $3|a''+b''$, *where* $a'' = a/\gcd(a,b)$ *and* $b'' = b/\gcd(a,b)$.

Finally, for testing the completeness of the template we state the following result.

Proposition 4.5

$$\begin{aligned} Dif(k,l,m,n,o) &= \{0,1,\ldots,k-1\} \cup \{0,1,\ldots,m-1\} \cup \{0,1,\ldots,o-1\} \\ &\cup \{l+1,\ldots,k+l+m-1\} \\ &\cup \{n+1,\ldots,m+l+n-1\} \\ &\cup \{l+m+n+1,\ldots,k+l+m+n+o-1\} \end{aligned}$$

It is easy to see whether the intervals identified in proposition 4.5 together form $\{0, 1, \ldots, k + l + m + n + o - 1\}$ or not, and hence to test a (k, l, m, n, o) template for completeness.

5 Tessellating Z with arbitrary templates

Given a template $T = \{x_0, x_1, \ldots, x_{c-1}\}$, we want to decide whether it tessellates $\mathbf{Z}$ or not. Some useful observations can be made. As usual we assume $x_0 = 0$. Further we can assume that $\gcd(x_0, \ldots, x_{c-1}) = 1$. Write kT for the set $\{kx_0, kx_1, \ldots, kx_{c-1}\}$.

Lemma 5.1 *T tessellates $\mathbf{Z}$ if and only if kT does, for any $k \geq 1$.*

Proof. Suppose T tessellates the line and $\{T(e) : e \in E\}$ is a partition of $\mathbf{Z}$. Then $\{kT(ke) : e \in E\}$ is a partition of $k\mathbf{Z}$ and hence we see that kT tessellates $k\mathbf{Z}$. As $k\mathbf{Z}$ tessellates $\mathbf{Z}$ (base copies at $0, 1, \ldots, k-1$), T does so by theorem 2.8.

Suppose that kT tessellates $\mathbf{Z}$. Because $k|x$ for each $x \in kT$, each instance of kT is either a subset of $k\mathbf{Z}$ or disjoint with it. The instances that are a subset of $k\mathbf{Z}$ completely tessellate it. But if $\{kT(e) : e \in E\}$ is a partition of $k\mathbf{Z}$ then all elements of E are multiples of k and $\{T(e) : ke \in E\}$ is a partition of $\mathbf{Z}$, hence T tessellates $\mathbf{Z}$. □

Corollary 5.2 *Let $d = \gcd(x_0, x_1, \ldots, x_{c-1})$. Then $\{x_0, x_1, \ldots, x_{c-1}\}$ tessellates the line iff $\{x_0/d, x_1/d, \ldots, x_{c-1}/d\}$ does.*

Of course we have $\gcd(x_0/d, x_1/d, \ldots, x_{c-1}/d) = 1$. We will now define the decomposition of a template.

Definition 5.3 *Write $x_i = q_i \times c + r_i$, with $r_i < c$. So $T = \{q_0 \times c + r_0, q_1 \times c + r_1, \ldots, q_{c-1} \times c + r_{c-1}\}$. Let $T_i = \{q_j : r_j = i\}$. The* decomposition *of T is $(T_0, \ldots, T_{c-1})$.*

As an example, consider $T = \{0,2,7,12,14,19\}$. Write T as $\{0 \times 6 + 0, 0 \times 6 + 2, 1 \times 6 + 1, 2 \times 6 + 0, 2 \times 6 + 2, 3 \times 6 + 1\}$ and in this case we have $T_0 = \{0,2\}$, $T_1 = \{1,3\}$, $T_2 = \{0,2\}$, and $T_3 = T_4 = T_5 = \emptyset$. The decompositions can be used to find a tessellation according to the following lemma (decomposition lemma).

Lemma 5.4 *Let $(T_0, \ldots, T_{c-1})$ be the decomposition of T and*

1. *there is a set E such that for all $i \in 0 \ldots c-1$ either T_i is empty or $\{T_i(e) : e \in E\}$ is a partition of $\mathbf{Z}$, and*

2. *the set $\{\bar{i} : T_i$ is not empty$\}$ tiles $\mathbf{Z}_c$.*

($\bar{i}$ denotes the equivalence class of i modulo c. Note that the first condition is stronger than the condition that each of the T_i tessellates the line. Each of the T_i must tessellate the line with *the same* set of basepoints.)

Proof. For all nonempty T_i, cT_i tiles $c\mathbf{Z}$ using the basepoint set cE. So, using cE as the basepoint set, T tiles $\{x : (x \bmod c) \in \{i : T_i \text{ is nonempty}\}\}$. The second condition implies that this infinite "supertile" tiles the line. So (by 2.8) T tessellates the line. □

The conditions of the lemma are true for the abovementioned example (see figure 22). For (1) take $E = \{\ldots, 0, 1, 4, 5, 8, 9, 12, 13, \ldots\}$. For (2), note that two copies of $\{\bar{0}, \bar{1}, \bar{2}\}$,

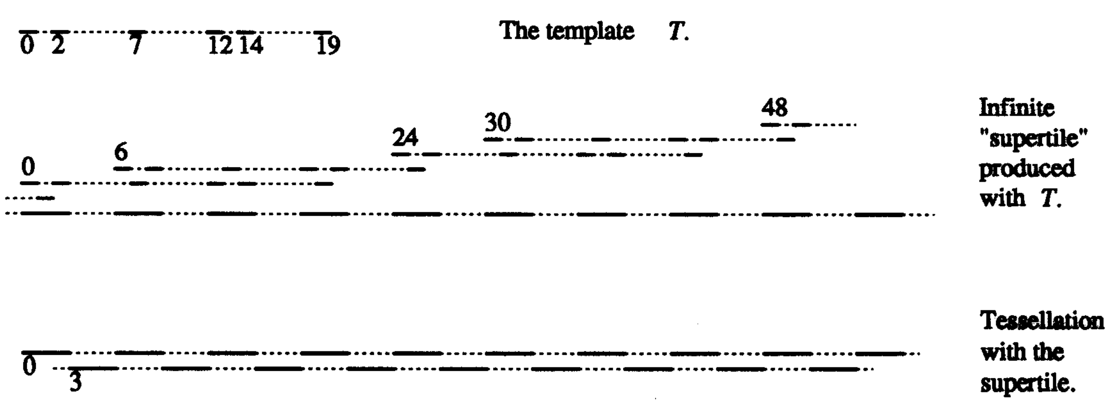

Figure 22: Example to the decomposition lemma.

based at $\bar{0}$ and $\bar{3}$, tile $\mathbf{Z}_6$. Now using $6E = \{\ldots, 0, 6, 24, 30, 48, 54, \ldots\}$ as a basepoint set for T yields the supertile as described in the lemma and two copies of it (based at 0 and 3) tile the line.

Lemma 5.4 states a sufficient condition for tessellation, but we do not know whether this condition is necessary. It is necessary for all classes of templates for which the tiling problem was solved in this paper, as is shown by the following propositions.

Proposition 5.5 *If $T = \{0, a\}$ tessellates the line then conditions (1) and (2) of lemma 5.4 hold.*

Proof. A two-cell template always tessellates the line. Thus we must show that the conditions are always true. If a is even then $T_0 = \{0, a\text{div}2\}$ and $T_1 = \emptyset$. T_0 tessellates the line and $\{\bar{0}\}$ tessellates $\mathbf{Z}_2$, so (1) and (2) hold. If a is odd, $T_0 = \{0\}$ and $T_1 = \{a\text{div}2\}$. Both tessellate the line with $E = \mathbf{Z}$ as basepoint set and $\{\bar{0}, \bar{1}\}$ tessellates $\mathbf{Z}_2$. □

Proposition 5.6 *If $T = (k, l, m)$ tessellates the line then conditions (1) and (2) of lemma 5.4 hold.*

Proof. We know from theorem 3.3 that T tessellates iff $k + m | l$ or $k = m \wedge k | l$. First suppose $c | l$, say $l = qc$. Then $T_i = \{0\}$ if $i < k$ and $T_i = \{q\}$ if $i \geq k$. For condition (1), take $E = \mathbf{Z}$. For condition (2), note that $\{\bar{i} : T_i \text{ is nonempty}\} = \mathbf{Z}_c$. Next suppose $k = m$ and $k | l$. If l/k is even then $c | l$ and it follows from the above that the conditions are true. So assume that $l = (2q+1)k$. Then $T_i = \{0, q+1\}$ if $i < k$ and $T_i = \emptyset$ if $i \geq k$, and again the conditions hold. □

Proposition 5.7 *If $T = \{0, a, b\}$ tessellates the line then conditions (1) and (2) of lemma 5.4 hold.*

Proof. If $3|a$ and $3|b$ then $T_0 = \{0, a/3, b/3\}$ tessellates the line (see lemma 5.1) and $T_1 = T_2 = \emptyset$. So the conditions (1) and (2) hold.

If $3 \not| a$ or $3 \not| b$ then $3 | a + b$ (see lemma 4.3) and hence a and b are 1 and 2 modulo 3. Thus every component of the decomposition contains one element, and the conditions are true. □

Proposition 5.8 *If $T = (k, l, m, n, o)$ tessellates the line then conditions (1) and (2) of lemma 5.4 hold.*

Proof. Use theorem 4.4 and reason as in the proof of proposition 5.6. □

If the elements in T have a factor in common T may tessellate while the conditions are not true. For example, we saw that $T = \{0, 2, 7, 12, 14, 19\}$ tessellates the line, and hence (by 5.1) $3T = \{0, 6, 21, 36, 42, 57\}$ does. However $3T$ has only two nonempty components in its decomposition namely $T_0 = \{0, 1, 6, 7\}$ and $T_3 = \{3, 9\}$. As T_0 and T_3 contain a different number of cells, condition (1) does not hold.

The discussion so far can be summarized in the following tessellation test:

ALGORITHM A: Start with template $T = \{x_0, x_1, \ldots, x_{c-1}\}$.

1. Subtract x_0 from all x_i. Divide all x_i by $\gcd(x_0, x_1, \ldots, x_{c-1})$. (This produces a "standard" template with $x_0 = 0$ and $\gcd(x_0, x_1, \ldots, x_{c-1}) = 1$.)
2. Compute the decomposition $(T_0, T_1, \ldots, T_{c-1})$.
3. Test whether there is a set E such that for all $i \in 0 \ldots c-1$ either T_i is empty or $\{T_i(e) : e \in E\}$ is a partition of $\mathbf{Z}$.
4. Test whether the set $\{\bar{i} : T_i \text{ is not empty}\}$ tiles $\mathbf{Z}_c$.

Output "yes" and a tessellation if a positive answer is found in steps (3) and (4), and "no" otherwise.