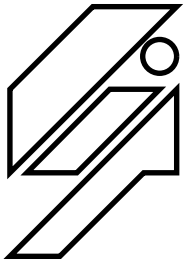

Laboratoire de l'Informatique du Parallélisme

Ecole Normale Supérieure de Lyon
Unité de recherche associée au CNRS n°1398

A Characterization of One-to-One Modular Mappings

Alain Darte
Michèle Dion
Yves Robert

April 1995

Research Report N^{O} 95-09

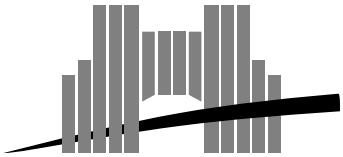

Ecole Normale Supérieure de Lyon
46 Allée d'Italie, 69364 Lyon Cedex 07, France
Téléphone : (+33) 72.72.80.00 Télécopieur : (+33) 72.72.80.80
Adresse électronique : lip@lip.ens−lyon.fr

A Characterization of One-to-One Modular Mappings

Alain Darte
Michèle Dion
Yves Robert

April 1995

Abstract

In this paper, we deal with *modular mappings* as introduced by Lee and Fortes [14, 13, 12], and we build upon their results. Our main contribution is a characterization of one-to-one modular mappings that is valid even when the source domain and the target domain of the transformation have the same size but not the same shape. This characterization is constructive, and a procedure to test the injectivity of a given transformation is presented.

Keywords: automatic parallelization, loop nests, time-space transformation, modular mapping, injectivity, characterization

Résumé

Nous étudions dans ce rapport les placements modulaires tels qu'ils ont été introduits par Lee et Fortes [14, 13, 12], et nous développons les résultats qu'ils ont obtenus. Notre apport principal consiste en une caractérisation des placements modulaires bijectifs qui reste valide même lorsque les domaines source et cible de la transformation contiennent le même nombre de points mais n'ont pas la même forme. La caractérisation est constructive et nous présentons une procédure qui permet de tester l'injectivité d'une transformation.

Mots-clés: parallélisation automatique, nids de boucles, transformation temps-espace, placement modulaire, injectivité, caractérisation

A Characterization of One-to-One Modular Mappings

Alain Darte, Michèle Dion and Yves Robert *

Laboratoire LIP, URA CNRS 1398
Ecole Normale Supérieure de Lyon, F - 69364 LYON Cedex 07
e-mail: [Alain.Darte,Michele.Dion,Yves.Robert]@lip.ens-lyon.fr

Abstract

In this paper, we deal with *modular mappings* as introduced by Lee and Fortes [14, 13, 12], and we build upon their results. Our main contribution is a characterization of one-to-one modular mappings that is valid even when the source domain and the target domain of the transformation have the same size but not the same shape. This characterization is constructive, and a procedure to test the injectivity of a given transformation is presented.

1 Introduction

Recently, Lee and Fortes [14, 13, 12] have introduced *modular mappings* in the context of systolic array design methodologies and parallelizing compilation. Their idea is to extend affine mapping techniques by using linear transformations modulo a constant vector. Affine mappings are time-space transformations that have been used extensively by a variety of researchers to derive efficient time-space transformations for loop nest programs (see [1, 17, 6, 5, 7, 8, 20, 11, 15, 19, 21] among others).

However, the systematic derivation of programs that can take advantage of wraparound connectivity in networks such as rings and 2D- or 3D-torus remains out of the scope of affine mappings. A typical example is Cannon's matrix-matrix product algorithm on a 2D-torus of processors [3]: this well-known algorithm (whose counterpart in the systolic field is the Preparata-Vuillemin 2D-systolic array [18]) cannot be synthesized using affine transformations, whereas Lee and Fortes [14, 13] demonstrate how to synthesize it, as well as many interesting variants, using one-to-one modular mappings. We point out that many other BLAS3-like kernels have been implemented onto 2D processor meshes using wraparound connections (e.g. the scientific library of the MasPar [2, 4]). We refer to Section 2 for the automatic synthesis of Cannon's algorithm using modular mappings, thereby providing the reader with a complete example to demonstrate the usefulness of modular mappings.

This paper deals with the automatic derivation of one-to-one modular mappings. We build upon the results of Lee and Fortes, which we summarize in Section 3. In a word, Lee and Fortes give several sufficient conditions for a modular mapping to be *one-to-one*. Injectivity plays a key role as modular mappings represent a time-space transformation from an index domain (computation points) to a target domain: clearly, the number of computation points must be preserved by the mapping. There are two major limitations in the results of Lee and Fortes:

*Supported by the ESPRIT Basic Research Action 6632 "NANA2" of the European Economic Community.

i/j	0	1	2	3	4
0	$d_{0,0}/e_{0,0}$	$d_{0,1}/e_{1,1}$	$d_{0,2}/e_{2,2}$	$d_{0,3}/e_{3,3}$	$d_{0,4}/e_{4,4}$
1	$d_{1,1}/e_{1,0}$	$d_{1,2}/e_{2,1}$	$d_{1,3}/e_{3,2}$	$d_{1,4}/e_{4,3}$	$d_{1,0}/e_{0,4}$
2	$d_{2,2}/e_{2,0}$	$d_{2,3}/e_{3,1}$	$d_{2,4}/e_{4,2}$	$d_{2,0}/e_{0,3}$	$d_{2,1}/e_{1,4}$
3	$d_{3,3}/e_{3,0}$	$d_{3,4}/e_{4,1}$	$d_{3,0}/e_{0,2}$	$d_{3,1}/e_{1,3}$	$d_{3,2}/e_{2,4}$
4	$d_{4,4}/e_{4,0}$	$d_{4,0}/e_{0,1}$	$d_{4,1}/e_{1,2}$	$d_{4,2}/e_{2,3}$	$d_{4,3}/e_{3,4}$

Figure 1: Initial data alignement

- they only deal with modular transformations that map an index domain onto itself. In other words, the target domain is assumed to be the same as the index domain. Clearly, if the transformation is one-to-one, the index domain and the target domain should have the same size, but not necessarily the same shape,
- they only give sufficient conditions for a transformation to be one-to-one. Given an arbitrary modular mapping (possibly given by the programmer), it is not always possible to decide from their results whether the transformation is one-to-one or not. Necessary and sufficient conditions would be necessary. Also, a procedure to determine whether a given transformation is one-to-one would be highly desirable.

Our paper overcomes both limitations. Our main result is a necessary and sufficient condition for a modular mapping to be one-to-one. The condition is rather technical, but the proof is constructive, hence a procedure to accompany the systematic derivation of one-to-one modular mappings. The condition extends to mappings for which the index domain and the target domain have the same size but not the same shape.

The rest of the paper is organized as follows: in Section 2 we detail the use of modular mappings through the matrix-matrix product example. In Section 3 we formally define modular mappings, and we review the results obtained by Lee and Fortes. In Section 4 we give a necessary and sufficient condition for a modular mapping to be one-to-one. We discuss several extensions in Section 4.2. Section 5 is devoted to some final remarks and conclusions.

2 Why modular mappings ?

Several basic computational kernels require other type of transformations than affine time-space mappings. A well-known example is Cannon's algorithm for matrix-matrix multiplication (see [3]):

DO $i = 0, 4$
  DO $j = 0, 4$
    DO $k = 0, 4$
      $c(i, j) = c(i, j) + d(i, k) \times e(k, j)$
CONTINUE

In Cannon's algorithm, the data arrays d and e are first aligned and multiplied (elementwise) by each other as shown in Figure 1. The result of each mutiplication is stored in $c(i, j)$. At the next step, matrix d is shifted to the left and matrix e is shifted up. Elementwise multiplication takes

place and the result is added to the values of $c(i,j)$. The processus is repeated until all elements in a row of d are multiplied by all elements in a column of e.

Let us consider the following transformation T_b:

$$T_b((i,j,k)^t) = (\begin{pmatrix} -1 & -1 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} i \\ j \\ k \end{pmatrix})_{\mathrm{mod}(5,5,5)} = \begin{pmatrix} -i-j+k \bmod 5 \\ i \bmod 5 \\ j \bmod 5 \end{pmatrix}$$

This transformation, called modular mapping in [14, 13], transforms the previously described program into an equivalent one:

```
DO t = 0,4
   DOALL p1 = 0,4
      DOALL p2 = 0,4
         i = p1
         j = p2
         k = (t + p1 + p2)mod5
         c(i,j) = c(i,j) + d(i,k) × e(k,j)
         MOVE_WEST(d)
         MOVE_NORTH(e)
CONTINUE
```

Cannon's algorithm (except data movement) can therefore be described by a modular transformation applied to the original program. We refer the reader to the original papers of Lee and Fortes [14, 13] for several interesting variants of this standard parallelization, as well as for a method to derive data communications.

3 Review of Lee and Fortes results

3.1 Definitions

In this section, we use the same definitions and notations as in Lee and Fortes [14, 13]. Let $u = (u_1, \ldots, u_n)^t \in \mathbb{Z}^n$ be a vector with n integer components, and let $m = (m_1, \ldots, m_n)^t \in (\mathbb{N}^*)^n$ be a vector with n positive integer components. The notation $u_{\mathrm{mod}m}$ denotes the vector $(u_1 \bmod m_1, \ldots, u_n \bmod m_n)^t$.

Definition 1 (Modular function) *A modular function $T_m : \mathbb{Z}^n \longrightarrow \mathbb{Z}^n$ is defined as $T_m(p) = (Tp)_{\mathrm{mod}m}$ for $p \in \mathbb{Z}^n$, where T is a $n \times n$ integer matrix (the* transformation matrix*) and $m \in (\mathbb{N}^*)^n$ is a n-vector (the* modulus vector*).*

Definition 2 (Modular time-space transformation of an index domain) *A* modular time-space transformation*, T_m, is a modular function that is* injective *when its domain is restricted to the index set J of an algorithm, i.e, $T_m : J \to \mathbb{Z}^n$ is injective.*

Definition 3 (Rectangular index set and boundary vector) *An index set J is rectangular and denoted J_b if $J = \{p \in \mathbb{Z}^n, 0 \leq p < b\}$, where inequalities between n-vectors are taken componentwise. The vector b is called the* boundary vector *of J_b.*

Definition 4 (Smith normal form) *For every matrix A of $\mathbb{Z}^n$, there exist two unimodular matrices Q_1, Q_2 and a diagonal matrix S such that:*

- $S = diag(s_1, s_2, \cdots, s_r, 0, 0, \cdots, 0)$

  where r is the rank of A, $s_1, s_2, \cdots, s_r$ are non-zero elements of $\mathbb{Z}$ and $s_i | s_{i+1}, 1 \leq i < r$.

- $A = Q_1 S Q_2$

The matrix S is then denoted by $S(A)$.

Definition 5 (Left Hermite form) *For every non singular matrix A of $\mathbb{Z}^n$, there exist a unimodular matrix Q and a lower triangular matrix H such that:*

- $\forall (i,j), h_{ij} >= 0$,
- *each non-diagonal element is lower than the diagonal element of the same row,*
- $A = HQ$.

Besides, this decomposition is unique up to a permutation of the rows. In fact, the row order used to "triangularize" A into H is arbitrary, hence there are $n!$ left Hermite forms.

An important remark Consider the modular transformation T_m with transformation matrix T and modulus vector m. It is important to point out that the coefficients of T are defined only up to a modulus operation. More precisely, let T' be the new transformation matrix defined by $T' = (t'_{ij})$ where $t'_{ij} = t_{ij} \bmod m_i$: then $T'_m = T_m$. The proof is immediate: $\forall (x_1, \cdots, x_n) \in \mathbb{Z}^n, \sum_j t_{ij} x_j \bmod m_i = \sum_j t'_{ij} x_j \bmod m_i$.

Similarly, the determinant of T is defined modulo the product $d = \prod_{i=1}^n m_i$. In particular, we can always assume that T is non singular (add a suitable multiple of d to each diagonal element t_{ii}, say, to get an equivalent non singular transformation matrix).

3.2 Main results of Lee and Fortes

In [14, 13], Lee and Fortes restrict themselves to the study of modular mappings for which the *modulus vector* is equal to the *boundary vector*, i.e. $m = b$. The case where $m = b$ is very important in practice, as the matrix-matrix product example demonstrates.

Lee and Fortes start with the following lemma:

Lemma 1 *Let $J_b = \{p \in \mathbb{Z}^n,\ 0 \leq j < b\}$ be a rectangular domain and define $\hat{J}_b = \{p \in \mathbb{Z}^n,\ -b < p < b\}$. A modular function $T_b : J_b \longrightarrow \mathbb{Z}^n$ is injective if and only if $T_b(p) \neq 0$ for all $p \in \hat{J}_b$ except $p = 0$.*

Proof See [13]. If T_b is not injective, there exist two distinct points $p, q \in J_b$ such that $T_b(p) = T_b(q)$. Then $r = p - q \in \hat{J}_b$, $r \neq 0$ and $T_b(r) = 0$.

Conversely, if there exists $r \in \hat{J}_b$, $r \neq 0$ and $T_b(r) = 0$, let p be defined as follows: $p_i = r_i$ if $0 \leq r_i < b_i$ and $p_i = 0$ if $-b_i < r_i < 0$. Let $q = p - r$. Then $p, q \in J_b$, $p \neq q$ and $T_b(p) = T_b(q)$, hence T_b is not injective. ■

Then, Lee and Fortes deal with *generator matrices*. They consider the set of integer points that are equivalent to zero, i.e, the equivalence class

$$S^0 = \{p \in \mathbb{Z}^n,\ T_b(p) = 0\}.$$

They prove that S^0 is a module, and that there exists a $n \times n$ integer matrix G that generates S^0: this means that every element of S^0 can be represented as an integer linear combination of the columns of G. Of course, there are several matrices that generate S^0, but they all are right equivalent. Indeed, let G be a generator matrix, then a matrix G' will generate S^0 if and only if there exists a $n \times n$ unimodular matrix U such that $G' = GU$.

The main contribution of [14, 13] on generators is a sufficient condition on the generators of S^0 that guarantees the injectivity of the transformation:

Lemma 2 *Let J_b be a rectangular index set with boundary vector b. Let T_b be a modular mapping and let G be a generator of S^0. Let $\succ$ be an arbitrary order on the set $\{1, 2, \cdots, n\}$. T_b is injective if G satisfies the following equations:*

1. $g_{ii} = b_{ii}$,
2. $g_{ij} = 0$ *if* $i \succ j$.

From this sufficient condition on generators, Lee and Fortes investigate the relationship between generator matrices G and transformation matrices T. They deduce the following sufficient conditions for a modular mapping T_b to be injective:

Theorem 1 *Let J_b be a rectangular index set with boundary vector b. Let T_b be a modular mapping. Let $\succ$ be an arbitrary order on the set $\{1, 2, \cdots, n\}$. $T_b : J_b \longrightarrow \mathbb{Z}^n$ is injective if the matrix T satisfies to the following equations:*

1. $t_{ii} \wedge b_i = 1$ [1]
2. $t_{ij} = 0$ *if* $i \succ j$

We restate Theorem 1 as follows: if T is triangular up to a permutation, and if its i-th diagonal entry is relatively prime with b_i for all $1 \leq i \leq n$, then T_b is a time-space transformation of J_b. It turns out that Theorem 1 can be proven without making use of generators, as shown by the following direct proof.

Another proof of Theorem 1 Let T be upper triangular (without loss of generality). We solve the system $Tx = 0 \bmod b$ for $x \in \hat{J}_b$, where T is upper triangular. The last equation is

$$t_{nn} \times x_n = 0 \bmod b_n,$$

hence b_n divides $t_{nn}x_n$. Since b_n is relatively prime with t_{nn}, b_n divides x_n, which implies $x_n = 0$ as $x \in \hat{J}_b$. The $(n-1)$-th equation gives

$$t_{n-1,n-1} \times x_{n-1} + t_{n-1,n} \times 0 = 0 \bmod b_{n-1},$$

hence $x_{n-1} = 0$ just as before, and continuing the process we find $x = 0$. ■

Therefore we do not need to make use of generator matrices to prove Theorem 1. However, generator matrices will enable us to characterize one-to-one modular mappings, i.e. to give a necessary and sufficient condition for injectivity, as we show below in Section 4.

Finally, Lee and Fortes give a necessary and sufficient condition in the case where all entries of the boundary vector b have the same value β: in this particular case, they show that T_b is injective on J_b if and only if the determinant of T and β are relatively prime. We extend this result in Section 4.2.1.

[1] We write $\gcd(u, v) = u \wedge v$

4 New results

4.1 Characterization when $m = b$

In this section, we consider as Lee and Fortes that $m = b$ (the *modulus vector* is equal to the *boundary vector*). For the general case, see Section 4.2.

As we have seen in Lemma 2, Lee and Fortes [14] exhibit sufficient conditions on generator matrices G to obtain one-to-one modular mappings. We prove here that these conditions are also sufficient, and we give a constructive method to check the injectivity of a given modular mapping. In the following, we denote by Θ the matrix $diag(b_1, \cdots, b_n)$.

Lemma 3 *If G is a generator matrix of S^0, then* $\det(G)$ *divides* $\det(\Theta) = b_1 b_2 \cdots b_n$.

Proof We are going to exhibit a generator matrix for S^0. Let us consider a point $p \in S^0 : \exists k \in \mathbb{Z}^n$ such that $Tp = \Theta k$. Let $\overline{\Theta} = diag(\overline{\Theta}_i)$ be the comatrix of Θ and $d = \det \Theta = \prod_{i=1}^{n} b_i$ ($\overline{\Theta}$ is defined such that $\Theta\overline{\Theta} = d \times I_n$ where I_n is the identity matrix of order n, $\overline{\Theta}_i = \prod_{j \neq i} b_j$). We have:

$$\begin{aligned} \overline{\Theta}Tp &= \overline{\Theta}\Theta k \\ \overline{\Theta}Tp &= dk \\ Q_1 S(\overline{\Theta}T) Q_2 p &= dk \end{aligned}$$

where Q_1, Q_2 are two unimodular matrices and $S(\overline{\Theta}T)$ is the Smith normal form of matrix $\overline{\Theta}T$.

$$\begin{aligned} S(\overline{\Theta}T) Q_2 p &= d{Q_1}^{-1} k \\ S(\overline{\Theta}T) Q_2 p &= dk' \end{aligned}$$

where $k' \in \mathbb{Z}^n$ (Q_1 is unimodular).

Let $S(\overline{\Theta}T) = diag(s_i)$:

$$s_i (Q_2 p)_i = dk'_i$$

We want only integer values for the components of p, therefore

$$(Q_2 p)_i = \frac{d}{\gcd(d, s_i)} k''_i$$

where $k'' \in \mathbb{Z}^n$.

$$p = {Q_2}^{-1} S' k''$$

where $S' = diag(\frac{d}{\gcd(d,s_i)})$. The matrix ${Q_2}^{-1} S'$ generates S^0.

Besides, if we let $S(\Theta) = diag(\Theta'_i)$ and $S(\overline{\Theta}) = diag(\overline{\Theta}'_i)$, we know that $\overline{\Theta}'_i \Theta'_{n-i+1} = d$ (see [16] p.40) and that $\overline{\Theta}'_i$ divides s_i (if A and B are two nonsingular integer $n \times n$ matrices, then the k-th element $s_k(AB)$ of the Smith normal form of AB is divisible by $s_k(A)$ and $s_k(B)$, see [16] p.33). $\overline{\Theta}'_i$ divides s_i and d, so $\gcd(d, s_i) = \overline{\Theta}'_i u_i$ where $u_i \in \mathbb{Z}$ and $s'_i = \frac{d}{\overline{\Theta}'_i u_i}$. Therefore

$$\begin{aligned} \det(S') &= \prod \frac{d}{\overline{\Theta}'_i u_i} \\ \det(S') &= \frac{d^n}{\det(\overline{\Theta}) \prod u_i} \end{aligned}$$

Besides, $\det(\overline{\Theta}) = d^{n-1}$. Thus,

$$\det(S') = \frac{d}{\prod u_i}$$

Since all generator matrices are right equivalent, they all have the same determinant as ${Q_2}^{-1}S'$, hence as S'. ■

Lemma 4 *Let G be a finite abelian group. Let $[g]_q$ be the subset $\{0, g, \ldots, (q-1)g\}$ with $g \in G$ and $1 < q \leq order(g)$. Let $S_1, \ldots, S_k$ be k subsets of G. G is said to be the direct sum of the S_i, which we denote as $G = S_1 \oplus \cdots \oplus S_k$, if the mapping $(g_1, \ldots, g_k) \mapsto g_1 + \cdots + g_k$ from $S_1 \times \cdots \times S_k$ to G is one-to-one. If $G = [g_1]_{k_1} \oplus \cdots \oplus [g_r]_{k_r}$ then at least one of the $[g_i]_{k_i}$ is a subgroup of G.*

Proof This result has been proved by Hajós in its works on one of the Minkowski's conjecture, see [9].

Lemma 5 *If T_b is a one-to-one modular mapping on J_b, then $\forall x \in \mathbb{Z}^n$, there exists $(x_1, x_2) \in S^0 \times J_b$ such that $x = x_1 + x_2$, and this decomposition is unique.*

Proof We first prove the existence of such a decomposition and then its uniqueness.

Existence Let us consider the finite abelian group $A = \mathbb{Z}^n / S_0$. The matrix ${Q_2}^{-1}S'$ generates S^0 (see Lemma 3), so the number of elements of A is $\det({Q_2}^{-1}S') = \det(S')$. For $x \in \mathbb{Z}^n$, we denote by $\overline{x}$ the canonical image of x in A.

Let us consider two distinct elements of J_b, x and y, then $\overline{x} \neq \overline{y}$. Indeed, $x - y \in \hat{J}_b$, if $\overline{x} = \overline{y}$, $x - y \in S^0$ and T_b would be not injective (see Lemma 1). All elements of J_b have distinct canonical images in A.

The number of elements in J_b is $\det(\Theta)$, there are more elements in J_b than in A ($\det(S') \leq \det(\Theta)$, see Lemma 3). So, for all $x \in A$, there exists $y \in J_b$ such that x is the canonical image of y (otherwise, two elements in J_b would have the same canonical image in A, and this is impossible). Besides, this also means that if T_b is injective, we have $\det S' = \det \Theta$.

Consider $x \in \mathbb{Z}^n$. $\overline{x}$ is the canonical image of x in A; there exists $x_2 \in J_b$ such that $\overline{x}_2 = \overline{x}$, $x - x_2 \in S_0$. So, there exists $(x_1, x_2) \in S^0 \times J_b$ such that $x = x_1 + x_2$.

Uniqueness If there exists (x_1, x_2) and (x'_1, x'_2) in $S^0 \times J_b$ such that $x = x_1 + x_2 = x'_1 + x'_2$ then, $x_2 - x'_2 = x_1 - x'_1 \in S^0$. $x_2 - x'_2 \in \hat{J}_b$ and T_b is injective, so $x_2 = x'_2$ and $x_1 = x'_1$. ■

Lemma 6 *If T_b is a one-to-one modular transformation then there exists i, $1 \leq i \leq n$, such that $O + b_i \vec{e_i} \in S^0$ (where e_i is the i-th vector of the canonical basis of $\mathbb{Z}^n$).*

Proof Let $f_1, f_2, ..., f_n$ be the canonical images of $O + \vec{e}_1, O + \vec{e}_2, \ldots, O + \vec{e}_n$ in A. If T_b is injective, for all $x \in \mathbb{Z}^n$, there exist $a \in S^0$ and integers α_i, $0 \leq \alpha_i < b_i$, such that $x = a + \alpha_1 \vec{e}_1 + \cdots + \alpha_n \vec{e}_n$ and this decomposition is unique, see lemma 5. This means exactly that $A = [f_1]_{b_1} \oplus \cdots \oplus [f_n]_{b_n}$. Lemma 4 shows that one of the f_i satisfies $b_i f_i = 0$, i.e, $0 + b_i \vec{e}_i \in S^0$. ■

Theorem 2 *A transformation T_b is one-to-one if and only if there exists a left Hermite form of a generator matrix G of S^0 with diagonal $b_1, b_2, ..., b_n$.*

Proof The sufficient condition has already been proved in [14], see Lemma 2.

Necessary condition Let us consider a one-to-one modular mapping T_b. Lemma 6 gives an index i such that $O + b_i\vec{e_i} \in S^0$.

Let us consider $[f_i]_{b_i}$, subgroup of A. Let B be the finite abelian group $B = A/[f_i]_{b_i}$. We also know that $A = [f_1]_{b_1} \oplus \cdots \oplus [f_n]_{b_n}$. Let $f'_1, f'_2, \cdots, f'_{i-1}, f'_{i+1}, \cdots, f'_n$ be the canonical images of $f_1, f_2, \cdots, f_{i-1}, f_{i+1}, \cdots, f_n$ in B (the canonical image of f_i in B is 0).

Let us consider $x \in B$, there exists $y \in A$ such that x is the canonical image of y in B. Besides, $y = m_1 f_1 + \cdots + m_n f_n$ and this decomposition is unique. So, $x = m_1 f'_1 + \cdots + m_{i-1} f'_{i-1} + m_{i+1} f'_{i+1} + \cdots + m_n f'_n$ and the decomposition is also unique. This means that $B = [f'_1]_{b_1} \oplus \cdots \oplus [f'_{i-1}]_{b_{i-1}} \oplus [f'_{i+1}]_{b_{i+1}} \oplus \cdots \oplus [f'_n]_{b_n}$. There exists $j \neq i$ such that $b_j f'_j = 0$, i.e, there exists t, $0 \leq t < b_i$ such that $0 + b_j\vec{e_j} + t\vec{e_i} \in S^0$.

By repeating this process, we obtain n vectors in S^0 and the matrix H formed by these vectors is upper triangular with diagonal $b_1, b_2, \cdots, b_n$ (up to a permutation of indices).

Furthermore, these vectors form a basis of the module generated by G. By adding suitable combinations of the vector columns of H to any vector $x \in S^0$, we get $x = H\lambda + a$ where $\forall a_i, 0 \leq a_i < b_i$. Because there is only one element of S^0 in J_b (which is 0), we have $a = 0$ and thus, x is a linear combination of the column vectors of H. Furthermore, this decomposition is unique because H is non-singular. So, H is the matrix of a basis of S^0, this means that there exists a unimodular matrix Q such that $G = HQ$, and this completes the proof. ■

Given a transformation matrix T and a modulus vector b, Theorem 2 gives a constructive method to check whether T_b is a time-space transformation or not. We sketch the procedure and run it on an example.

Procedure From Theorem 2, a procedure to know whether a modular transformation is injective or not can be deduced:

1. Calculate the Smith normal form $\overline{\Theta}T$ and then deduce the matrix $Q_2^{-1}S'$ that generates S^0 (we use the same notations as in lemma 3).
2. Calculate the $n!$ left Hermite normal forms (by permuting the rows) of $Q_2^{-1}S'$.
3. If there exists a left Hermite normal form of $Q_2^{-1}S'$ with diagonal $b_1, b_2, \cdots, b_n$, the transformation T_b is injective.

Example Let us consider the matrix $T = \begin{pmatrix} 1 & 0 & 3 \\ 1 & 1 & 2 \\ 3 & 3 & 1 \end{pmatrix}$ and the vector $b = \begin{pmatrix} 5 \\ 4 \\ 6 \end{pmatrix}$.

We calculate the Smith normal form of $\overline{\Theta}T$ and we deduce the two matrices S' and Q_2^{-1}: $S' = \begin{pmatrix} 60 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ and $Q_2^{-1} = \begin{pmatrix} 3 & 152 & -613 \\ 0 & 0 & 1 \\ 1 & 51 & -204 \end{pmatrix}$, $Q_2^{-1}S' = \begin{pmatrix} 180 & 304 & -613 \\ 0 & 0 & 1 \\ 60 & 102 & -204 \end{pmatrix}$.

We calculate the 6 left Hermite forms of $Q_2^{-1}S'$ by permuting the rows. The left Hermite form of $\begin{pmatrix} 60 & 102 & -204 \\ 180 & 304 & -613 \\ 0 & 0 & 1 \end{pmatrix}$ is $\begin{pmatrix} 6 & 0 & 0 \\ 2 & 5 & 0 \\ 2 & 3 & 4 \end{pmatrix}$. T_b is injective.

4.2 Extensions

In this section, we start by proving a useful property that allows to restrict the search of one-to-one modular mappings to a more restricted set: we show that a transformation $T_{\lambda b}$ is injective on $J_{\lambda b}$ if and only if T_b in injective on J_b and $det(T) \wedge \lambda = 1$. Then, we consider the particular case where $\forall (i,j), b_i \wedge b_j = 1$. In this particular case, we have a necessary and sufficient condition directly with the transformation matrix. Finally, we extend the results given in Section 4 to a more general case: $m \neq b$, but $\prod m_i = \prod b_i$.

4.2.1 Injectivity of $T_{\lambda b}$

In this section, let T be a transformation matrix and b a modulus vector. We still assume that the source domain and the target domain are the same. We will prove a "scalability property". Beforehand, we prove the following lemma as a prerequisite for Theorem 3.

Lemma 7 $\det(T) \wedge \lambda \neq 1 \Rightarrow T_{\lambda b}$ *is not injective on* $J_{\lambda b}$.

Proof We use the notations of lemma 3. Let $d = \det \Theta = \prod_{i=1}^{n} b_i$, $S(\overline{\Theta}T) = diag(s_i)$, $S(\Theta) = diag(\Theta'_i)$ and $S(\overline{\Theta}) = diag(\overline{\Theta'}_i)$. Let Q_2 and S' be the matrices such that ${Q_2}^{-1}S'$ generates the module S^0 for T_b, Q'_2 and S'' the matrices such that ${Q'_2}^{-1}S''$ generates the module S^0 for $T_{\lambda b}$ (Q_2, S', Q'_2 and S'' are calculated as in the proof of lemma 3).

Let $S' = diag(s'_i)$ and $S'' = diag(s''_i)$. $S(\lambda^{n-1}\overline{\Theta}T) = \lambda^{n-1}S(\overline{\Theta}T)$, so we have:

$$s''_i = \frac{\lambda^n d}{\gcd(\lambda^n d, \lambda^{n-1} s_i)}$$

$$s''_i = \frac{\lambda d}{\gcd(\lambda d, s_i)}$$

We have seen in the proof of lemma 3 that $\overline{\Theta}'_i$ divides s_i. So, we can write $s_i = \overline{\Theta}'_i x_i$ with $\prod x_i = det(T)$ ($\prod s_i = \det(\overline{\Theta}T) = d^{n-1} \times \det(T) = \prod \overline{\Theta}'_i \prod x_i = d^{n-1} \prod x_i$).

Besides, $\overline{\Theta'}_i = \frac{d}{\Theta'_{n-i+1}}$. Thus,

$$s''_i = \frac{\lambda d}{\overline{\Theta}'_i \gcd(\lambda \Theta'_{n-i+1}, x_i)}$$

There exists i such that $\gcd(\lambda, x_i) \neq 1$ ($\prod x_i = det(T)$ and $\det(T) \wedge \lambda \neq 1$). Thus, $\prod s''_i < \prod \frac{\lambda b}{\overline{\Theta}'_i} = \lambda^n d$. Hence, $T_{\lambda b}$ cannot be injective (we see in the proof of Lemma 5 that if $T_{\lambda b}$ is injective, we must have $det(S'') = \det(\lambda\Theta) = \lambda^n d$). ■

Theorem 3 *Let* $\lambda \in \mathbb{N}$. *The modular mapping* $T_{\lambda b}$ *is a time-space transformation of* $J_{\lambda b}$ *if and only if* T_b *is a time-space transformation of* J_b *and* $\det(T) \wedge \lambda = 1$.

Proof Assume that $T_{\lambda b}$ is injective. Let $p \in \hat{J}_b$ such that $T_b(p) = 0$. Equivalently, $Tp = \Theta k$ for some $k \in \mathbb{Z}^n$. Then $T\lambda p = \lambda\Theta k$ and $\lambda p \in \hat{J}_{\lambda b}$. As $T_{\lambda p}$ is injective, Lemma 1 implies that $p = 0$. Hence, T_b is injective. Besides, we know from Lemma 7 that $\det(T) \wedge \lambda = 1$.

Conversely, assume now that T_b is injective and $\det(T) \wedge \lambda = 1$. If $\det(T) = 0$, the proof is immediate ($\det(T) = 0$ and $\det(T) \wedge \lambda = 1 \Rightarrow \lambda = 1$). Consider now the case $\det(T) \neq 0$.

Let p in $\hat{J}_{\lambda b}$ such that

$$Tp = \lambda\Theta k,\ k \in \mathbb{Z}^n$$

Let U be the comatrix of T ($UT = \det(T)I_n$). We have:

$$\det(T)p = \lambda U\Theta k$$

$$\det(T)p_i = \lambda(U\Theta k)_i$$

So, λ divides $\det(T)p_i$. But, we also have $\det(T) \wedge \lambda = 1$. Hence, λ divides p_i. Let us consider q such that $p_i = \lambda q_i$.

$$\det(T)\lambda q = \lambda U\Theta k$$

$$\det(T)q = U\Theta k$$

So, $\det(T)Tq = \det(T)\Theta k$ and $Tq = \Theta k$. Besides, $-\lambda b < p < \lambda b$ implies $-b < q < b$. T_b is injective, thus $q = 0$ and $p = 0$. ■

Remark The previous theorem leads to another proof of the following result of Lee and Fortes: in the case where all entries of the boundary vector b have the same value β, T_b is injective on J_b if and only if the determinant of T and β are relatively prime. Indeed, if $b = (1, \cdots, 1)^t$, J_b contains only the point $(0, \cdots, 0)$ and T_b is always injective. Therefore $T_{\beta b}$ is injective iff $\det(T) \wedge \beta = 1$.

4.2.2 When $\forall(i,j), b_i \wedge b_j = 1$

We know (see Section 3.1) that for any transformation T_m, there exists an equivalent transformation denoted as T'_m and such that $\forall(i,j), 0 \leq t'_{ij} < b_i$. We still assume $m = b$ here. We prove that if $\forall(i,j), b_i \wedge b_j = 1$, then T_b is one-to-one iff T is triangular (up to a permutation) with "good" diagonal coefficients: in this particular case, the characterization of one-to-one mappings is quite simple.

Theorem 4 *If $\forall(i,j), b_i \wedge b_j = 1$, then T_b is injective on J_b if and only if T'_b is an upper triangular matrix (up to a permutation on row and column indices) with $\forall i, t_{ii} \wedge b_i = 1$.*

Proof The sufficient condition has been proved in [13] (see Theorem 1) .

Necessary condition Assume that the transformation T_b is injective. The proof uses the same lemma as Theorem 2. Let us consider $t'_j, 1 \leq j \leq n$, the columns of T' and let A be the group $\mathbb{Z}/b_1\mathbb{Z} \times \mathbb{Z}/b_2\mathbb{Z} \times \cdots \times \mathbb{Z}/b_n\mathbb{Z}$. The restriction of T_b to J_b defines an injective application. The two sets have the same number of elements, so it is also bijective, *i.e*, $\forall x \in A, \exists! y \in J_b \ / \ x = T_b(y) = (Ty)_{\bmod b}$. This means exactly that $A = [t'_1]_{b_1} \oplus_b [t'_2]_{b_2} \oplus_b \cdots \oplus_b [t'_n]_{b_n}$.

Lemma 4 shows that there exists j such that $[t'_j]_{b_j}$ is a subgroup of A. There exists j such that $t'_j b_j = 0_{\mathrm{mod} b}$, *i.e* $\forall i, t'_{ij} b_j = 0 \bmod b_i$. So, we must have $\forall i \neq j, t'_{ij} = 0$ since $b_i \wedge b_j = 1$ and $t'_{jj} \wedge b_j = 1$ (otherwise the transformation would not be injective).

Up to a permutation on rows and columns, the matrix T' is now: $\left(\begin{array}{c|c} t'_{jj} & u \\ \hline 0 & \\ . & T'' \\ 0 & \end{array}\right)$ (where u is a row vector of $n-1$ elements). Consider the new modular transformation $T''_{b'}$, where $b' = (b_1, \cdots, b_{j-1}, b_{j+1}, \cdots, b_n)^t$. Let us prove that $T''_{b'}$ is an injective modular transformation on $J_{b'}$.

Let us consider $x = (x_1, \cdots, x_{j-1}, x_{j+1}, \cdots, x_n) \in \hat{J}_{b'}$ such that $T''x = 0_{\mathrm{mod} b'}$.

The element t'_{jj} has an inverse in $\mathbb{Z}/b_j\mathbb{Z}$ ($t'_{jj} \wedge b_j = 1$). Let $\alpha \in \mathbb{Z}/b_j\mathbb{Z}$ be the value $-{t'_{jj}}^{-1} u.x$, where ${t'_{jj}}^{-1}$ is the inverse of t'_{jj} in $\mathbb{Z}/b_j\mathbb{Z}$. We have $t'_{jj}\alpha + u.x = 0 \bmod b_j$. So, the vector $(x_1, \cdots, x_{j-1}, \alpha, x_{j+1}, \cdots, x_n) \in S^0$ and we have $\forall i \neq j, x_i = 0$ (T_b is injective).

We have just proved that $T''_{b'}$ is injective. So, in the same way, there exists k such that $T'' = \left(\begin{array}{c|ccc} t'_{kk} & . & . & . \\ \hline 0 & & & \\ . & & T''' & \\ 0 & & & \end{array}\right)$ with $t'_{kk} \wedge b_k = 1$, up to a permutation.

By repeating this process, we obtain that T' is triangular up to a permutation, and each element t'_{ii} on the diagonal satisfies $t'_{ii} \wedge b_i = 1$. ■

4.2.3 Extension to the general case: $m \neq b$

In this section, we consider a modular transformation T_m and a rectangular index set J_b and we prove that general results can be easily derived from the particular case $m = b$.

First, Lemma 1 remains satisfied in the general case.

Lemma 8 *A modular function $T_m : J_b \longrightarrow \mathbb{Z}^n$ is injective if and only if $T_m(p) \neq 0$ for all $p \in \hat{J}_b$ except $p = 0$.*

Proof The proof is immediate from the proof of Lemma 1. ■

Besides, if we consider the set of integer points that are equivalent to zero, $S^0 = \{p \in \mathbb{Z}^n,\ T_m(p) = 0\}$, we can find in the same way a generator matrix for S^0. Let $\Theta = diag(m_i)$. As in Section 4 and with the notations of lemma 1, we obtain a matrix ${Q_2}^{-1} S'$ that generates S^0 and that satisfies $\det(Q_2^{-1} S') | \det(\Theta)$ (simply replace b by m in Lemma 3).

Lemma 9 *If $\prod_{i=1}^n m_i = \prod_{i=1}^n b_i$, then Theorem 2 remains valid: a transformation T_m is one-to-one if and only if there exists a left Hermite form of a generator G of S^0 with diagonal $b_1, \cdots, b_n$.*

Proof The proof is immediate from the proof of lemma 2. The condition that $\prod_{i=1}^n m_i = \prod_{i=1}^n b_i$ is needed to prove that the sum of subsets used in Lemma 5 is a direct sum. Of course, if T_m is one-to-one from J_b onto J_m, both domains must have the same number of integer points. ■

In [13], Lee and Fortes dealt with the particular case when the modulus vector results from a permutation of the entries of the boundary vector. Lemma 9 is an extension of this particular case.

Example Let us consider the matrix transformation $T = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$ and the modulus vector $m = \begin{pmatrix} 2 \\ 3 \end{pmatrix}$. We have ${Q_2}^{-1}S' = \begin{pmatrix} 6 & 1 \\ 0 & 1 \end{pmatrix}$. Thus, T_m is injective on the rectangular index set $J_{(6,1)^t}$ but is not injective on $J_{(2,3)^t}$.

Lemma 9 is very useful as it enables to check injectivity for transformations that map a given rectangular domain onto a domain of different shape (but of the same size).

5 Conclusion

In this paper, we have considered *modular mappings* as introduced by Lee and Fortes [14, 13, 12]. Our main contribution is a characterization of one-to-one modular mappings that is valid even when the source domain and the target domain of the transformation have the same size but not the same shape. This characterization is constructive, and a procedure to test the injectivity of a given transformation has been presented.

We believe the study of modular mappings to be very promising in the context of automatic parallelization techniques. Indeed, mapping techniques usually proceed in two steps: first the input domain (computation points) is mapped onto a time-space domain where a virtual processor is assigned to each computation. Then virtual processors are mapped onto physical processors, most often using a block-cyclic allocation *à la* HPF [10]. Characterizing valid modular mappings from input domains onto target domains of larger dimension would enable to fully automatize the mapping procedure.

References

[1] Jennifer M. Anderson and Monica S. Lam. Global optimizations for parallelism and locality on scalable parallel machines. *ACM Sigplan Notices*, 28(6):112–125, June 1993.

[2] T. Blank. The MasPar MP-1 architecture. In *Compcon Spring*, pages 20–24. IEEE Press, February 1990.

[3] L.E. Cannon. *A cellular computer to implement the Kalman filter algorithm.* PhD thesis, Montana State University, 1969.

[4] P. Christy. Software to support massively parallel computing on the MasPar MP-1. In *Compcon Spring*, pages 29–33. IEEE Press, February 1990.

[5] Alain Darte and Yves Robert. Constructive methods for scheduling uniform loop nests. *IEEE Trans. Parallel Distributed Systems*, 5(8):814–822, 1994.

[6] Alain Darte and Yves Robert. Mapping uniform loop nests onto distributed memory architectures. *Parallel Computing*, 20:679–710, 1994.

[7] Paul Feautrier. Some efficient solutions to the affine scheduling problem, part I, one-dimensional time. *Int. J. Parallel Programming*, 21(5):313–348, October 1992. Available as Technical Report 92-28, Laboratoire MASI, Université Pierre et Marie Curie, Paris, May 1992.

[8] Paul Feautrier. Towards automatic distribution. *Parallel Processing Letters*, 4(3):233–244, 1994.

[9] G. Hajos. Uber einfache und mehrfache bedeckung des n-dimensionalen raumes mit einem wurfelgitter. *Math. Zeitschrift*, 47:427–467, 1942.

[10] Charles H. Koelbel, David B. Loveman, Robert S. Schreiber, Guy L. Steele Jr., and Mary E. Zosel. *The High Performance Fortran Handbook*. The MIT Press, 1994.

[11] S.Y. Kung. *VLSI array processors*. Prentice-Hall, 1988.

[12] Hyuk J. Lee and José A.B. Fortes. Data distribution independent parallel programs for matrix multiplication. Technical Report TR-EE 94-32, School of Electrical Engineering, Purdue University, October 1994.

[13] Hyuk J. Lee and José A.B. Fortes. Modular mappings of rectangular algorithms. Technical Report TR-EE 94-22, School of Electrical Engineering, Purdue University, June 1994.

[14] Hyuk J. Lee and José A.B. Fortes. On the injectivity of modular mappings. In Peter Cappello, Robert M. Owens, Jr Earl E. Swartzlander, and Benjamin W. Wah, editors, *Application Specific Array Processors*, pages 237–247, San Francisco, California, August 1994. IEEE Computer Society Press.

[15] D.I. Moldovan and J.A.B. Fortes. Partitioning and mapping algorithms into fixed-size systolic arrays. *IEEE Transactions on Computers*, 35(1):1–12, January 1986.

[16] Morris Newman. *Integral Matrices*. Academic Press, 1972.

[17] Michael O'Boyle and G.A. Hedayat. Data alignment: Transformations to reduce communications on distributed memory architectures. In *Scalable High-performance Computing Conference SHPCC-92*, pages 366–371. IEEE Computer Society Press, 1992.

[18] F.P. Preparata and J.E. Vuillemin. Area-time optimal VLSI networks for multiplying matrices. *Information Processing Letters*, 11(2):77–80, 1980.

[19] Patrice Quinton and Yves Robert. *Systolic Algorithms and Architectures*. Prentice Hall, 1991. Translated from French, Masson (1989).

[20] Weijia Shang and A.B. Fortes. Time optimal linear schedules for algorithms with uniform dependencies. *IEEE Transactions on Computers*, 40(6):723–742, June 1991.

[21] Michael E. Wolf and Monica S. Lam. A loop transformation theory and an algorithm to maximize parallelism. *IEEE Trans. Parallel Distributed Systems*, 2(4):452–471, October 1991.