


TR-93-006

Mapping Nested Loop Algorithms into Grid-Connected Systolic
Arrays without Data Collisions in the Data Links*

PeiZong Lee and Cheng-Fong Chen



中研院資訊所圖書室



3 0330 03 000366 4

Mapping Nested Loop Algorithms into Grid-Connected Systolic Arrays without Data Collisions in the Data Links*

PeiZong Lee[†] and Cheng-Fong Chen

Institute of Information Science
Academia Sinica
#128, Sec. 2, Yen-Chiu-Yun Road
Nankang, Taipei, Taiwan, R. O. C.

July 18, 1993

Abstract

Systolic arrays, which are made out of simple processing elements connected by data links, have made significant improvements in speeding up computation in comparison to conventional computers. Although systolic arrays belong to distributed memory parallel devices, they adopt systolic communications instead of message passing communications for sending data between neighboring processing elements. Therefore, it is important to provide necessary and sufficient conditions to avoid data collisions in the data links for a correct algorithm design. In this paper, we present necessary and sufficient conditions for mapping the class of shift-invariant uniform-dependence algorithms structured as nested loops into grid-connected systolic arrays of arbitrary dimensions. The proposed conditions, which are based on the *ZERO-ONE-INFINITE* property of tokens' behavior that describes how many times tokens are used and generated during the computation, can allow us to generate all feasible solutions.

Index Terms— algorithm transformation, data collision, data link, grid-connected systolic array, nested-loop algorithm.

*** A preliminary version of this technical report is accepted to be presented at the ISCA Sixth International Conference on Parallel and Distributed Computing Systems, Louisville, Kentucky, U.S.A., October 14-16, 1993.

*This work was partially supported by the NSC under Grants NSC 80-0408-E-001-02, NSC 81-0408-E-001-505, and 82-0408-E-001-016.

[†]PeiZong Lee: leepe@iis.sinica.edu.tw, TEL: +886 (2) 788-3799, FAX: +886 (2) 782-4814.

1 Introduction

It is significant that many time-intensive scientific algorithms are formulated as nested loops, which are inherently regularly structured. In this paper we study the mapping of shift-invariant uniform-dependence algorithms structured as nested loops into grid-connected systolic arrays without data collisions in the data links. Although previously there has been a significant amount of research studying the transformations of nested-loop algorithms into systolic arrays [3] [4] [12]–[19] [22] [23], the proposed methods either are not general enough for solving the mapping of nested-loop algorithms into arbitrary dimensional systolic arrays or cannot avoid the possibility of data collisions in data links. A complete survey of the systolic array transformations can be seen in [5] [10] [11].

In general, the transformations presented in these papers are all based on the hyperplane method [9]. This technique finds the data-dependence vectors of an algorithm first, and then finds a non-singular linear mapping which preserves the data-dependence ordering of the original algorithm. This technique is also called space-time mapping. However, previous attempts did not consider the cases when data might collide in the data links. As we will show in Section 5, transformations which satisfy conditions presented in other research papers may still incur data collisions in data links.

Although Lee and Kedem [11] have introduced a set of necessary and sufficient conditions for correctly transforming a p -nested-loop algorithm into a q -D systolic array, where D stands for dimensional and $1 \leq q < p$, yet in their model, a data link, in which a set of data tokens with a specific data-dependence vector flow through, has only one direction. Therefore, if a data token wants to flow to the diagonal PE (for example, the north-eastern processing element), there must be a data link connecting the current PE and the diagonal PE. This design will of course increase the complexity of the hardware and the fabrication cost of the target systolic array, because additional data links and their corresponding I/O pins are very expensive to be included in a chip with a limited area. Similarly, in Ganapathy and Wah's model, a data link also has only one direction [6], and they incur the same problem.

It is our goal in this paper to present a complete set of necessary and sufficient conditions for mapping nested-loop algorithms with uniform-dependence relations into grid-connected systolic arrays of arbitrary dimensions, including 1-D linear systolic arrays or 2-D mesh-connected systolic arrays or high dimensional ones. It is of practical importance to develop such necessary and sufficient conditions for a correct mapping. Because the topology of a grid-connected systolic array is regular, in addition, each PE of a q -D grid-connected systolic array has only $2q$ data links. This architecture is thus especially suitable for the VLSI

(Very Large Scale Integration) implementations, such as the Warp computer [1], which is a linear systolic array, and the iWarp computer [2], which is a mesh-connected systolic array, designed in CMU.

The rest of this paper is organized as follows. In Section 2, we relate general nested-loop algorithms with uniform-dependence relations to the grid-connected systolic arrays. We also explain data collisions with an example. In Section 3, we present a complete set of necessary and sufficient conditions for a correct mapping under the constraint that at each time step, in each data link between two neighboring PEs there is at most one data token which can flow through. In Section 4, we generalize the conditions in Section 3 by allowing to shuffle traveling tokens. In Section 5, we compare our model with related research works. Finally, some concluding remarks and the directions of future works are given in Section 6.

2 Relating Nested-Loop Algorithms to Systolic Arrays

A p -nested-loop algorithm Ag is comprised of three parts: (1) the *loop index set*, $IP = \{(i_1, i_2, \dots, i_p) \mid l_j \leq i_j \leq u_j, 1 \leq j \leq p\}$, which is also called the *problem space*, where l_j and u_j are respectively the lower bound and the upper bound for each level j of the loop; (2) the set of *variables* V_{Ag} , including input variables, output variables, and temporary variables; and (3) the sequence of *statements* F_{Ag} in the loop body.

Given a p -nested-loop algorithm, we can find data-dependence vectors [20] [21] and classify them according to the *ZERO-ONE-INFINITE* property [10]. In general, data-dependence vectors can be defined in four ways: (1) modify-modify dependency, which is the relation that a variable is generated (modified) in different indices; (2) use-use dependency, which is the relation that a variable is used in different indices; (3) modify-use dependency, which is the relation that a variable is generated in one index and will be used in the other index; and (4) use-modify dependency, which is the relation that a variable is used in one index and will be generated in the other index.

Each different variable token symbol in the loop body causes one data-dependence vector of either modify-modify dependency or use-use dependency. If each variable token is either generated only once or used only once, the data-dependence vector is of type ZERO. If each variable token is generated or used periodically during the computation, the data-dependence vector is of type INFINITE. On the other hand, each pair of a modified token symbol (on the left-hand-side of $:=$) and a used token symbol (on the right-hand-side of $:=$); or each pair of a modified token symbol and the other modified token symbol; or each pair of a used token symbol and the other used token symbol of the same array variable, causes one "type

variable token	data-dependence vector	vector type	token type	dependency relation
$A[i, k]$	$d_1 = (0, 1, 0)^t$	INFINITE	input	use-use
$B[k, j]$	$d_2 = (1, 0, 0)^t$	INFINITE	input	use-use
$C[i, j]$	$d_3 = (0, 0, 1)^t$	INFINITE	output	modify-modify

Table 1: All three data-dependence vectors in the matrix multiplication algorithm are of type INFINITE.

ONE" data-dependence vector. This vector will be of either: modify-use dependency (flow dependence); or use-modify dependency (anti-dependence); or modify-modify dependency (output dependence); or use-use dependency. Tokens with data-dependence vectors of type ZERO or of type INFINITE are input or output variable tokens. Tokens with data-dependence vectors of type ONE are temporary variable tokens.

The ZERO-ONE-INFINITE classification can help understand whether tokens can be destroyed or not; whether tokens can be pipelined or not; and whether a PE needs additional I/O ports or not. This classification is the key to formulating the necessary and sufficient conditions for mapping nested-loop algorithms into the target systolic arrays.

Definition: The mathematical structure of a p -nested-loop algorithm A_g is a 5-tuple $M_{A_g} = (I^p, V_{A_g}, F_{A_g}, D_{A_g}, DT_{A_g})$, where:

1. I^p is the loop index set;
2. V_{A_g} is the set of variables;
3. F_{A_g} is the sequence of statements in the loop body;
4. D_{A_g} is the sequence of data-dependence vectors;
5. DT_{A_g} is the sequence of types for each data-dependence vector.

We will say that a nested-loop algorithm is shift-invariant if the dependence relations corresponding to all loop iterations in the index space are independent of their positions. We will also say that a nested-loop algorithm is of uniform-dependence, if all of the data-dependence vectors are constant. In this paper, we only consider shift-invariant uniform-dependence algorithms.

Example 1: The 3-nested-loop matrix multiplication algorithm

```

for  $i, j, k := 0$  to  $n$  do
   $C[i, j] := C[i, j] + A[i, k] * B[k, j];$ 

```

has three different variable token symbols, which cause three data-dependence vectors and all of the three data-dependence vectors are of type INFINITE as shown in Table 1. \square

Example 2: The 3-nested-loop algorithm

variable token	data-dependence vector	vector type	token type	dependency relation
$\langle A[i, j, k], A[i, j - 1, k] \rangle$	$d_1 = (0, 1, 0)^t$	ONE	temporary	modify-use
$\langle B[i, j, k], B[i - 1, j, k] \rangle$	$d_2 = (1, 0, 0)^t$	ONE	temporary	modify-use
$\langle C[i, j, k], C[i, j, k - 1] \rangle$	$d_3 = (0, 0, 1)^t$	ONE	temporary	modify-use
$A[i, j, k]$	$d_4 = (0, 0, 0)^t$	ZERO	output	modify-modify
$B[i, j, k]$	$d_5 = (0, 0, 0)^t$	ZERO	output	modify-modify
$C[i, j, k]$	$d_6 = (0, 0, 0)^t$	ZERO	output	modify-modify
$A[i, j - 1, k]$	$d_7 = (0, 0, 0)^t$	ZERO	input	use-use
$B[i - 1, j, k]$	$d_8 = (0, 0, 0)^t$	ZERO	input	use-use
$C[i, j, k - 1]$	$d_9 = (0, 0, 0)^t$	ZERO	input	use-use

Table 2: In Example 2, six different variable token symbols cause nine data-dependence vectors.

for $i, j, k := 0$ to n do

$A[i, j, k] := A[i, j - 1, k];$

$B[i, j, k] := B[i - 1, j, k];$

$C[i, j, k] := C[i, j, k - 1] + A[i, j, k] * B[i, j, k];$

has six different variable token symbols, which cause nine data-dependence vectors. Among them, three are of type ONE and six are of type ZERO as shown in Table 2.

Note that, $A[i, j, k]$, $B[i, j, k]$, and $C[i, j, k]$ are with data dependence vectors of type ZERO, and they are generated only once. We probably should call their dependency relations “modify”; however, for consistency with the dependency relations of data dependence vectors of type INFINITE and of type ONE, we call their dependency relations “modify-modify”. Similarly, $A[i, j - 1, k]$, $B[i - 1, j, k]$, and $C[i, j, k - 1]$ are with data dependence vectors of type ZERO, and they are used only once. Also, we call their dependency relations “use-use” instead of “use” for consistency with the dependency relations of data dependence vectors of type INFINITE and of type ONE. \square

We now briefly state our q -D grid-connected systolic array model. A PE in the q -D grid-connected systolic array is represented by the tuple $(pi_1, pi_2, \dots, pi_q)$, where $pl_j \leq pi_j \leq pu_j$ for $1 \leq j \leq q$, pl_j and pu_j are respectively the lower bound index and the upper bound index for each dimension j of the systolic array. In addition, it has only two physical data links connected to the neighboring PEs in each dimension. Therefore, there are in total $2q$ physical data links in each PE, or say, in each PE there are in total $2q$ interconnection primitives connected to $2q$ neighboring PEs. For convenience, we will say that P_i is the i th interconnection primitive, and $P_1 = [1, 0, \dots, 0]_{1 \times q}^t$; $P_2 = [-1, 0, \dots, 0]_{1 \times q}^t$; $P_{2i-1} = [0, \dots, 0, 1, 0, \dots, 0]_{1 \times q}^t$,

where the only non-zero entry in position i is 1; $P_{2i} = [0, \dots, 0, -1, 0, \dots, 0]_{1 \times q}^t$, where the only non-zero entry in position i is -1 ; $P_{2q-1} = [0, \dots, 0, 1]_{1 \times q}^t$; and $P_{2q} = [0, \dots, 0, -1]_{1 \times q}^t$.

We now explain how to relate nested-loop algorithms to systolic arrays. After finding and classifying all of the data-dependence vectors, we associate a single dedicated virtual data link with each data-dependence vector, which is with the corresponding set of data tokens. Note that, in the sequential algorithm there is only one copy of data tokens in the shared memory. However, in the systolic array algorithm, since dependent indices may be executed in different PEs, multiple copies of data tokens may be flowing through the systolic array. In general, if there are w data-dependence vectors with a specific array variable, then w copies of array variable tokens will be "traveling" through the systolic array, and each in a dedicated virtual data link. If a data-dependence vector has the condition $d_i = \vec{0}_p$, the corresponding virtual data link will be fixed in PEs, and therefore, the corresponding array variable tokens will be fixed in PEs. The purpose of classifying data-dependence vectors of type ZERO is to specify the requirements of local registers (or local memory) for those tokens with these zero data-dependence vectors.

A virtual data link may consist of several physical data links, therefore, a virtual data link may change its direction although each physical data link has only one direction. However, if a data token flows from $PE_{(p_1, p_2, \dots, p_i)}$ to $PE_{(p_1+j_1, p_2+j_2, \dots, p_i+j_i)}$, it must follow the following rule: it starts to flow along the first dimension j_1 PEs, then it flows along the second dimension j_2 PEs, and so on.

The number of virtual data links is usually more than the number of physical data links. In this case, the time-slicing method is used to assign different data token streams using different time slots flowing through the physical data links. As we will associate a single dedicated virtual data link with each data-dependence vector, we will say for convenience that data-dependence vector d_i will correspond to virtual data link i and data token stream i . A correct systolic array algorithm must manage the data token streams so that the data tokens' values are the same as those in the sequential algorithm.

To simplify the discussion, we assume throughout this paper that for each data-dependence vector $d_i = (d_{i1}, d_{i2}, \dots, d_{ip})^t$, $\gcd(d_{i1}, d_{i2}, \dots, d_{ip}) = 1$. The discussion for the general case when $\gcd(d_{i1}, d_{i2}, \dots, d_{ip}) \neq 1$, which requires additional virtual data links in order to avoid data collisions in the data links, is the same as that of [11].

The main effort of mapping a p -nested-loop algorithm into a q -D systolic array is to find a feasible space and time mapping, which maps an index iteration to be executed in a specific PE at a specific time. At the end of this section, we present a linear systolic array algorithm which Lee and Kedem claimed incurs

data collisions in the data link [10]. Let $\mathbf{H} \in Z^{1 \times p}$ be the time mapping vector and $\mathbf{S} \in Z^{1 \times p}$ be the space mapping vector for mapping p -nested loop algorithms into linear systolic arrays. Then the computation indexed by \bar{I} is executed at time $\mathbf{H}\bar{I}$ in $PE_{\mathbf{S}\bar{I}}$.

Example 3: Consider the 3-nested-loop matrix multiplication algorithm in Example 1. The linear systolic array implementation, where the time mapping $\mathbf{H} = (2, 1, 2)$ and the space mapping $\mathbf{S} = (1, 1, -2)$ for the problem size $n = 3$ will generate data collisions in data link 3 as mentioned in [10]. This is because $C[0, 3]$ collides with $C[2, 0]$ and $C[1, 3]$ collides with $C[3, 0]$ in data link 3 during the computation. For instance, in Fig. 1, there is only one shift register in data link 3. It is true that $C[0, 3]$ is first used and generated in index $(0, 3, 0)^t$ in PE_3 at time step 3, and $C[0, 3]$ will be used and generated next time in index $(0, 3, 1)^t$ in PE_1 at time step 5. Therefore, at time step 4, $C[0, 3]$ will be in PE_2 . However, at time step 4, $C[2, 0]$ also reaches PE_2 along data link 3. It is a data collision. \square

3 Synthesizing Grid-Connected Systolic array Algorithms under a Restricted Model

In this section, we show conditions for mapping nested loop algorithms into grid-connected systolic arrays with the constraint that at each time step, in each virtual data link between two neighboring PEs there is at most one data token which can flow through. Before introducing any conditions, we show an example of data collisions in a mesh-connected systolic array.

Example 4: The 3-nested-loop algorithm

$$\begin{aligned} & \text{for } i := 0 \text{ to } 15, j := 0 \text{ to } 15, k := 0 \text{ to } 13 \text{ do} \\ & \quad A[i, j, k] := F_1(A[i, j-4, k-3], C[-3j+2k, i]); \\ & \quad C[-3j+2k, i] := F_2(A[i-1, j, k-2], B[3i-j+k, 3i-j]); \end{aligned}$$

has five different variable token symbols, which cause seven data-dependence vectors as shown in Table 3. Note that F_1 and F_2 are two functions.

Let $\mathbf{H} = (1, 1, 1)$ and $\mathbf{S} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$. Then \mathbf{H} preserves data-dependence relations; in addition, \mathbf{H} and \mathbf{S} satisfy the non-singularity condition. However, as shown in Table 4 and Figure 2: tokens $A[0, 5, 5]$ (which is generated in index $(0, 5, 5)^t$ in $PE_{5,5}$ at time 10 and will be used in index $(0, 9, 8)^t$ in $PE_{9,8}$ at time 17), $A[0, 6, 5]$, $A[0, 7, 5]$, and $A[0, 8, 5]$, are with the data-dependence vector $d_1 = (0, 4, 3)^t$ and will pass through virtual data link 1 between $PE_{8,5}$ and $PE_{9,5}$ at time 13 simultaneously. Clearly, they cause

variable token	data-dependence vector	type	token type	dependency relation
$\langle A[i, j, k], A[i, j-4, k-3] \rangle$	$d_1 = (0, 4, 3)^t$	ONE	temporary	modify-use
$\langle A[i, j, k], A[i-1, j, k-2] \rangle$	$d_2 = (1, 0, 2)^t$	ONE	temporary	modify-use
$B[3i-j+k, 3i-j]$	$d_3 = (1, 3, 0)^t$	INFINITE	input	use-use
$C[-3j+2k, i]$	$d_4 = (0, 2, 3)^t$	INFINITE	output	modify-modify
$A[i, j, k]$	$d_5 = (0, 0, 0)^t$	ZERO	output	modify-modify
$A[i, j-4, k-3]$	$d_6 = (0, 0, 0)^t$	ZERO	input	use-use
$A[i-1, j, k-2]$	$d_7 = (0, 0, 0)^t$	ZERO	input	use-use

Table 3: In Example 4, there are seven data-dependence vectors.

data token	generated index	PE_{SI}	$H\bar{I}$	used index	PE_{SP}	$H\bar{I}'$
$A[0, 5, 5]$	$(0, 5, 5)^t$	$PE_{5,5}$	10	$(0, 9, 8)^t$	$PE_{9,8}$	17
$A[0, 6, 5]$	$(0, 6, 5)^t$	$PE_{6,5}$	11	$(0, 10, 8)^t$	$PE_{10,8}$	18
$A[0, 7, 5]$	$(0, 7, 5)^t$	$PE_{7,5}$	12	$(0, 11, 8)^t$	$PE_{11,8}$	19
$A[0, 8, 5]$	$(0, 8, 5)^t$	$PE_{8,5}$	13	$(0, 12, 8)^t$	$PE_{12,8}$	20

data token	time step										
	10	11	12	13	14	15	16	17	18	19	20
$A[0, 5, 5]$	$PE_{5,5}$	$PE_{6,5}$	$PE_{7,5}$	$PE_{8,5}$	$PE_{9,5}$	$PE_{9,6}$	$PE_{9,7}$	$PE_{9,8}$			
$A[0, 6, 5]$		$PE_{6,5}$	$PE_{7,5}$	$PE_{8,5}$	$PE_{9,5}$	$PE_{10,5}$	$PE_{10,6}$	$PE_{10,7}$	$PE_{10,8}$		
$A[0, 7, 5]$			$PE_{7,5}$	$PE_{8,5}$	$PE_{9,5}$	$PE_{10,5}$	$PE_{11,5}$	$PE_{11,6}$	$PE_{11,7}$	$PE_{11,8}$	
$A[0, 8, 5]$				$PE_{8,5}$	$PE_{9,5}$	$PE_{10,5}$	$PE_{11,5}$	$PE_{12,5}$	$PE_{12,6}$	$PE_{12,7}$	$PE_{12,8}$

Table 4: Tokens $A[0, 5, 5]$, $A[0, 6, 5]$, $A[0, 7, 5]$, and $A[0, 8, 5]$ are with the data-dependence vector $d_1 = (0, 4, 3)^t$ and will pass through virtual data link 1 between $PE_{8,5}$ and $PE_{9,5}$ at time 13 simultaneously. They cause data collisions in virtual data link 1.

data collisions in virtual data link 1. \square

In the following, we show conditions for a correct mapping. Let $\mathbf{H} \in Z^{1 \times p}$ be the time mapping vector and $\mathbf{S} \in Z^{q \times p}$ be the space mapping matrix, where $1 \leq q < p$. Then the computation indexed by \bar{I} is executed at time $\mathbf{H}\bar{I}$ in $PE_{\mathbf{S}\bar{I}}$. A correct q -D systolic array algorithm (\mathbf{H}, \mathbf{S}) that maps a p -nested-loop algorithm into a q -D grid-connected systolic array must preserve data-dependence relations, and the right tokens must be in the right place at the right time, and in addition, data tokens must not collide in the data links. We define $sign(\delta) = 1, -1, \text{ or } 0$, if $\delta > 0, \delta < 0, \text{ or } \delta = 0$, respectively.

Theorem 1 : *Suppose that at each time step, in each virtual data link between two neighboring PEs there is at most one data token which can flow through. Then, a q -D systolic array algorithm (\mathbf{H}, \mathbf{S}) maps a p -nested-loop algorithm correctly into a q -D grid-connected systolic array if and only if it satisfies the following four conditions.*

1. $\mathbf{H}d_i > 0$ for each non-zero data-dependence vector d_i .
2. $\forall \bar{I}_1, \bar{I}_2 \in I^p$, if $\bar{I}_1 \neq \bar{I}_2$, then either $\mathbf{H}\bar{I}_1 \neq \mathbf{H}\bar{I}_2$ or $\mathbf{S}\bar{I}_1 \neq \mathbf{S}\bar{I}_2$.
3. If $\mathbf{S}d_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \vec{0}_q$, then $\mathbf{H}d_i = b_i \sum_{j=1}^q |sd_{ij}|$, where b_i is a positive integer.
4. Three cases must be true to avoid data collisions in the data links.

Case 1: / for type ONE data-dependence vector d_i . */*

- If*
- (1) the data-dependence vector d_i is of type ONE,
 - (2) $\mathbf{S}d_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \vec{0}_q$,
 - (3) $\mathbf{H}(\bar{I}_2 - \bar{I}_1) > 0$, and
 - (4) $\mathbf{S}(\bar{I}_2 - \bar{I}_1) = \alpha sign(sd_{ij}) P_{2j-1} \neq \vec{0}_q$, where α is a positive integer and $0 < \alpha < |sd_{ij}|$,
- then* $\mathbf{H}(\bar{I}_2 - \bar{I}_1) \neq b_i \alpha$.

Case 2: / for type INFINITE data-dependence vector d_i and $\mathbf{S}d_i = sd_{ij} P_{2j-1} \neq \vec{0}_q$. */*

- If*
- (1) the data-dependence vector d_i is of type INFINITE,
 - (2) $\mathbf{S}d_i = sd_{ij} P_{2j-1} \neq \vec{0}_q$,
 - (3) $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ ,
 - (4) $\mathbf{H}(\bar{I}_2 - \bar{I}_1) > 0$, and
 - (5) $\mathbf{S}(\bar{I}_2 - \bar{I}_1) = \alpha sign(sd_{ij}) P_{2j-1} \neq \vec{0}_q$, where α is a positive integer,
- then* $\mathbf{H}(\bar{I}_2 - \bar{I}_1) \neq b_i \alpha$.

Case 3: / for type INFINITE data-dependence vector d_i and $\mathbf{S}d_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \vec{0}_q$. */*

- If
- (1) the data-dependence vector d_i is of type INFINITE,
 - (2) $Sd_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \bar{0}_q$,
 - (3) $Sd_i \neq sd_{ij}P_{2j-1}$, for every j ,
 - (4) $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ ,
 - (5) $H(\bar{I}_2 - \bar{I}_1) > 0$, and
 - (6) $S(\bar{I}_2 - \bar{I}_1) = \beta Sd_i + \alpha \text{sign}(sd_{ij})P_{2j-1} \neq \bar{0}_q$,
 where β is a non-negative integer and $-|sd_{ij}| < \alpha < |sd_{ij}|$,
- then $H(\bar{I}_2 - \bar{I}_1) \neq b_i(\beta \sum_{j=1}^q |sd_{ij}| + \alpha)$. \square

Note that, formally, the PEs in the q -D systolic array should be numbered from $(1, 1, \dots, 1)$ to $(pu_1, pu_2, \dots, pu_q)$, and S should map the indices to $\{PE_{(1,1,\dots,1)}, \dots, PE_{(pu_1, pu_2, \dots, pu_q)}\}$. However, it will be convenient to number the PE's from $\min\{S\bar{I}_1 | \bar{I}_1 \in I^p\}$ to $\max\{S\bar{I}_2 | \bar{I}_2 \in I^p\}$ in this presentation.

In the following, we describe the theorem instead of giving a formal proof. Condition 1 ensures that data-dependence relations are preserved according to the original sequential algorithm. Condition 2 insists that (H, S) is a non-singular mapping. That is, no two indices can be mapped to the same PE at the same time. Condition 3 specifies that tokens flow through data links at a constant speed. A benefit of this condition is that we can assign a constant amount of shift registers or data buffers for each of the virtual data links in each PE so that data streams can flow at a constant speed. Therefore, data congestion can be avoided. This condition also guarantees that data tokens will flow to the right place at the right time.

Suppose that a variable with the data-dependence vector d_i is generated in index \bar{I} and will be used next time in index $\bar{I} + d_i$. Index \bar{I} is executed in $PE_{S\bar{I}}$ at time $H\bar{I}$, and the index $\bar{I} + d_i$ will be executed in $PE_{S(\bar{I}+d_i)}$ at time $H(\bar{I} + d_i)$. Suppose that $Sd_i \neq \bar{0}_q$. Then the Manhattan distance between $PE_{S\bar{I}}$ and $PE_{S(\bar{I}+d_i)}$ is equal to $\sum_{j=1}^q |sd_{ij}|$, because $Sd_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t$. Therefore, each token in virtual data link i is delayed by $b_i = Hd_i / \sum_{j=1}^q |sd_{ij}|$ units of time in each PE, where b_i must be a positive integer. Thus, if there are r_i non-zero entries of sd_{ij} , where $1 \leq j \leq q$, we require in total $b_i r_i$ shift registers for virtual data link i in each PE.

The purpose of Condition 4 is to avoid data collisions in the data links. If $Sd_i = \bar{0}_q$, the virtual data link i is fixed in each PE and tokens with the data-dependence vector d_i are stored in local registers (or local memory). Therefore, in this case, tokens will not collide with each other in virtual data link i . If $Sd_i \neq \bar{0}_q$, there are three cases for data stream i , depending on the type of the data-dependence vector d_i and the value of Sd_i .

Case 1: d_i is of type ONE. The active life of each data token, which is with the data-dependence vector d_i , has only Hd_i units of time, because a token is generated only once and will be used only once. Therefore,

after that token is used, it can be destroyed. Suppose that a token is generated in index \bar{I}_1 at time $H\bar{I}_1$ in $PE_{S\bar{I}_1}$ and is used in index $\bar{I}_1 + d_i$ at time $H(\bar{I}_1 + d_i)$ in $PE_{S(\bar{I}_1 + d_i)}$. As $Sd_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t$, the token starts to flow along the first dimension sd_{i1} PEs in P_1 direction if $sd_{i1} > 0$, or $-sd_{i1}$ PEs in P_2 direction if $sd_{i1} < 0$; then, it flows along the second dimension sd_{i2} PEs in P_3 direction if $sd_{i2} > 0$, or $-sd_{i2}$ PEs in P_4 direction if $sd_{i2} < 0$; and so on. Data collisions occur only if there are other tokens using part of the path when the token flows through the path.

For instance, in Figure 2 of Example 4, token $A[0, 9, 5]$, which is with the data-dependence vector $d_1 = (0, 4, 3)^t$ of type ONE, is generated in index $(0, 9, 5)^t$ in $PE_{9,5}$, and will be used in index $(0, 13, 8)^t$ in $PE_{13,8}$. It may collide with tokens generated in $PE_{6,5}$, $PE_{7,5}$, $PE_{8,5}$, $PE_{10,5}$, $PE_{11,5}$, $PE_{12,5}$, $PE_{9,3}$, $PE_{9,4}$, $PE_{9,6}$, and $PE_{9,7}$ which are marked by blank circles. This is because all of these tokens will flow through part of the path from $PE_{9,5}$ to $PE_{13,8}$. Similarly, token $A[0, 15, 10]$, which is with the data-dependence vector $d_2 = (1, 0, 2)^t$ of type ONE, is generated in index $(0, 15, 10)^t$ in $PE_{15,10}$, and will be used in index $(1, 15, 12)^t$ in $PE_{15,12}$. It may collide with tokens generated in $PE_{15,9}$ and $PE_{15,11}$ which are marked by dotted circles.

The constraint in Case 1 of Condition 4 allows us to avoid data collisions. That is, if two distinct indices are executed in the same horizontal PEs, or in the same vertical PEs, or in the same perpendicular PEs within a certain distance, and the direction of these two PEs ($S(\bar{I}_2 - \bar{I}_1)$) is consistent with one of the basic directions $sign(sd_{ij})P_{2j-1}$ of the data path, then the difference of the execution time steps between these two indices cannot be equal to the time that a token flows from one PE to the other PE along the interconnection primitive $sign(sd_{ij})P_{2j-1}$ at speed $1/b_i$.

Case 2: d_i js of type INFINITE and $Sd_i = sd_{ij}P_{2j-1}$. If a token is with a data-dependence vector of type INFINITE, this token will be generated or used periodically during the computation. Therefore, throughout the computation, we cannot destroy this token. The data path in which tokens flow through has only one direction, $sign(sd_{ij})P_{2j-1}$. For instance, in Figure 2, $B[0, 0]$, which is with the data-dependence vector $d_3 = (1, 3, 0)^t$ of type INFINITE, is generated in index $(0, 0, 0)^t$ in $PE_{0,0}$, and will be used again in all indices $(\gamma, 3\gamma, 0)^t$ in $PE_{3\gamma,0}$ for integers $\gamma > 0$. It may collide with other tokens generated in $PE_{i,0}$ for all $i > 0$ which are marked by netted circles. Note that if $\bar{I}_2 - \bar{I}_1 = \gamma d_i$, where γ is an integer, the same token is generated or used, and therefore, it does not cause any data collision.

We conclude the condition in Case 2 (which is slightly different from the condition in Case 1): if two distinct indices are executed in the same horizontal PEs, or in the same vertical PEs, or in the same

perpendicular PEs without any constraint of the distance between these two execution PEs, and the direction of these two PEs is consistent with the direction $sign(sd_{ij})P_{2j-1}$ of the data path, then the difference of the execution time steps between these two indices cannot be equal to the time that a token flows from one PE to the other PE along the interconnection primitive $sign(sd_{ij})P_{2j-1}$ at speed $1/b_i$.

Case 3: d_i is of type INFINITE and $Sd_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq sd_{ij}P_{2j-1}$ for every j . Similar to Case 2, if a token is with such a data-dependence vector of type INFINITE, then this token cannot be destroyed during the computation. However, unlike Case 2, the data path in which the token flows through zigzags. For instance, in Figure 2, $C[8,0]$, which is with the data-dependence vector $d_4 = (0, 2, 3)^t$ of type INFINITE, is generated in index $(0, 0, 4)^t$ in $PE_{0,4}$, and will be used and generated again in indices $(0, 0, 4)^t + \beta d_i$ in $PE_{S((0,0,4)^t + \beta d_i)}$ for $\beta > 0$. The data path which token $C[8,0]$ flows through has a zigzag pattern. In addition, all tokens generated in the PEs marked by black circles will use part of the path, and they may cause data collisions.

The constraint in Case 3 of Condition 4 states that only the tokens, (which are generated or used in index \bar{I}_2 in $PE_{S\bar{I}_2}$ that are located horizontally, vertically, or perpendicularly within a certain distance to $PE_{S(\bar{I}_1 + \beta d_i)}$, for $\beta \geq 0$, and the direction of $PE_{S\bar{I}_2}$ and $PE_{S(\bar{I}_1 + \beta d_i)}$ is consistent with one of the basic directions $sign(sd_{ij})P_{2j-1}$ of the data path), may collide with the token generated or used in index \bar{I}_1 in $PE_{S\bar{I}_1}$ in virtual data link i . Therefore, the difference of the execution time steps between indices \bar{I}_2 and \bar{I}_1 cannot be equal to the time that a token flows from $PE_{S\bar{I}_1}$ to $PE_{S\bar{I}_2}$ along the composite data path " $\sum_{h=1}^q sd_{ih}P_{2h-1}$ " β times and then along the interconnection primitive P_{2j-1} or P_{2j} at speed $1/b_i$ depending on the location of $PE_{S\bar{I}_2}$.

Corollary 2 : If the systolic array algorithm (H, S) satisfies Condition 1 through Condition 3, and the data-dependence vector d_i is of type ONE, and $|sd_{ij}| \leq 1$ in Condition 3, then there is no data collision along the direction $sign(sd_{ij})P_{2j-1}$ in virtual data link i .

Proof : If $|sd_{ij}| \leq 1$ in Condition 3, then there does not exist any positive integer α such that $0 < \alpha < |sd_{ij}| \leq 1$ in Case 1 of Condition 4. Therefore, Condition 4 is satisfied automatically. \square

4 Synthesizing Grid-Connected Systolic Array Algorithms under a General Model

Condition 4 in the last section is over-restricted. It eliminates some solutions which allow more than one data token flowing through a specific data link and can be handled by employing a cyclic shuffle of traveling tokens. For instance, in Example 4, the mapping $\mathbf{H} = (1, 1, 1)$ and $\mathbf{S} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ violates Condition 4 of Theorem 1, this is because at each time step three tokens will reach each PE via virtual data link 1 from south and four tokens will reach each PE via virtual data link 1 from west (left). However, these seven tokens in effect are sent to seven distinct PEs. The first token from the southern neighboring PE will be used to perform the iteration loop body; the second and the third tokens from the southern neighboring PE will be sent to the northern neighboring PE; the fourth token from the western neighboring PE will be sent to the northern neighboring PE; and the remaining three tokens from the western neighboring PE will be sent to the eastern neighboring PE all via virtual data link 1. Therefore, a feasible sequence of microcodes in $PE_{(j,k)}$ at time step $i + j + k$ for handling tokens traveling through virtual data link 1 can be as follows.

1. Get a token $A[i, j - 4, k - 3]$ from the southern neighboring PE;
/* $A[i, j - 4, k - 3]$ is used to compute the iteration loop body. */
2. get a token $A[i, j - 4, k - 2]$ from the southern neighboring PE and then send to the northern neighboring PE;
3. get a token $A[i, j - 4, k - 1]$ from the southern neighboring PE and then send to the northern neighboring PE;
4. get a token $A[i, j - 4, k]$ from the western neighboring PE and then send to the northern neighboring PE;
5. get a token $A[i, j - 3, k]$ from the western neighboring PE and then send to the eastern neighboring PE;
6. get a token $A[i, j - 2, k]$ from the western neighboring PE and then send to the eastern neighboring PE;
7. get a token $A[i, j - 1, k]$ from the western neighboring PE and then send to the eastern neighboring PE;
8. compute the iteration loop body;
9. send a generated token $A[i, j, k]$ which is with data dependence vector d_1 to the eastern neighboring PE. \square

In general, for an arbitrary data stream i which is with data dependence vector d_i , the following algorithm can be used to handle the cyclic shuffle of traveling tokens in virtual data link i .

Algorithm for shuffling traveling tokens in virtual data link i :

Suppose that $\mathbf{S}d_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \vec{0}_q$. Then the generated microcodes in PE_{S_I} at time step $H\bar{i}$ for handling virtual data link i can be as follows.

Step 0: Starting from index $j = q$ of sd_{ij} , where $q \geq j \geq 1$, and using a decreasing order,

Step 1: /* obtain a token which is used to compute the iteration loop body */

search the first non-zero entry of sd_{ij} ;

get a token from $PE_{S\bar{I}-sign(sd_{ij})P_{2j-1}}$;

Step 2: /* transfer other tokens from $PE_{S\bar{I}-sign(sd_{ij})P_{2j-1}}$ to $PE_{S\bar{I}+sign(sd_{ij})P_{2j-1}}$ */

for $h = 2$ to $|sd_{ij}|$ do

get a token from $PE_{S\bar{I}-sign(sd_{ij})P_{2j-1}}$ and then send to $PE_{S\bar{I}+sign(sd_{ij})P_{2j-1}}$;

end_for

Step 3: /* transfer remaining tokens between neighboring PEs */

let $last = j$ and let $j = j - 1$;

while ($j \geq 1$) do

if ($sd_{ij} \neq 0$) then

get a token from $PE_{S\bar{I}-sign(sd_{ij})P_{2j-1}}$ and then send to $PE_{S\bar{I}+sign(sd_{i_{last}})P_{2last-1}}$;

for $h = 2$ to $|sd_{ij}|$ do

get a token from $PE_{S\bar{I}-sign(sd_{ij})P_{2j-1}}$ and then send to $PE_{S\bar{I}+sign(sd_{ij})P_{2j-1}}$;

end_for

let $last = j$;

end_if

let $j = j - 1$;

end_while

Step 4: after computing the iteration loop body, send a generated token with data-dependence vector d_i to $PE_{S\bar{I}+sign(sd_{i_{last}})P_{2last-1}}$. \square

Using the technique of shuffling traveling tokens, we can generalize Theorem 1 as the following theorem.

Theorem 3 : *A q -D systolic array algorithm (H, S) maps a p -nested-loop algorithm correctly into a q -D grid-connected systolic array if and only if it satisfies the following four conditions.*

1. $Hd_i > 0$ for each non-zero data-dependence vector d_i .
2. $\forall \bar{I}_1, \bar{I}_2 \in I^p$, if $\bar{I}_1 \neq \bar{I}_2$, then either $H\bar{I}_1 \neq H\bar{I}_2$ or $S\bar{I}_1 \neq S\bar{I}_2$.
3. If $Sd_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \vec{0}_q$, then $Hd_i = b_i \sum_{j=1}^q |sd_{ij}|$, where b_i is a positive integer.

4. /* for type INFINITE data-dependence vector d_i . */
 If (1) the data-dependence vector d_i is of type INFINITE,
 (2) $Sd_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \vec{0}_q$,
 (3) $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ ,
 (4) $H(\bar{I}_2 - \bar{I}_1) > 0$, and
 (5) $S(\bar{I}_2 - \bar{I}_1) = \beta Sd_i$, where β is a positive integer,
 then $H(\bar{I}_2 - \bar{I}_1) \neq b_i(\beta \sum_{j=1}^q |sd_{ij}|)$.

Proof: in Appendix. \square

Corollary 4 : *If the systolic array algorithm (H, S) satisfies Condition 1 through Condition 3, and the data-dependence vector d_i is of type ONE, and the algorithm for shuffling traveling tokens in virtual data link i is used, then there is no data collision along virtual data link i .*

Proof: Because the data-dependence vector d_i is of type ONE, Condition 4 in Theorem 3 is satisfied automatically. \square

A systolic array algorithm based on Theorem 3, which employs the technique of shuffling traveling tokens, is more general than a systolic array algorithm based on Theorem 1, which restricts at most one data token flowing from one PE to a neighboring PE in each virtual data link at each time step. It is easy to check that $H = (2, 1, 2)$ and $S = (1, 1, -2)$ in Example 3 and $H = (1, 1, 1)$ and $S = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ in Example 4 are infeasible systolic array algorithms under Theorem 1; however, they are feasible systolic array algorithms under Theorem 3.

Although systolic array algorithms based on Theorem 3 are more general than systolic array algorithms based on Theorem 1, the former systolic array algorithms require more shift registers than the latter ones. Suppose that $Sd_i^t = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t \neq \vec{0}_q$. A systolic array algorithm based on Theorem 3 requires $b_i(\sum_{j=1}^q |sd_{ij}|)$ shift registers in virtual data link i in each PE. However, a systolic array algorithm based on Theorem 1 only requires $b_i r_i$ shift registers in virtual data link i in each PE, where r_i is the number of non-zero entries of sd_{ij} .

Therefore, when evaluating whether (H, S) is a feasible systolic array algorithm, we first check whether data dependence vector d_i satisfies Theorem 1; if it does not, then we check whether d_i satisfies Theorem 3.

Algorithm for evaluating whether (H, S) is a feasible systolic array implementation :

Step 1: Check Conditions 1 through 3 in Theorem 1;

Step 2: for each data dependence vector d_i do Steps 3 through 5:

Step 3: if d_i satisfies Condition 4 of Theorem 1, then as at each time step there is at most one data token flowing from one PE to a neighboring PE via virtual data link i , thus, virtual data link i only requires $b_i r_i$ shift registers in each PE, where r_i is the number of non-zero entries of sd_{ij} ;

Step 4: else if d_i satisfies Condition 4 of Theorem 3, then because the technique of shuffling traveling tokens is used, virtual data link i requires $b_i(\sum_{j=1}^q |sd_{ij}|)$ shift registers in each PE;

Step 5: else data collisions occur in virtual data link i . \square

4.1 Other Properties for Evaluating Feasible Systolic Array Implementations

In general, the four necessary and sufficient conditions in Theorem 3 are exclusive. That is, none of them can be derived from other conditions. However, under certain constraints, either Condition 2 or Condition 4 would be redundant. In the following, we derive two theorems which allow us to reduce the computation complexity of evaluating feasible systolic array implementations for some special cases. First, for the special case when $q = p - 1$, Condition 4 can be derived from Condition 2.

Theorem 5 : Let $q = p - 1$. A q -D systolic array algorithm (\mathbf{H}, \mathbf{S}) from a p -nested-loop algorithm Ag that satisfies Conditions 1 through 3 maps Ag correctly into a q -D grid-connected systolic array.

Proof : We show that Condition 4 can be derived from Condition 1 through Condition 3. This is because from Corollary 4, there is no data collision along virtual data link i if the data dependence vector d_i is of type ONE. In the following we only consider the case when d_i is a type INFINITE data dependence vector. We want to show that if $\mathbf{S}d_i \neq \vec{0}_q$, and $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ , and $\mathbf{H}(\bar{I}_2 - \bar{I}_1) > 0$, and $\mathbf{S}(\bar{I}_2 - \bar{I}_1) = \beta \mathbf{S}d_i$ for some positive integer β , then $\mathbf{H}(\bar{I}_2 - \bar{I}_1) \neq b_i(\beta \sum_{j=1}^q |sd_{ij}|)$.

Assume by contradiction that $\mathbf{H}(\bar{I}_2 - \bar{I}_1) = b_i(\beta \sum_{j=1}^q |sd_{ij}|)$. Then, from Condition 3, because $\mathbf{H}d_i = b_i \sum_{j=1}^q |sd_{ij}|$, we have $\mathbf{H}(\bar{I}_2 - \bar{I}_1) = \beta \mathbf{H}d_i$. Next, from the assumption, $\mathbf{S}(\bar{I}_2 - \bar{I}_1) = \beta \mathbf{S}d_i$. Thus, we have $\mathbf{H}((\bar{I}_2 - \bar{I}_1) - \beta d_i) = 0$ and $\mathbf{S}((\bar{I}_2 - \bar{I}_1) - \beta d_i) = \vec{0}_{p-1}$. That is, $\begin{pmatrix} \mathbf{H} \\ \mathbf{S} \end{pmatrix} ((\bar{I}_2 - \bar{I}_1) - \beta d_i) = \vec{0}_p$.

However, from Condition 2, \mathbf{H} and \mathbf{S} are non-singular, and in addition, \mathbf{H} is of rank 1 and \mathbf{S} is of rank $p - 1$ (because $q = p - 1$). Therefore, $\begin{pmatrix} \mathbf{H} \\ \mathbf{S} \end{pmatrix}$ is a basis. As a basis will not map a non-zero vector to a zero vector, $(\bar{I}_2 - \bar{I}_1) - \beta d_i = \vec{0}_p$ must be true. Therefore, $\bar{I}_2 - \bar{I}_1$ is equal to βd_i . From the assumption that $\gcd(d_{i1}, d_{i2}, \dots, d_{ip}) = 1$, where $d_i = (d_{i1}, d_{i2}, \dots, d_{ip})^t$, it follows that β is an integer. However, this contradicts our assumption that $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ . \square

Second, if there exists a type INFINITE data dependence vector d_i such that $Sd_i \neq \vec{0}_q$, then Condition 2 is implicitly satisfied by Condition 1, Condition 3, and Condition 4.

Lemma 6 : *Suppose that (\mathbf{H}, \mathbf{S}) is a feasible systolic array algorithm which satisfies Theorem 3. Then, for any type INFINITE data dependence vector d_i , if $Sd_i \neq \vec{0}_q$, and $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ , and $S(\bar{I}_2 - \bar{I}_1) = \beta Sd_i$ for some positive integer β , we have $\mathbf{H}(\bar{I}_2 - \bar{I}_1)Sd_i \neq \mathbf{H}d_i S(\bar{I}_2 - \bar{I}_1)$.*

Proof : Assume by contradiction that $\mathbf{H}(\bar{I}_2 - \bar{I}_1)Sd_i = \mathbf{H}d_i S(\bar{I}_2 - \bar{I}_1)$. $\mathbf{H}(\bar{I}_2 - \bar{I}_1)$ and $\mathbf{H}d_i$ are integers, $S(\bar{I}_2 - \bar{I}_1)$ and Sd_i are q -tuple vectors. Because $S(\bar{I}_2 - \bar{I}_1) = \beta Sd_i$, $\mathbf{H}(\bar{I}_2 - \bar{I}_1) = \beta \mathbf{H}d_i = \beta(b_i \sum_{j=1}^q |sd_{ij}|)$ by Condition 3. However, this contradicts Condition 4 that $\mathbf{H}(\bar{I}_2 - \bar{I}_1) \neq b_i(\beta \sum_{j=1}^q |sd_{ij}|)$. \square

Theorem 7 : *Suppose that (\mathbf{H}, \mathbf{S}) is a systolic array algorithm. If there exists at least one type INFINITE data-dependence vector d_i such that $Sd_i \neq \vec{0}_q$, then Condition 2 is implicitly satisfied by Condition 1, Condition 3, and Condition 4 in Theorem 3.*

Proof : We want to show that for any two distinct indexes \bar{I}_1 and \bar{I}_2 either $\mathbf{H}\bar{I}_1 \neq \mathbf{H}\bar{I}_2$ or $\mathbf{S}\bar{I}_1 \neq \mathbf{S}\bar{I}_2$. Consider two cases:

1. $\mathbf{H}\bar{I}_1 = \mathbf{H}\bar{I}_2$.

If $(\bar{I}_2 - \bar{I}_1) = \gamma d_i$ for some integer γ , then $\gamma \mathbf{H}d_i = \mathbf{H}\bar{I}_2 - \mathbf{H}\bar{I}_1 = 0$ which contradicts Condition 1.

If $(\bar{I}_2 - \bar{I}_1) \neq \gamma d_i$ for all integers γ , then because $Sd_i \neq \vec{0}_q$ and $\mathbf{H}d_i \neq 0$, from Lemma 6, $\vec{0}_q = \mathbf{H}(\bar{I}_2 - \bar{I}_1)Sd_i \neq \mathbf{H}d_i S(\bar{I}_2 - \bar{I}_1)$ which implies $\mathbf{S}\bar{I}_1 \neq \mathbf{S}\bar{I}_2$.

2. $\mathbf{S}\bar{I}_1 = \mathbf{S}\bar{I}_2$.

If $(\bar{I}_2 - \bar{I}_1) = \gamma d_i$ for some integer γ , then $\mathbf{H}\bar{I}_2 - \mathbf{H}\bar{I}_1 = \gamma \mathbf{H}d_i \neq 0$ by Condition 1.

If $(\bar{I}_2 - \bar{I}_1) \neq \gamma d_i$ for all integers γ , then because $Sd_i \neq \vec{0}_q$ and $\mathbf{H}d_i \neq 0$, from Lemma 6, $\mathbf{H}(\bar{I}_2 - \bar{I}_1)Sd_i \neq \mathbf{H}d_i S(\bar{I}_2 - \bar{I}_1) = \vec{0}_q$ which implies $\mathbf{H}\bar{I}_1 \neq \mathbf{H}\bar{I}_2$. \square

5 Comparisons with Related Works

As mentioned in Section 1, the problem of finding feasible or optimal mappings for systolic array implementations has a rich history with a large number of transformations proposed by various researchers. However, most of them only considered the case when mapping p -nested-loop algorithms into $(p-1)$ -dimensional systolic arrays, in addition, they did not consider the cases when data might collide in the data links [4] [7]

[13]–[17] [19] [23]. In effect, as shown in Theorem 5, when $q = p - 1$, if the non-singularity condition holds, the condition for avoiding data collision also holds when the algorithm for shuffling traveling tokens is used. However, the condition for avoiding data collision is still necessary, if we restrict at most one data token flowing from one PE to a neighboring PE in each virtual data link at each time step, such as the case in Example 4.

In the following, we compare our method with three related works published in the relevant literature. Lee and Kedem first proposed a set of necessary and sufficient conditions for mapping nested-loop algorithms into linear systolic arrays [10]. They then generalized their method for mapping nested-loop algorithms into systolic arrays of arbitrary dimensions [11]. However, their model restricts each virtual data link to have only one direction. Therefore, the systolic arrays they design may require more physical data links than ours. Certainly, their design may increase the complexity of the hardware and the fabrication cost of the target systolic arrays. Next, Lee and Kedem's model only allows at most one data token flowing from one PE to a neighboring PE in each virtual data link at each time step. Therefore, some interesting systolic array algorithms which can be handled by shuffling traveling tokens are eliminated by their constraints, such as the systolic array algorithm in Example 3.

Recently, Shang and Fortes [22] presented conditions for the space-time mapping. However, under our systolic array model, their methods are not sufficient to avoid data collisions in the data links. Shang and Fortes proposed conditions for mapping p -nested-loop algorithms into q -D systolic arrays, where $1 \leq q < p - 1$. However, as we will show in Example 5, the linear array implementation for the matrix multiplication they derived will generate data collisions in the data links.

Example 5: Consider the 3-nested-loop matrix multiplication algorithm in Example 1 in Section 2. The linear systolic array implementation, where the time mapping $H = (1, 2, 2)$ and the space mapping $S = (1, 1, -1)$ for the problem size $n = 3$ are derived by [22], will generate data collisions in data link 2. This is because H and S do not satisfy Condition 4 in either Theorem 1 or Theorem 3.

In Figure 3, $B[k, j]$, which are with the data-dependence vector $d_2 = (1, 0, 0)^t$ of type INFINITE, are input variables for $0 \leq j, k \leq 3$, and therefore, $B[k, j]$ cannot be destroyed during the computation. We now examine the behavior of two tokens $B[0, 3]$ and $B[1, 0]$. $B[0, 3]$ is first used in index $\bar{I}_2 = (0, 3, 0)^t$ in PE_3 at time step 6; $B[1, 0]$ is first used in index $\bar{I}_1 = (0, 0, 1)^t$ in PE_{-1} at time step 2. It is true that $\bar{I}_2 - \bar{I}_1 = (0, 3, -1)^t \neq \gamma d_2$ for all integers γ , $Sd_2 = 1$, $b_2 = 1$, and $S(\bar{I}_2 - \bar{I}_1) = 4 = \alpha$; however, $H(\bar{I}_2 - \bar{I}_1) = 4 = b_2\alpha$ violates the constraint at Case 2 of Condition 4 in Theorem 1, or the constraint of

Condition 4 in Theorem 3. Therefore, $B[0, 3]$ must collide with $B[1, 0]$ in virtual data link 2.

This is because $B[0, 3]$ is an input token, and $B[0, 3]$ enters into the linear array from PE_{-3} . As $Sd_2 = 1$ and $b_2 = Hd_2/Sd_2 = 1$, there is only one shift register in virtual data link 2 in each PE. In addition, $B[0, 3]$ is delayed by only one unit of time in each PE. Therefore, if $B[0, 3]$ is used in PE_3 at time step 6, it is in PE_2 at time step 5, in PE_1 at time step 4, in PE_0 at time step 3, and in PE_{-1} at time step 2. However, at these times, $B[1, 0]$ also passes through virtual data link 2, and therefore, $B[0, 3]$ collides with $B[1, 0]$. Similarly, $B[1, 3]$ collides with $B[2, 0]$, and $B[2, 3]$ collides with $B[3, 0]$ in virtual data link 2. \square

Example 6: Consider the 3-nested-loop algorithm in Example 2 in Section 2. Then, the linear array implementation $H = (1, 2, 2)$ and $S = (1, 1, -1)$ for the problem size $n = 3$ is feasible. This is because H and S satisfy Conditions 1 through 3 of Theorem 1; all of the non-zero data-dependence vectors d_1 , d_2 , and d_3 are of type ONE; and in addition, $Sd_1 = 1$, $Sd_2 = 1$, and $Sd_3 = -1$. Thus, from Corollary 2 or from Corollary 4, there is no data collision in virtual data links 1, 2, and 3. \square

We have shown two examples in this section which illustrate that our necessary and sufficient conditions are superior to the ones proposed by Shang and Fortes [22]. In effect, the semantic meaning of the 3-nested-loop matrix multiplication algorithm in Example 1 is different from that of the 3-nested-loop algorithm in Example 2. In the former case, all data dependence vectors are of type INFINITE. Therefore, if $Sd_i \neq \bar{0}_q$, input data tokens can be pipelined entering into the systolic arrays, output data tokens can be pipelined extracted from the systolic arrays, the I/O time can be overlaid by the computation time, and in addition, each processing element can contain only a constant number of shift registers for data link i . On the other hand, in the latter case, all data dependence vectors are of type ZERO and of type ONE. Therefore, input data tokens with data dependence vectors of type ZERO must be loaded into the systolic arrays before the computation, and output data tokens with data dependence vectors of type ZERO must be unloaded from the systolic array after the computation. The I/O time and the computation time cannot be overlaid; moreover, each processing element must contain a sufficiently large local memory for storing data tokens with data dependence vectors of type ZERO.

The conditions proposed by Shang and Fortes can only deal with algorithms whose data-dependence vectors are all of type ONE or of type ZERO. Therefore, their method, in general, is not sufficient to deal with algorithms which include data-dependence vectors of type INFINITE, such as the matrix-multiplication algorithm.

Ganapathy and Wah [6], who based their work on Li and Wah's method [12], proposed the parameter-




Figure 4: The relations among variant solution space.

based approach to map p -dimensional uniform recurrences to any q -dimensional processor arrays, where $q < p$. As mentioned in Section 1, in their model, tokens flowing through systolic arrays cannot change directions as the case in [11]. In addition, their model does not classify the tokens' behavior. And finally, their method can only deal with algorithms whose data dependence vectors are all of type INFINITE; they do not consider algorithms whose data dependence vectors are all of type ZERO and of type ONE such as the 3-nested-loop algorithm in Example 2. However, they provided certain conditions for dealing with shift-varying algorithms such as the reindexed transitive closure algorithm [8].

In Fig. 4, we show the relations among variant solution space based on published results by Lee and Kedem [10] [11], Shang and Fortes [22], Ganapathy and Wah [6], and our two models: Theorem 1 and Theorem 3. Note that, although the solution space by Shang and Fortes covers all other solution space, their feasible systolic array implementations may incur data collisions in data links. Unlike them, feasible systolic array implementations in all other solution space are collision-free.

6 Conclusions and Future Works

We have presented in this paper two sets of necessary and sufficient conditions for mapping nested-loop algorithms with shift-invariant uniform-dependence relations into grid-connected systolic arrays of arbitrary dimensions without data collisions in the data links. The first set of conditions is under the constraint that at each time step, in each virtual data link between two neighboring PEs there is at most one data token which can pass through. The second set of conditions is under a general model in which a cyclic shuffling algorithm is used to handle traveling tokens, and therefore, more than one data token is allowed to flow from one PE to a neighboring PE at each time step. These two sets of conditions, which are based on the ZERO-ONE-INFINITE classification of tokens' behavior, can be used to generate all feasible grid-connected

systolic array implementations. In addition, when the size of the problem to be solved grows and the size of the provided systolic array is less than the size which a space mapping needs, then the partition method due to Moldovan and Fortes can be used to handle this case [15]. However, although systolic array algorithms based on the second set of conditions are more general than systolic array algorithms based on the first set of conditions, the formal systolic array algorithms require more shift registers than the latter ones.

The future research works will include two directions. First, it would be interesting to find optimal systolic array implementations when provided some performance criteria. Although an ad-hoc method also can find optimal systolic array implementations, it may use a lot of computation time. Therefore, fast algorithms or heuristic algorithms for finding optimal systolic array implementations for a class of nested loop algorithms may be of interest. Second, it would be interesting to extend our method to deal with shift-varying uniform-dependence algorithms. Although our conditions are sufficient for mapping shift-varying uniform-dependence algorithms into systolic arrays, the conditions are over-restricted, because some loop iterations really do not have certain dependence relations. We may consider the specific indices where dependence relations are variant. However, as different shift-varying algorithms may contain different sets of indices whose dependence relations are variant, therefore, the problem of how to deal with general shift-varying algorithms is still an open question.

Appendix

Theorem 8 : *A q -D systolic array algorithm (H, S) that maps correctly a p -nested-loop algorithm into a q -D grid-connected systolic array must satisfy Conditions 1 through 4.*

Proof: Condition 1 follows from Lamport's result [9]. Condition 2 is related to the non-singularity condition of Lee and Kedem [11], Moldovan and Fortes [15], and others. We now show that Condition 3 is necessary. From the restriction of systolic arrays, tokens must flow through data links at a constant speed. Suppose that a variable with the data-dependence vector d_i is generated in index \bar{I} and will be used next time in index $\bar{I} + d_i$. Index \bar{I} is executed in PE_{S_I} at time $H\bar{I}$, and the index $\bar{I} + d_i$ will be executed in $PE_{S(\bar{I}+d_i)}$ at time $H(\bar{I} + d_i)$. Suppose that $Sd_i \neq \bar{0}_q$. Then the Manhattan distance from PE_{S_I} to $PE_{S(\bar{I}+d_i)}$ is equal to $\sum_{j=1}^q |sd_{ij}|$, because $Sd_i = (sd_{i1}, sd_{i2}, \dots, sd_{iq})^t$. Therefore, each token in virtual data link i is delayed by $b_i = Hd_i / \sum_{j=1}^q |sd_{ij}|$ units of time in each PE, where b_i must be a positive integer.

We now show that Condition 4 is necessary. Let x_1 and x_2 be two tokens of the variable X corresponding to d_i that is of type INFINITE, such that x_1 is the token generated in \bar{I}_1 and x_2 is the token generated

in \bar{I}_2 . Then, x_1 is in $PE_{S_{\bar{I}_1}}$ at time $H\bar{I}_1$, x_2 is in $PE_{S_{\bar{I}_2}}$ at time $H\bar{I}_2$, and they cannot be destroyed during the computation. Since $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ , from Lemma 7 in [10], $x_1 \neq x_2$.¹ Since $S(\bar{I}_2 - \bar{I}_1) = \beta S d_i \neq \bar{0}_q$, $\beta \sum_{j=1}^q |sd_{ij}|$ is the Manhattan distance from $PE_{S_{\bar{I}_1}}$ to $PE_{S_{\bar{I}_2}}$ along the composite interconnection primitives " $\sum_{j=1}^q sd_{ij} P_{2j-1}$ " β times.

Since $H(\bar{I}_2 - \bar{I}_1) > 0$, $H\bar{I}_1 < H\bar{I}_2$. Let c be the index such that x_1 is in PE_c at time $H\bar{I}_2$. Then, from Condition 3, and the assumption $S d_i = \sum_{j=1}^q sd_{ij} P_{2j-1} \neq \bar{0}_q$, $H(\bar{I}_2 - \bar{I}_1) = b_i \times$ (the Manhattan distance from $PE_{S_{\bar{I}_1}}$ to PE_c along the composite interconnection primitives $\sum_{j=1}^q sd_{ij} P_{2j-1}$).

Assume by contradiction that $H(\bar{I}_2 - \bar{I}_1) = b_i(\beta \sum_{j=1}^q |sd_{ij}|)$. Then, the Manhattan distance from $PE_{S_{\bar{I}_1}}$ to PE_c along the composite interconnection primitives $\sum_{j=1}^q sd_{ij} P_{2j-1}$ is equal to $\beta \sum_{j=1}^q |sd_{ij}|$ (the Manhattan distance from $PE_{S_{\bar{I}_1}}$ to $PE_{S_{\bar{I}_2}}$ along the composite interconnection primitives $\sum_{j=1}^q sd_{ij} P_{2j-1}$). We have $c = S\bar{I}_2$. Thus, the tokens x_1 and x_2 both use virtual data link i , and appear in $PE_{S_{\bar{I}_2}}$ at time $H\bar{I}_2$. And in addition, because β is a positive integer, even if we apply the algorithm for shuffling traveling tokens in virtual data link i , tokens x_1 and x_2 both must be handled by the first step of the algorithm. This is a "data collision", which is not allowed. Therefore, $H(\bar{I}_2 - \bar{I}_1) \neq b_i(\beta \sum_{j=1}^q |sd_{ij}|)$. \square

Theorem 9 : *A q -D systolic array algorithm (H, S) from a p -nested-loop algorithm Ag that satisfies Conditions 1 through 4 maps Ag correctly into a q -D grid-connected systolic array.*

Proof: First, from Condition 1, H preserves the data-dependence ordering.

Second, we will show that the "needed" tokens will flow to the right place at the right time. Since all the statements in F_{Ag} can be handled in the same way, in the following we only consider a single representative statement. Such statement can be rewritten so that it has m variables and each variable corresponds to a data-dependence vector, this is because the p -nested-loop algorithm has m data-dependence vectors. Let the single statement be $\hat{F}_{Ag}(v_1, v_2, \dots, v_m)$, where \hat{F}_{Ag} is a statement in F_{Ag} and v_i , which is with the data-dependence vector d_i and was generated in \bar{I}_i ($\bar{I}_i \in I^p$), is an entry of variable V_i . For simplicity, we assume that only one time unit is needed to execute the loop body. We now show by induction on $H\bar{I}$ that all tokens v_i arriving at $PE_{S_{\bar{I}}}$ at time $H\bar{I}$ will have correct values.

Basic step: $H\bar{I} = \min\{H\bar{I}_1 | \bar{I}_1 \in I^p\}$. All tokens v_i are initial input/output tokens (variables). From Condition 3 and the assumption that every variable enters into the systolic array at a correct time, v_i , which flows through virtual data link i , will reach $PE_{S_{\bar{I}}}$ at time $H\bar{I}$. Thus, the systolic array algorithm

¹Lemma 7 in [10]: Let x_1 and x_2 be tokens of variable X in the data stream corresponding to d_i that is of type INFINITE. Let x_1 be used in some index \bar{I}_1 and let x_2 be used in some index \bar{I}_2 . Then, $x_1 \neq x_2$ if and only if $\bar{I}_2 - \bar{I}_1$ is not an integer multiple of d_i .

performs the same function as \hat{F}_{Ag} in \bar{I} .

Induction step: $H\bar{I} > \min\{H\bar{I}_1 | \bar{I}_1 \in I^P\}$. From the definition of the data-dependence vectors we have

$$\bar{I} = \bar{I}_1 + d_1 = \bar{I}_2 + d_2 = \dots = \bar{I}_m + d_m, \quad (1)$$

for appropriate $\bar{I}_1, \bar{I}_2, \dots, \bar{I}_m$. By induction, v_i was regenerated with the correct value in $PE_{S\bar{I}_1}$ at time $H\bar{I}_1$. However, v_i was not used (and was not modified) during the time instances $H\bar{I}_1 + 1, H\bar{I}_1 + 2, \dots, H(\bar{I}_1 + d_1) - 1$. Then from Condition 3, v_i , which flows through virtual data link i , will reach $PE_{S\bar{I}_1 + Sd_1}$ at time $H\bar{I}_1 + Hd_1$. But from (1), $S\bar{I} = S(\bar{I}_1 + d_1)$ and $H\bar{I} = H(\bar{I}_1 + d_1)$. Therefore, all of the tokens v_1, v_2, \dots, v_m arriving at $PE_{S\bar{I}}$ at time $H\bar{I}$ will have correct values. The systolic array algorithm will thus perform the same function as \hat{F}_{Ag} in \bar{I} .

Third, we will show that no two tokens (variables) collide in any data link. If $Sd_i = \bar{0}_g$, the virtual data link i is fixed in each PE and tokens with the data-dependence vector d_i are stored in local registers (or local memory). Therefore, in this case, tokens will not collide with each other in virtual data link i . In the following we consider the case when $Sd_i \neq \bar{0}_g$. Assume by contradiction that virtual data link i has two distinct tokens x_1 and x_2 of the variable X that collide in PE_a during the execution. (Thus, of course, d_i is with X .) Suppose that x_1 was just generated in index \bar{I}_1 in $PE_{S\bar{I}_1}$ at time $H\bar{I}_1$ and x_2 was just generated in index \bar{I}_2 in $PE_{S\bar{I}_2}$ at time $H\bar{I}_2$. (If x_1 or x_2 was not generated before, then x_1 is then immediately used in index \bar{I}_1 or x_2 is then immediately used in index \bar{I}_2 .) Consider two cases:

1. $H\bar{I}_1 = H\bar{I}_2$.

Our assumption is that at time instance $H\bar{I}_1 = H\bar{I}_2$ x_1 is in $PE_{S\bar{I}_1}$ and x_2 is in $PE_{S\bar{I}_2}$, and then at some time instance x_1 and x_2 collide in PE_a . However, since all tokens of data stream i flow through virtual data link i at the same speed, we have $PE_{S\bar{I}_1} = PE_{S\bar{I}_2}$ which contradicts Condition 2.

2. $H\bar{I}_1 \neq H\bar{I}_2$.

Without loss of generality, $H\bar{I}_1 < H\bar{I}_2$. Since x_1 and x_2 flow at the same speed, x_1 flowing from $PE_{S\bar{I}_1}$ at time $H\bar{I}_1$ will reach PE_a at some time instance t and x_2 flowing from $PE_{S\bar{I}_2}$ at time $H\bar{I}_2$ will reach PE_a at the time instance t . Therefore, x_1 will reach $PE_{S\bar{I}_2}$ at time $H\bar{I}_2$. As all tokens flow with non-zero velocity, $S\bar{I}_1 \neq S\bar{I}_2$. Consider two subcases:

2(a). $\bar{I}_2 - \bar{I}_1 = \gamma d_i$ for some integer γ .

Consider two cases:

- i. d_i is a type ONE data-dependence vector.

Since x_1 is generated in \bar{I}_1 , x_1 is only used in $\bar{I}_1 + d_i$. However, since x_1 is generated in $PE_{S_{\bar{I}_1}}$ and $\bar{I}_2 \neq \bar{I}_1$ and $\bar{I}_2 - \bar{I}_1 = \gamma d_i$ for some integer γ , we have $\bar{I}_2 = \bar{I}_1 + d_i$ and $PE_a = PE_{S_{\bar{I}_2}}$. From the definition of type ONE data-dependence vector, x_1 will not be used again after it is used in \bar{I}_2 . Therefore, when x_2 is generated in $PE_{S_{(\bar{I}_1+d_i)}} = PE_{S_{\bar{I}_2}}$, x_1 is destroyed. Thus, no collisions will occur.

ii. d_i is a type INFINITE data-dependence vector.

From Lemma 7 in [10], $x_1 = x_2$. However, this contradicts our assumption that $x_1 \neq x_2$.

2(b). $\bar{I}_2 - \bar{I}_1 \neq \gamma d_i$ for all integers γ .

Consider two cases:

i. $S(\bar{I}_2 - \bar{I}_1) = \beta S d_i$ for some integer β .

Since $S\bar{I}_1 \neq S\bar{I}_2$ and $S d_i \neq \bar{0}_q$, β is non-zero. Suppose that β is a negative integer. Because our assumption is that $H\bar{I}_1 < H\bar{I}_2$, thus at time instance $H\bar{I}_2$, the Manhattan distance between tokens x_1 and x_2 is at least $\beta \sum_{j=1}^q |s_{d_{ij}}|$. However, this contradicts our derivation result that x_1 will reach $PE_{S_{\bar{I}_2}}$ at time $H\bar{I}_2$.

Suppose that β is a positive integer. Consider two subcases:

i(a). d_i is a type ONE data-dependence vector.

Since d_i is a type ONE data-dependence vector, token x_1 , which is generated in $PE_{S_{\bar{I}_1}}$, can be destroyed after it is used in $PE_{S_{(\bar{I}_1+d_i)}}$. On the other hand, token x_2 is generated in $PE_{S_{(\bar{I}_1+\beta d_i)}}$ and is used in $PE_{S_{(\bar{I}_1+(\beta+1)d_i)}}$, where $\beta \geq 1$. Therefore, tokens x_1 and x_2 , which flow through different PEs, will not collide with each other in virtual data link i .

i(b). d_i is a type INFINITE data-dependence vector.

Because d_i is a type INFINITE data-dependence vector, tokens x_1 and x_2 cannot be destroyed during the execution. Since x_1 is in $PE_{S_{\bar{I}_1}}$ at time $H\bar{I}_1$ and is in $PE_{S_{\bar{I}_2}}$ at time $H\bar{I}_2$, and therefore, $H(\bar{I}_2 - \bar{I}_1) = b_i(\beta \sum_{j=1}^q |s_{d_{ij}}|)$. However, this contradicts Condition 4.

ii. $S(\bar{I}_2 - \bar{I}_1) \neq \beta S d_i$ for all integers β .

Although both x_1 and x_2 may be in the same PE at the same time, we can apply the algorithm for shuffling traveling tokens. Then, tokens x_1 and x_2 are handled at different steps of the algorithm. Thus, tokens x_1 and x_2 will not collide with each other in virtual data link i .

We conclude that (H, S) can be performed in parallel. \square

References

- [1] M. A. Annaratone, E. Arnould, T. Gross, H. T. Kung, M. Lam, O. Menzilcioglu, and J. A. Webb. The Warp computer: Architecture, implementation, and performance. *IEEE Transactions on Computers*, C-36:1523-1538, December 1987.
- [2] S. Borkar, R. Cohn, G. Cox, S. Gleason, T. Gross, H. T. Kung, et al. iwarp: An integrated solution to high-speed parallel computing. In *Supercomputing '88*, pages 330-339, Orlando, Florida, November 1988. IEEE/ACM.
- [3] P. R. Cappello and K. Steiglitz. *Advances in Computing Research*, volume 2: VLSI Theory, chapter Unifying VLSI Array Design with Linear Transformations of Space-Time, pages 23-65. JAI Press, Greenwich, CT, 1984.
- [4] M. C. Chen. The generation of a class of multipliers: Synthesizing highly parallel algorithms in VLSI. *IEEE Transactions on Computers*, C-37:329-338, March 1988.
- [5] J. A. B. Fortes, K. S. Fu, and B. W. Wah. Systematic approaches to the design of algorithmic specified systolic array. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, pages 300-303, Tampa, FL, March 1985. IEEE.
- [6] K. N. Ganapathy and B. W. Wah. Synthesizing optimal lower dimensional processor arrays. In *Proceedings of the International Conference on Parallel Processing*, pages III-96-103, August 1992.
- [7] S. Y. Kung. On supercomputing with systolic/wavefront array processors. *Proceedings of the IEEE*, 72:867-884, July 1984.
- [8] S. Y. Kung, S. C. Lo, and P. S. Lewis. Optimal systolic design for the transitive closure and the shortest path problems. *IEEE Transactions on Computers*, C-36:603-614, May 1987.
- [9] L. Lamport. The parallel execution of Do loops. *Communications of the ACM*, pages 83-93, Feb 1974.
- [10] P.-Z. Lee and Z. M. Kedem. Synthesizing linear-array algorithms from nested For loop algorithms. *IEEE Transactions on Computers*, C-37:1578-1598, December 1988.
- [11] P.-Z. Lee and Z. M. Kedem. Mapping nested loop algorithms into multi-dimensional systolic arrays. *IEEE Transactions on Parallel and Distributed Systems*, 1:64-76, January 1990.
- [12] G. Li and B. W. Wah. The design of optimal systolic arrays. *IEEE Transactions on Computers*, pages 66-77, January 1985.
- [13] W. L. Miranker and A. Winkler. Spacetime representations of computational structures. *Computing*, 32:93-114, 1984.
- [14] D. I. Moldovan. ADVIS: A software package for the design of systolic arrays. *IEEE Transactions on CAD*, CAD-6:33-40, January 1987.
- [15] D. I. Moldovan and J. A. B. Fortes. Partitioning and mapping algorithms into fixed size systolic arrays. *IEEE Transactions on Computers*, C-35:1-12, January 1986.
- [16] P. Quinton. Automatic synthesis of systolic arrays from uniform recurrence equations. In *11th Annu. Symp. Comput. Architecture*, pages 208-214, 1984.
- [17] P. Quinton. Mapping recurrences on parallel architectures. In *Third International Conference on Supercomputing*, Boston, Massachusetts, May 15-20 1988.
- [18] I. V. Ramakrishnan, D. Fussell, and A. Silberschatz. Mapping homogeneous graphs on linear arrays. *IEEE Transactions on Computers*, C-35(3):198-209, March 1986.
- [19] S. K. Rao and T. Kailath. Regular iterative algorithms and their implementation on processor arrays. *Proceedings of the IEEE*, 76(3):259-269, March 1988.
- [20] H. B. Ribas. *Automatic Generation of Systolic Programs from Nested Loops*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, June 1990.
- [21] H. B. Ribas. Obtaining dependence vectors for nested-loop computations. In *Proceedings of the International Conference on Parallel Processing*, pages II-212-219, 1990.
- [22] W. Shang and J. A. B. Fortes. On time mapping of uniform dependence algorithms into lower dimensional processor arrays. *IEEE Transactions on Parallel and Distributed Systems*, 3(3):350-363, May 1992.
- [23] Y. Wong and J. Delosme. Optimal systolic implementations of n-dimensional recurrences. In *Proceedings of the International Conference on Computer Design*, pages 618-621. IEEE, 1985.




Figure 2: In Example 4, the space mapping matrix $S = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$. Then tokens, which are generated or used in PEs marked by blank (dotted, netted, black) circles, may collide with the token generated or used in the PE marked by the blank (dotted, netted, black) square, respectively.





Figure 3: In Example 5, the PE and the space-time mapping during the infeasible computation with $H = (1, 2, 2)$ and $S = (1, 1, -1)$ for the matrix multiplication. Note that $B03$ collides with $B10$, $B13$ collides with $B20$, and $B23$ collides with $B30$ in the virtual data link 2.