


## 5 Code Parallelization


Processor with **instruction level parallelism (ILP)** executes several instructions in parallel.

Classes of processors and parallelism:  
 VLIW, super scalar  
 Pipelined processors  
 Data parallel processors




Compiler **analyzes sequential programs to exhibit potential parallelism** on instruction level;

model **dependences between computations**



Compiler arranges instructions for shortest execution time: **instruction scheduling**

Compiler **analyzes loops** to execute them in parallel  
**loop transformation**  
**array transformation**



## 5.1 Instruction Scheduling Data Dependence Graph

Exhibit potential **fine-grained parallelism** among operations.  
 Sequential code is over-specified!

**Data dependence graph (DDG)** for a basic block:


**Node:** operation;

**Edge**  $a \rightarrow b$ : operation  $b$  uses the result of operation  $a$

**Example for a basic block:**

- 1:  $t1 := a$
- 2:  $t2 := b$
- 3:  $t3 := t1 + t2$
- 4:  $x := t3$
- 5:  $t4 := c$
- 6:  $t5 := t3 + t4$
- 7:  $y := t5$
- 8:  $t6 := d$
- 9:  $t7 := e$
- 10:  $t8 := t6 + t7$
- 11:  $z := t8$

**data dependence graph**




**ti** are symbolic registers, store intermediate results, obey single assignment rule

## List Scheduling

**Input:** data dependence graph

**Output:** a schedule of **at most k operations per cycle**, such that all **dependences point forward**; DDG arranged in levels

**Algorithm:** A **ready list** contains all operations that are **not yet scheduled**, but whose **predecessors are scheduled**  
 Iterate: **select** from the ready list up to  $k$  operations for the next cycle (heuristic), **update** the ready list



- Algorithm is **optimal** only for **trees**.
- **Heuristic:** Keep ready list sorted by distance to an end node, e. g.

(1 2 5) (8 9 3) (6 10 4) (7 11)


without this heuristic:  
 (1 8 9) (2 5 10) (3 11) (6 4) (7)

( ) operations in one cycle

**Critical paths** determine minimal schedule length: e. g.  $1 \rightarrow 3 \rightarrow 6 \rightarrow 7$

## Variants and Restrictions for List Scheduling


- Allocate **as soon as possible**, ASAP (C-5.3); as **late** as possible, ALAP
- Operations have **unit execution time** (C-5.3); **different execution times:** selection avoids conflicts with already allocated operations
- Operations only on **specific functional units** (e. g. 2 int FUs, 2 float FUs)
- **Resource restrictions** between operations, e. g.  $\leq 1$  load or store per cycle



Scheduled DDG models **number of needed registers:**  
 • arc represents the use of an intermediate result

• **cut width** through a level gives the number of **registers needed**

The tighter the schedule the more registers are needed (*register pressure*).



### Instruction Scheduling for Pipelining

Instruction pipeline with 3 stages:



Dependent instructions may not follow one another immediately.

Schedule rearranges the operation sequence, to minimize the number of delays:

**without scheduling:**

- 1: t1 := a
- 2: t2 := b
- nop
- 3: t3 := t1 + t2
- nop
- 4: x := t3
- 5: t4 := c
- nop
- 6: t5 := t3 + t4
- nop
- 7: y := t5
- 8: t6 := d
- 9: t7 := e
- nop
- 10: t8 := t6 + t7
- nop
- 11: z := t8

- 1: t1 := a
  - 2: t2 := b
  - 5: t4 := c
  - 3: t3 := t1 + t2
  - 8: t6 := d
  - 9: t7 := e
  - 6: t5 := t3 + t4
  - 10: t8 := t6 + t7
  - 4: x := t3
  - 7: y := t5
  - 11: z := t8
- with scheduling**  
**no delays**

### Instruction Scheduling Algorithm for Pipelining

Algorithm: modified list scheduling:

Select from the ready list such that the selected operation


- has a sufficient **distance to all predecessors** in DDG
- has **many successors** (heuristic)
- has a **long path to the end node** (heuristic)

Insert an empty operation if none is selectable.

Ready list with additional information:

opr.	1	2	5	8	9	3	6	4	10	7	11
succ #	1	1	1	1	1	2	1	0	1	0	0
to end	3	3	2	2	2	2	1	1	1	0	0
sched. cycle	1	2	3	5	6	4	7	9	8	10	11

**data dependence graph**




**cycle**

- 1: t1 := a
  - 2: t2 := b
  - 3: t4 := c
  - 4: t3 := t1 + t2
  - 5: t6 := d
  - 6: t7 := e
  - 7: t5 := t3 + t4
  - 8: t8 := t6 + t7
  - 9: x := t3
  - 10: y := t5
  - 11: z := t8
- with scheduling**


### Reused registers: anti- and output-dependences

- u → v **flow-dependence:**  
u writes before v uses
- u  $\xrightarrow{a}$  v **anti-dependence:**  
u uses a value before v overwrites it
- u  $\xrightarrow{o}$  v **output-dependence:**  
u writes before v overwrites

**DDG with symbolic registers ti flow-dependences only**



**DDG with reused registers ti flow, anti-, and output-dependences**



### DDG with Loop Carried Dependences

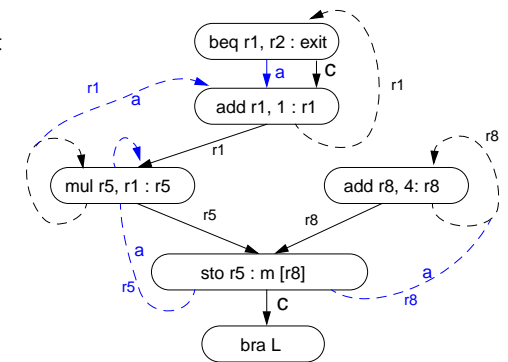
Factorial computation:

```

program:
i = 0; f = 1;
while ( i != n)
{
  i = i + 1;
  f = f * i;
  m[i] = f;
}

seq. machine code:
L: beq r1, r2 : exit
  add r1, 1 : r1
  mul r5, r1 : r5
  add r8, 4 : r8
  sto r5 : m[r8]
  bra L
    
```

Data dependence graph:



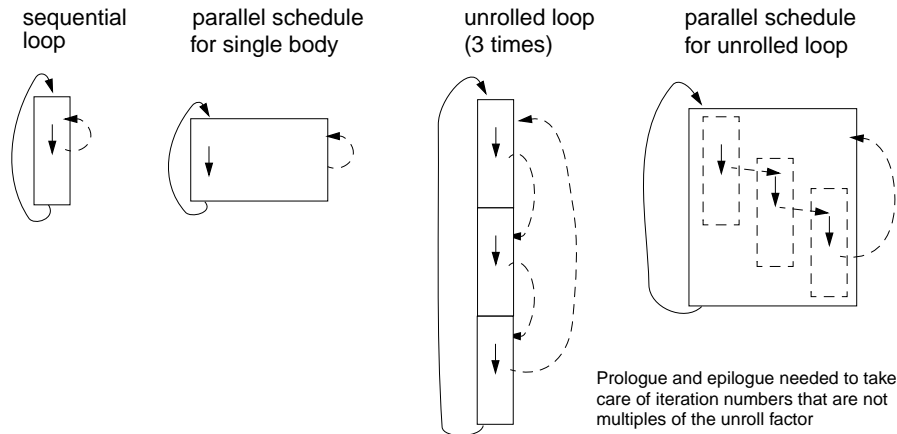
- u → v **flow-dependence:**  
u writes before v uses
- u - - - -> v **flow-dependence** into subsequent iteration
- u  $\xrightarrow{a}$  v **anti-dependence:**  
u uses a value before v overwrites it
- u  $\xrightarrow{o}$  v **output-dependence:**  
u writes before v overwrites
- u  $\xrightarrow{c}$  v **control-dependence:**  
u has to be executed before v (u or v may branch)

## Loop unrolling

C-5.6u

Loop unrolling: A technique for parallelization of loops.

A single loop body does not exhibit enough parallelism => sparse schedule.  
**Schedule the code (copies) of several adjacent iterations together** => more compact schedule



© 2009 bei Prof. Dr. Uwe Kastens

## Software Pipelining

C-5.7


Software Pipelining: A technique for parallelization of loops.

A single loop body does not exhibit enough parallelism => sparse schedule.  
**Overlap the execution of several adjacent iterations** => compact schedule

### The pipelined loop body

has **each operation** of the original sequential body, they belong to **several iterations**, they are **tightly scheduled**, its length is the **initiation interval II**, is **shorter** than the original body.

**Prologue, epilogue:** initiation and finalization code




© 2011 bei Prof. Dr. Uwe Kastens

## Transform Loops by Software Pipelining

C-5.8


### Technique:

1. **Data dependence graph** for the loop body, include **loop carried dependences**.
2. Chose a **small initiation interval II** - not smaller than  $\#instructions / \#FUs$
3. Make a „**Modulo Schedule**“  $s$  for the loop body: Two instructions can not be scheduled on the same FU,  $i_1$  in cycle  $c_1$  and  $i_2$  in cycle  $c_2$ , if  $c_1 \bmod II = c_2 \bmod II$
4. If (3) does not succeed without conflict, increase II and repeat from 3
5. Allocate the instructions of  $s$  in the new loop of length II:  $i_j$  scheduled in cycle  $c_j$  is allocated to  $c_j \bmod II$
6. Construct prologue and epilogue.



Modulo schedule for a loop body

cycle	0	1	2	3	4	5
0	11					
1						
2		12				
3	13			14		
4						
5		15				



© 2011 bei Prof. Dr. Uwe Kastens

## Result of Software Pipelining

C-5.10

t	t <sub>m</sub>	ADD	MUL	MEM	CTR
0	0	L:			beq r1, r2:exit
1	1	add r1, 1: r1			
2	0	add r8, 4: r8	mul r5, r1: r5		
3	1		... mul		
4	0			sto r5: m r8	
5	1			... sto	
6	0				
7	1				bra L

4 dedicated FUs schedule of the loop body for  $II = 2$   
 mul and sto need 2 cycles  
 add and sto in  $t_m=0$ , sto reads r8 before add writes it  
 bra not in cycle 6, it collides with beq:  $t_m=0$

t	t <sub>m</sub>	ADD	MUL	MEM	CTR
0	0				beq r1; r2: exit
1	1	add r1, 1: r1			
2	0	add r8, 4: r8	mul r5, r1: r5		beq r1; r2: ex
3	1	add r1, 1: r1	... mul		
4	0	L:	add r8, 4: r8	mul r5, r1: r5	sto r5: m r8
5	1		add r1, 1: r1	... mul	... sto
6	1	ex:	... mul	... sto	
7	0			sto r5: m r8	
8	1			... sto	
9	0				bra exit

**prologue**  
**software pipeline with  $II = 2$**   
**epilogue**

© 2009 bei Prof. Dr. Uwe Kastens

## 5.2 / 6. Data Parallelism: Loop Parallelization


**Regular loops** on orthogonal data structures - parallelized for **data parallel** processors

Development steps (automated by compilers):

- **nested loops** operating on **arrays**, sequential execution of iteration space
- analyze **data dependences**  
data-flow: definition and use of array elements
- **transform loops**  
keep data dependences forward in time
- **parallelize inner loop(s)**  
map to field or vector of processors
- **map arrays to processors**  
such that many accesses are local, transform index spaces

```

DECLARE B[0..N,0..N+1]
FOR I := 1 .. N
  FOR J := 1 .. I
    B[I,J] :=
      B[I-1,J]+B[I-1,J-1]
  END FOR
END FOR
    
```



## Iteration space of loop nests


**Iteration space** of a loop nest of depth n:

- **n-dimensional space of integral points** (polytope)
- each point  $(i_1, \dots, i_n)$  represents an execution of the innermost loop body
- loop bounds are in general not known before run-time
- iteration need not have orthogonal borders
- iteration is elaborated sequentially


example:  
computation of Pascal's triangle

```

DECLARE B[-1..N,-1..N]
FOR I := 0 .. N
  FOR J := 0 .. I
    B[I,J] :=
      B[I-1,J]+B[I-1,J-1]
  END FOR
END FOR
    
```




## Examples for Iteration spaces of loop nests




```

FOR I := 0 .. N
  FOR J := 0 .. I
    
```




```

FOR I := 0 .. N
  FOR J := 0..I BY 2
    
```



```


FOR I := 0..N BY 2
  FOR J := 0 .. I
    
```



```

FOR I := 0 .. N
  FOR J := I..I+M
    
```

M = 3, N = 4




```

FOR I := 0 .. M+N
  FOR J := max(0, I-M)..
    min(I, N)
    
```

## Data Dependences in Iteration Spaces

**Data dependence from iteration point  $i_1$  to  $i_2$ :**


- Iteration  $i_1$  computes a value that is used in iteration  $i_2$  (flow dependence)
- relative **dependence vector**  
 $d = i_2 - i_1 = (i_{2_1} - i_{1_1}, \dots, i_{2_n} - i_{1_n})$   
holds for all iteration points except at the border
- Flow-dependences can **not be directed against the execution order**, can not point backward in time: each dependence vector must be **lexicographically positive**, i. e.  $d = (0, \dots, 0, d_j, \dots), d_j > 0$



Example:  
Computation of Pascal's triangle

```

DECLARE B[-1..N,-1..N]
FOR I := 0 .. N
  FOR J := 0 .. I
    B[I,J] :=
      B[I-1,J]+B[I-1,J-1]
  END FOR
END FOR
    
```



## Loop Transformation

The **iteration space** of a loop nest is transformed to **new coordinates**. Goals:

- **execute innermost loop(s) in parallel**
- improve **locality** of data accesses;  
**in space**: use storage of executing processor,  
**in time**: reuse values stored in cache
- **systolic** computation and communication scheme

Data dependences must **point forward in time**, i.e. **lexicographically positive** and **not within parallel dimensions**


**linear basic transformations:**

- **Skewing**: add iteration count of an outer loop to that of an inner one
- **Reversal**: flip execution order for one dimension
- **Permutation**: exchange two loops of the loop nest


**SRP transformations** (next slides)

**non-linear transformations**, e. g.

- **Scaling**: stretch the iteration space in one dimension, causes gaps
- **Tiling**: introduce **additional inner loops** that **cover tiles** of fixed size



## Transformations of



## Transformations defined by matrices

Transformation matrices: systematic transformation, check dependence vectors

Reversal 
$$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} = \begin{pmatrix} i \\ -j \end{pmatrix} = \begin{pmatrix} i' \\ j' \end{pmatrix}$$

Skewing 
$$\begin{pmatrix} 1 & 0 \\ f & 1 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} = \begin{pmatrix} i \\ f*i+j \end{pmatrix} = \begin{pmatrix} i' \\ j' \end{pmatrix}$$

Permutation 
$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} = \begin{pmatrix} j \\ i \end{pmatrix} = \begin{pmatrix} i' \\ j' \end{pmatrix}$$

## Reversal


**Iteration count of one loop is negated**, that dimension is enumerated backward

general transformation matrix

$$\begin{pmatrix} 1 & & & \\ \dots & & & 0 \\ & 1 & & \\ & & -1 & \\ 0 & & & 1 \dots \end{pmatrix}$$

```

for i = 0 to M
  for j = 0 to N
    ...
    
```




2-dimensional:

loop variables  
old                      new

$$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} = \begin{pmatrix} i \\ -j \end{pmatrix} = \begin{pmatrix} i' \\ j' \end{pmatrix}$$

```

for ir = 0 to M
  for jr = -N to 0
    ...
    
```




### Skewing

The **iteration count** of an outer loop is **added to the count of an inner loop**; iteration space is shifted; **execution order** of iteration points **remains unchanged**

general transformation matrix:

$$\begin{pmatrix} 1 & & & 0 \\ \dots & & & \\ f & 1 & & \\ 0 & & 1 & \dots \\ & & & 1 \end{pmatrix}$$

```
for i = 0 to M
  for j = 0 to N
    ...
```




original

2-dimensional:

$$\begin{matrix} & & \text{loop variables} \\ & & \text{old} & \text{new} \\ \begin{pmatrix} 1 & 0 \\ f & 1 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} = \begin{pmatrix} i \\ f*i+j \end{pmatrix} = \begin{pmatrix} is \\ js \end{pmatrix} \end{matrix}$$

```
for is = 0 to M
  for js = f*is to N+f*is
    ...
```



transformed


### Permutation

**Two loops of the loop nest are interchanged**; the iteration space is flipped; the **execution order** of iteration points **changes**; new dependence vectors must be legal.

general transformation matrix:

$$\begin{matrix} i \\ j \end{matrix} \begin{pmatrix} 1 & & & 0 \\ 0 & 1 & & \\ & & 1 & 0 \\ 0 & & & 1 \end{pmatrix}$$

```
for i = 0 to M
  for j = 0 to N
    ...
```




original

2-dimensional:

$$\begin{matrix} & & \text{loop variables} \\ & & \text{old} & \text{new} \\ \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} = \begin{pmatrix} j \\ i \end{pmatrix} = \begin{pmatrix} ip \\ jp \end{pmatrix} \end{matrix}$$

```
for ip = 0 to N
  for jp = 0 to M
    ...
```

transformed



### Use of Transformation Matrices

- Transformation matrix **T** defines **new iteration counts** in terms of the old ones:  $T * i = i'$

e. g. Reversal  $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} = \begin{pmatrix} i \\ -j \end{pmatrix} = \begin{pmatrix} i' \\ j' \end{pmatrix}$

- Transformation matrix **T** transforms old **dependence vectors** into new ones:  $T * d = d'$

e. g.  $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} * \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$

- inverse Transformation matrix  $T^{-1}$  defines **old iteration counts** in terms of new ones, for transformation of index expressions in the loop body:  $T^{-1} * i' = i$

e. g.  $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} * \begin{pmatrix} i' \\ j' \end{pmatrix} = \begin{pmatrix} i' \\ -j' \end{pmatrix} = \begin{pmatrix} i \\ j \end{pmatrix}$

- concatenation of transformations** first  $T_1$  then  $T_2$ :  $T_2 * T_1 = T$

e. g.  $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} * \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$


### Inequalities Describe Loop Bounds

The bounds of a loop nest are described by a **set of linear inequalities**. Each **inequality separates the space** in „inside and outside of the iteration space“:

$$B * i \leq c$$


example 1  $\begin{pmatrix} -1 & 0 \\ 1 & 0 \\ 0 & -1 \\ 0 & 1 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} \leq \begin{pmatrix} 0 \\ M \\ 0 \\ N \end{pmatrix}$

- 1  $-i \leq 0$
- 2  $i \leq M$
- 3  $-j \leq 0$
- 4  $j \leq N$



example 2  $\begin{pmatrix} -1 & 1 \\ 1 & 0 \\ 0 & -1 \\ 0 & 1 \end{pmatrix} * \begin{pmatrix} i \\ j \end{pmatrix} \leq \begin{pmatrix} 0 \\ M \\ 0 \\ N \end{pmatrix}$

- 1  $-i + j \leq 0$
- 2  $i \leq M$
- 3  $-j \leq 0$
- 4  $j \leq N$



**positive** factors represent **upper** bounds  
**negative** factors represent **lower** bounds

1, 4:  $j \leq \min(i, N)$     1+ 3:  $0 \leq i$   
3:  $0 \leq j$                     2:  $i \leq M$

### Transformation of Loop Bounds

The inverse of a transformation matrix  $T^{-1}$  transforms a set of inequalities:  $B * T^{-1} i' \leq c$

skewing      inverse      B       $T^{-1}$        $B * T^{-1}$

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} -1 & 0 \\ 1 & 0 \\ 0 & -1 \\ 0 & 1 \end{pmatrix} * \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix} = \begin{pmatrix} -1 & 0 \\ 1 & 0 \\ 1 & -1 \\ -1 & 1 \end{pmatrix}$$

example 1  
new bounds:

$$\begin{pmatrix} -1 & 0 \\ 1 & 0 \\ 1 & -1 \\ -1 & 1 \end{pmatrix} * \begin{pmatrix} i' \\ j' \end{pmatrix} \leq \begin{pmatrix} 0 \\ M \\ 0 \\ N \end{pmatrix}$$

1  $-i' \leq 0$   
2  $i' \leq M$   
3  $i' - j' \leq 0$   
4  $-i' + j' \leq N$


### Example for Transformation and Parallelization of a Loop

```
for i = 0 to N
  for j = 0 to M
    a[i, j] = (a[i, j-1] + a[i-1, j]) / 2;
```

Parallelize the above loop.

1. Draw the iteration space.
2. Compute the dependence vectors and draw examples of them into the iteration space. Why can the inner loop not be executed in parallel?
3. Apply a skewing transformation and draw the iteration space.
4. Apply a permutation transformation and draw the iteration space. Explain why the inner loop now can be executed in parallel.
5. Compute the matrix of the composed transformation and use it to transform the dependence vectors.
6. Compute the inverse of the transformation matrix and use it to transform the index expressions.
7. Specify the loop bounds by inequalities and transform them by the inverse of the transformation matrix.
8. Write the complete loops with new loop variables ip and jp and new loop bounds.

### Solution of the Transformation and Parallelization Example



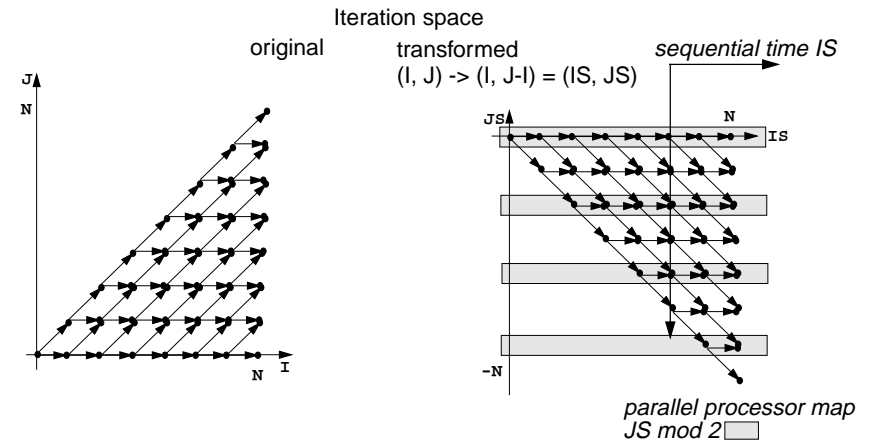
5.:  $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$       6.: Inverse  $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$        $\begin{pmatrix} 0 & 1 \\ 1 & -1 \end{pmatrix}$

7. Bounds: orig.:  $B \begin{pmatrix} -1 & 0 \\ 1 & 0 \\ 0 & -1 \\ 0 & 1 \end{pmatrix} c \begin{pmatrix} 0 \\ N \\ 0 \\ M \end{pmatrix}$  new:  $B * T^{-1} \begin{pmatrix} 0 & -1 \\ 0 & 1 \\ -1 & 1 \\ 1 & -1 \end{pmatrix}$

1  $-jp \leq 0$       1, 3  $\Rightarrow 0 \leq ip$   
2  $jp \leq N$       2, 4  $\Rightarrow ip \leq M+N$   
3  $-ip+jp \leq 0$       1, 4  $\Rightarrow \max(0, ip-M) \leq jp$   
4  $ip - jp \leq M$       2, 3  $\Rightarrow jp \leq \min(ip, N)$

```
8. for ip = 0 to M+N
  for jp = max(0, ip-M) to min(ip, N)
    a[jp, ip-jp] = (a[jp, ip-jp-1] + a[jp-1, ip-jp]) / 2;
```

### Transformation and Parallelization



```
DECLARE B[-1..N,-1..N]
FOR I := 0 .. N
  FOR J := 0 .. I
    B[I,J] :=
      B[I-1,J]+B[I-1,J-1]
  END FOR
END FOR
```

```
DECLARE B[-1..N,-1..N]
FOR IS := 0 .. N
  FOR JS := -IS .. 0
    B[IS,JS+IS] :=
      B[IS-1,JS+IS]+B[IS-1,JS-1+IS]
  END FOR
END FOR
```

## Data Mapping

### Goal:

**Distribute array elements** over processors, such that as many **accesses as possible are local**.

### Index space

of an array:  
n-dimensional space of integral index points (polytope)

- **same properties as iteration space**
- same mathematical model
- same **transformations** are applicable (Skewing, Reversal, Permutation, ...)
- **no restrictions** by data dependences

## Data distribution for parallel loops

