

Embedded Parallel Computing

Parallel programming with Ambric

EPC Lecture 2

- Parallel programming concepts
- Lab2 Introduction
- Matrix Multiplication

Parallel thinking

- Know your application
 - Natural ways to decompose your application into tasks
 - Data access patterns
 - Data dependencies
- Be prepared to step beyond what you know
 - Requires a mix of what are traditionally hardware parallelization techniques and software techniques
 - Mix of optimizations at various levels: algorithm, data flow, instruction sequence, etc.
 - One algorithm choice may fit more naturally
- Parallelizing for Ambric is not like
 - OpenMP or POSIX Threads (Pthreads)
 - Both rely on Shared Memory Parallelism
 - Single Program Multiple Data (i.e., Vector Parallel)

Compared to Developing Hardware

- | | |
|--|--|
| <ul style="list-style-type: none">■ Things which may be familiar<ul style="list-style-type: none">□ Partitioning of tasks bounded by minimally sized memories□ Assigning specific portions of the chip to specific tasks | <ul style="list-style-type: none">■ Things which may be unfamiliar<ul style="list-style-type: none">□ Interconnect is always 33 bits□ All blocks have the same resources□ No timing constraints |
|--|--|

Compared to Developing Software

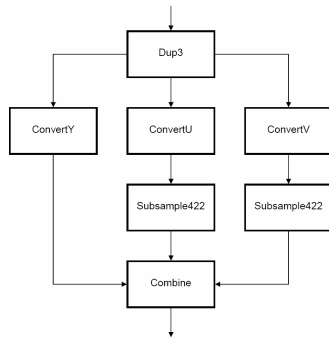
- | | |
|--|---|
| <ul style="list-style-type: none">■ Things which may be familiar<ul style="list-style-type: none">□ Analyzing your code to locate data dependencies and find available "threads"□ Writing and optimizing your code | <ul style="list-style-type: none">■ Things which may be unfamiliar<ul style="list-style-type: none">□ You require structure (aStruct) around code objects.□ More deterministic performance of "threads" because of encapsulation and lack of task switching□ Smaller available memories; generally local |
|--|---|

Parallelism - Introduction

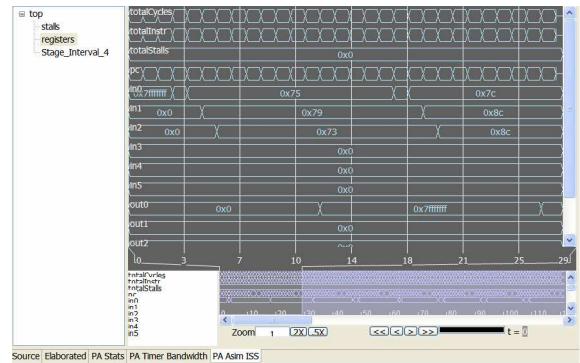
- **Task Level parallelism**
 - Blocks with clearly defined interfaces are good candidates.
 - Tasks may be reusable in different applications (i.e., DCT)
 - First consider your block diagram
 - Validate the I/O requirements of each block
- **Data Level Parallelism**
 - Select partitioning such that there are no (or limited) data dependencies between parallel elements!
 - Fanout data to multiple cores, process, and fanin
 - Fanout and fanin should use less valuable resources (SR's)
 - Overall resource utilization should be balanced

Lab2 – Parallel Color Space Conversion

- Block diagram
- Parallel execution
- Pipelining



Performance Analysis




Matrix Multiplication Example

- Matrix multiplication

```


for (i=0; i<n; i++){
  for (j=0; j<n; j++){
    C[i][j] = 0;
    for(k=0; k<n; k++){
      C[i][j] = C[i][j] + A[i][k]*B[k][j];
    }
  }
}
    
```

Matrix Multiplication On a Ring




Example of matrix mapping

Matrix Multiplication Cont.



In step 1 all processors compute a part of the final matrix

Matrix Multiplication Cont.



In step 2 the sub matrix B stored in each processor are rotated left, and a new part of the final matrix is computed.

Matrix multiplication on a Mesh

■ Mesh algorithm

- assume mesh with wraparound

Step 1: Initialize matrix A by shifting row[i] i steps left ($0 \leq i < n$)
 Initialize matrix B by shifting column[j] j steps up ($0 \leq j < n$)
 Step 2: Multiply elements of A and B
 Step 3: Shift matrix A one step left
 Shift matrix B one step up
 Step 4: Multiply elements of A and B and add to C
 Step 5: repeat step 3 and step 4 n-1 times (With n rows and columns)

This also works with blocks.
 Sub matrices are shifted instead of elements
 (n are the number of rows and columns of processors)

Matrix multiplication on a Mesh

A[0][0]	A[0][1]	A[0][2]	A[0][3]
A[1][0]	A[1][1]	A[1][2]	A[1][3]
A[2][0]	A[2][1]	A[2][2]	A[2][3]
A[3][0]	A[3][1]	A[3][2]	A[3][3]

B[0][0]	B[0][1]	B[0][2]	B[0][3]
B[1][0]	B[1][1]	B[1][2]	B[1][3]
B[2][0]	B[2][1]	B[2][2]	B[2][3]
B[3][0]	B[3][1]	B[3][2]	B[3][3]

Matrix multiplication on a Mesh

Initialization

A[0][0]	A[0][1]	A[0][2]	A[0][3]
A[1][0]	A[1][1]	A[1][2]	A[1][3]
A[2][2]	A[2][3]	A[2][0]	A[2][1]
A[3][3]	A[3][0]	A[3][1]	A[3][2]

B[0][0]	B[0][1]	B[0][2]	B[0][3]
B[1][0]	B[1][1]	B[1][2]	B[1][3]
B[2][0]	B[2][1]	B[2][2]	B[2][3]
B[3][0]	B[3][1]	B[3][2]	B[3][3]

Rotate A matrix left
 Rotate row[i] i steps

Rotate B matrix up
 Rotate col[j] j steps

Matrix multiplication on a Mesh

Initialization

A[0][0]	A[0][1]	A[0][2]	A[0][3]
A[1][0]	A[1][1]	A[1][2]	A[1][3]
A[2][2]	A[2][3]	A[2][0]	A[2][1]
A[3][0]	A[3][1]	A[3][2]	A[3][3]

B[0][0]	B[0][1]	B[0][2]	B[0][3]
B[1][0]	B[1][1]	B[1][2]	B[1][3]
B[2][0]	B[2][1]	B[2][2]	B[2][3]
B[3][0]	B[3][1]	B[3][2]	B[3][3]

Rotate A matrix left
 Rotate row[i] i steps

Rotate B matrix up
 Rotate col[j] j steps

Matrix multiplication on a Mesh


Multiplication loop

A[0][0]	A[0][1]	A[0][2]	A[0][3]
A[1][0]	A[1][1]	A[1][2]	A[1][3]
A[2][2]	A[2][3]	A[2][0]	A[2][1]
A[3][0]	A[3][1]	A[3][2]	A[3][3]

B[0][0]	B[0][1]	B[0][2]	B[0][3]
B[1][0]	B[1][1]	B[1][2]	B[1][3]
B[2][0]	B[2][1]	B[2][2]	B[2][3]
B[3][0]	B[3][1]	B[3][2]	B[3][3]


Rotate A matrix left 1 step Rotate B matrix up one step

Exercise




Start: Every Process contains a part of the total matrix.
 Rows of A and columns of B

Exercise




Step 1: Every process multiplies their local A and B matrices to form a part of the C matrix

Exercise




Step 2: Every process sends their B matrix one step left and multiplies the local matrices

Exercise




Step 2: Every process sends their B matrix one step left and multiplies the local matrices

Exercise




Step 2: Every process sends their B matrix one step left and multiplies the local matrices

Exercise




Start: Every process contains a part of the A matrix and one part of the B matrix

Exercise




Init: Every process shifts the A matrix left as many steps as its row coordinate. Every process shifts the B matrix up as many steps as its column coordinate.

Exercise




Step 1: Do a local matrix multiplication $C=AB$

Exercise



Step 2: Shift the A matrix one step left.
Shift the B matrix one step up

Exercise



Step 3: Do a local matrix multiplication $C=C+AB$

Exercise

	0	1	2	3		0	1	2	3
0	1	2	3	4		1	2	3	4
1	5	6	7	8		5	6	7	8
2	9	10	11	12		9	10	11	12
3	13	14	15	16		13	14	15	16
	A					B			

Step by step example of the matrix multiplication.
The matrix values in this example could be sub matrices instead.
Assume one value in each processor and the processors are connected in a Mesh with wraparound.

Exercise

	0	1	2	3		0	1	2	3
0	1	2	3	4		1	6	11	16
1	6	7	8	5		5	10	15	4
2	11	12	9	10		9	14	3	8
3	16	13	14	15		13	2	7	12
	A					B			

Init: Shift A matrix left as many steps as the mesh row index.
Shift B matrix as many step up as the mesh column index.

Exercise

	0	1	2	3		0	1	2	3
0	1	2	3	4		1	6	11	16
1	6	7	8	5		5	10	15	4
2	11	12	9	10		9	14	3	8
3	16	13	14	15		13	2	7	12
	A					B			

$C_{1,3}=5 \times 4$

Mult: Do a local matrix multiplication. Here we look at process 1,3.

Exercise

	0	1	2	3		0	1	2	3	
0	2	3	4	1		5	10	15	4	$C_{1,3}=5*4+6*8$
1	7	8	5	6		9	14	3	8	
2	12	9	10	11		13	2	7	12	
3	13	14	15	16		1	6	11	16	
	A					B				

Step 1: Shift A matrix one step left and B matrix one step up.
Do a local matrix multiplication $C=C+AB$.

Exercise

	0	1	2	3		0	1	2	3	
0	3	4	1	2		9	14	3	8	$C_{1,3}=5*4+6*8 +7*12$
1	8	5	6	7		13	2	7	12	
2	9	10	11	12		1	6	11	16	
3	14	15	16	13		5	10	15	4	
	A					B				

Step 2: Shift A matrix one step left and B matrix one step up.
Do a local matrix multiplication $C=C+AB$.

Exercise

	0	1	2	3		0	1	2	3	
0	4	1	2	3		13	2	7	12	$C_{1,3}=5*4+6*8 +7*12+8*16$
1	5	6	7	8		1	6	11	16	
2	10	11	12	9		5	10	15	4	
3	15	16	13	14		9	14	3	8	
	A					B				

Step 3: Shift A matrix one step left and B matrix one step up.
Do a local matrix multiplication $C=C+AB$.
Done!