Plan





1 Parallelism Complexity Measures

2 cilk_for Loops

3 Scheduling Theory and Implementation

4 Measuring Parallelism in Practice

5 Announcements



3

We shall also call this model **multithreaded parallelism**.

(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	3 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	4 / 62

Parallelism Complexity Measures

Terminology

(Moreno Maz



- a strand is is a maximal sequence of instructions that ends with a spawn, sync, or return (either explicit or implicit) statement.
- At runtime, the *spawn* relation causes procedure instances to be structured as a rooted tree, called spawn tree or parallel instruction stream, where dependencies among strands form a dag.



We define several performance measures. We assume an ideal situation: no cache issues, no interprocessor costs:

 T_p is the minimum running time on p processors

Parallelism Complexity Measures

- T_1 is called the **work**, that is, the sum of the number of instructions at each node.
- T_{∞} is the minimum running time with infinitely many processors, called

	$\bullet \Box \bullet \bullet$	@ + * 폰 + * 폰 + _ 폰	$\mathcal{O}\mathcal{Q}\mathcal{O}$	the span	< □ ▶ <	□ > < 글 > < 글 > _ 글	うくで
(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	5 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	6 / 62
Parallelism Complexity Measures				Paralle	lism Complexity Measures		
The critical path	n length			Work law			

VVOIK IAN

Work and span



Assuming all strands run in unit time, the longest path in the DAG is equal to T_{∞} . For this reason, T_{∞} is also referred to as the critical path length.



- We have: $T_p \geq T_1/p$.
- Indeed, in the best case, p processors can do p works per unit of time. ◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 善臣 - のへで

a)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	7 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	8 / 62

◆□ ▶ ◆□ ▶ ◆臣 ▶ ◆臣 ▶ ○臣 ○ のへで

Span law	Speedup on <i>p</i> processors
	 T₁/T_p is called the speedup on p processors A parallel program execution can have: linear speedup: T₁/T_P = Θ(p) superlinear speedup: T₁/T_P = ω(p) (not possible in this model, though it is possible in others) sublinear speedup: T₁/T_P = o(p)
• We have: $T_p \ge T_{\infty}$. • Indeed, $T_p < T_{\infty}$ contradicts the definitions of T_p and T_{∞} . (Moreno Maza) Multithreaded Parallelism and Performance M CS 4435 - CS 9624 9 / 62 Parallelism Complexity Measures	(Moreno Maza) Parallelism Complexity Measures Multithreaded Parallelism and Performance M CS 4435 - CS 9624 10 / 62 Parallelism Complexity Measures
Because the Span Law dictates that $T_p \ge T_{\infty}$, the maximum possible speedup given T_1 and T_{∞} is $T_1/T_{\infty} = parallelism$ = the average amount of work per step along the span.	• For Fib(4), we have $T_1 = 17$ and $T_\infty = 8$ and thus $T_1/T_\infty = 2.125$. • What about $T_1(\text{Fib}(n))$ and $T_\infty(\text{Fib}(n))$?
(Moreno Maza) Multithreaded Parallelism and Performance N CS 4435 - CS 9624 11 / 62	(Moreno Maza) Multithreaded Parallelism and Performance M CS 4435 - CS 9624 12 / 62

Parallelism Complexity Measures

Parallelism Complexity Measures

Series composition

- The Fibonacci example (2/2)
 - We have $T_1(n) = T_1(n-1) + T_1(n-2) + \Theta(1)$. Let's solve it.
 - One verify by induction that T(n) ≤ aF_n − b for b > 0 large enough to dominate Θ(1) and a > 1.
 - We can then choose *a* large enough to satisfy the initial condition, whatever that is.
 - On the other hand we also have $F_n \leq T(n)$.
 - Therefore $T_1(n) = \Theta(F_n) = \Theta(\psi^n)$ with $\psi = (1 + \sqrt{5})/2$.
 - We have $T_\infty(n) = \max(T_\infty(n-1), T_\infty(n-2)) + \Theta(1).$
 - We easily check $T_\infty(n-1) \geq T_\infty(n-2)$.
 - This implies $T_{\infty}(n) = T_{\infty}(n-1) + \Theta(1)$.
 - Therefore $T_{\infty}(n) = \Theta(n)$.
 - Consequently the parallelism is $\Theta(\psi^n/n)$.



• Work?

• Span?

	< □ ≻ <	(□) + (=) + (=) = =	500		< 🗆	→ < @ > < E > < E > < E < E < E < E < E < E <	- Dac
(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	13 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	14 / 62
Parallelism Complexity Measures			Parallelism Complexity Measures				
Series composition	on			Parallel composition	on		



- Work: $T_1(A \cup B) = T_1(A) + T_1(B)$
- Span: $T_{\infty}(A \cup B) = T_{\infty}(A) + T_{\infty}(B)$



- Work?
- Span?

-

Parallel composition

Some results in the fork-join parallelism model

A	
B	

• Work: $T_1(A \cup B) = T_1(A) + T_1(B)$ • Span: $T_{\infty}(A \cup B) = \max(T_{\infty}(A), T_{\infty}(B))$

Algorithm	Work	Span
Merge sort	Θ(n lg n)	Θ(lg³n)
Matrix multiplication	Θ(n ³)	Θ(lg n)
Strassen	Θ(n ^{lg7})	Θ(lg²n)
LU-decomposition	Θ(n ³)	Θ(n lg n)
Tableau construction	Θ(n ²)	$\Omega(n^{lg3})$
FFT	Θ(n lg n)	Θ(lg²n)
Breadth-first search	Θ(Ε)	Θ(d lg V)

We shall prove those results in the next lectures.



	4 L		~) Q (*		< □	지 문 지 문 지	= •) <	C.
(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	19 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	20 / 6	2

cilk_for Loops

Implementation of for loops in Cilk++

Up to details (next week!) the previous loop is compiled as follows, using a **divide-and-conquer implementation**:

<pre>void recur(int lo, int hi) { if (hi > lo) { // coarsen int mid = lo + (hi - lo)/2; cilk_spawn recur(lo, mid); recur(mid, hi); cilk_sync; } else for (int j=0; j<i; ++j)="" <="" [2438-2010]="" a[i][j]="temp;" compare="" comparent="" double="" pre="" temp="A[i][j];" {="" }=""></i;></pre>	<image/> <text><text><list-item><list-item><list-item><list-item><code-block><text></text></code-block></list-item></list-item></list-item></list-item></text></text>
Parallelizing the inner loop	Plan
<pre>cilk_for (int i=1; i<n; (int="" ++i)="" ++j)="" a[i][j]="A[j][i];" a[j][i]="temp;" cilk_for="" double="" j="0;" j<i;="" pre="" temp="A[i][j];" {="" }="" }<=""></n;></pre>	 Parallelism Complexity Measures cilk_for Loops Scheduling Theory and Implementation
 Span of outer loop control: Θ(log(n)) Max span of an inner loop control: Θ(log(n)) Span of an iteration: Θ(1) Span: Θ(log(n)) Work: Θ(n²) Parallelism: Θ(n²/log(n)) But! More on this next week 	 Measuring Parallelism in Practice Announcements
(Moreno Maza) Multithreaded Parallelism and Performance M CS 4435 - CS 9624 23 / 62	(Moreno Maza) Multithreaded Parallelism and Performance M CS 4435 - CS 9624 24 / 62

Analysis of parallel for loops

Scheduling



A **scheduler**'s job is to map a computation to particular processors. Such a mapping is called a **schedule**.

- If decisions are made at runtime, the scheduler is *online*, otherwise, it is *offline*
- Cilk++'s scheduler maps strands onto processors dynamically at runtime.

(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4			
Scheduling Theory and Implementation					

Greedy scheduling (2/2)



- In any *greedy schedule*, there are two types of steps:
 - **complete step**: There are at least *p* strands that are ready to run. The greedy scheduler selects any *p* of them and runs them.
 - **incomplete step**: There are strictly less than *p* threads that are ready to run. The greedy scheduler runs them all.

Scheduling Theory and Implementation

Greedy scheduling (1/2)



- A strand is **ready** if all its predecessors have executed
- A scheduler is **greedy** if it attempts to do as much work as possible at every step.

▶ ★ Ξ ▶ ★ Ξ ▶ = Ξ.	$\mathcal{O}\mathcal{Q}$		< □ ▶	→ 御 → → 注 → → 注 → → 注	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~		
CS 4435 - CS 9624	25 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	26 / 62		
Scheduling Theory and Implementation							
Theorem of Graham and Brent							

P = 3

For any greedy schedule, we have T_p $~\leq~~T_1/p~+~T_\infty$

- #complete steps $\leq T_1/p$, by definition of T_1 .
- #incomplete steps $\leq T_{\infty}$. Indeed, let G' be the subgraph of G that remains to be executed immediately prior to a incomplete step.
 - (i) During this incomplete step, all strands that can be run are actually run
 - (ii) Hence removing this incomplete step from G'_{\Box} reduces T_{∞} by one

(Moreno Maza)

27 / 62

Scheduling Theory and Implementation	Scheduling Theory and Implementation
orollary 1	Corollary 2

A greedy scheduler is always within a factor of 2 of optimal.

From the work and span laws, we have:

$$T_P \ge \max(T_1/p, T_\infty) \tag{1}$$

In addition, we can trivially express:

$$T_1/p \le \max(T_1/p, T_\infty) \tag{2}$$

$$T_{\infty} \leq \max(T_1/p, T_{\infty}) \tag{3}$$

From Graham - Brent Theorem, we deduce:

$$T_P \leq T_1/p + T_\infty \tag{4}$$

$$\leq \max(T_1/p, T_\infty) + \max(T_1/p, T_\infty)$$
(5)

$$\leq 2 \max(T_1/p, T_\infty)$$
 (6)

. . .

The greedy scheduler achieves linear speedup whenever $T_{\infty} = O(T_1/p)$.

From Graham - Brent Theorem, we deduce:

$$T_{p} \leq T_{1}/p + T_{\infty} \tag{7}$$

$$= T_1/p + O(T_1/p)$$
 (8)

$$= \Theta(T_1/p) \tag{9}$$

The idea is to operate in the range where T_1/p dominates T_{∞} . As long as T_1/p dominates T_{∞} , all processors can be used efficiently. The quantity T_1/pT_{∞} is called the **parallel slackness**.

which concludes the p	root.		500		< 🗆 ۱	 < □ > < Ξ > < Ξ > Ξ 	596
(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	29 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	30 / 62
Scheduling The	eory and Implementation			Scheduling The	ory and Implementation		
The work-stealing	scheduler $(1/11)$			The work-stealing	scheduler $(2/11)$		



イロト イポト イヨト イヨト

3

・ロト ・ 日 ・ ・ 日 ・ ・ 日 ・

Scheduling Theory and Implementation	Scheduling Theory and Implementation
he work-stealing scheduler $(3/11)$	The work-stealing scheduler $(4/11)$



	< □ ≻ <		-		< □ >		990
(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	33 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	34 / 62
Scheduling Theory and Implementation			Scheduling Theo	ry and Implementation			
The work-stealing scheduler $(5/11)$		The work-stealing	scheduler ($6/11$)				



Scheduling Theory and Implementation	Scheduling Theory and Implementation	
The work-stealing scheduler $(7/11)$	The work-stealing scheduler $(8/11)$	
spawn	spawn	

call

call

call

P

P





call

call

call

Ρ

P

spawn

call

-sall

Ρ

all

call

Ρ

CS 4435 - CS 9624 39 / 62

▲□▶ ▲圖▶ ▲臣▶ ▲臣▶ ―臣 – のへで

(Moreno Maza)

spawn

call

call

call

P

call

Ρ

Scheduling Theory and Implementation	Scheduling Theory and Implementation
The work-stealing scheduler $(11/11)$	Performances of the work-stealing scheduler
spawn call call spawn spawn spawn spawn pp P P P P P P P P P P P P P	 Assume that each strand executes in unit time, for almost all "parallel steps" there are at least <i>p</i> strands to run, each processor is either working or stealing. Then, the randomized work-stealing scheduler is expected to run in T_P = T₁/p + O(T_∞) During a steal-free parallel steps (steps at which all processors have work on their deque) each of the <i>p</i> processors consumes 1 work unit. Thus, there is at most T₁/p steal-free parallel steps. During a parallel step with steals each thief may reduce by 1 the running time with a probability of 1/p Thus, the expected number of steals is O(p T_∞). Therefore, the expected running time
(Moreno Maza) Multithreaded Parallelism and Performance N CS 4435 - CS 9624 41 / 62 Scheduling Theory and Implementation	$T_P = (T_1 + O(p \ T_\infty))/p = T_1/p + O(T_\infty).$ (Moreno Maza) Multithreaded Parallelism and Performance M Scheduling Theory and Implementation
Overheads and burden	Span overhead
	• Let T_1, T_{∞}, T_p be given. We want to refine the <i>randomized</i>

- Obviously $T_1/p + T_{\infty}$ will over-estimate T_p in practice.
- Many factors (simplification assumptions of the fork-join parallelism model, architecture limitation, costs of executing the parallel constructs, overheads of scheduling) will make T_p smaller in practice.
- One may want to estimate the impact of those factors:
 - In the estimate of the randomized work-stealing complexity result
 - 2 by comparing a Cilk++ program with its C++ elision
 - 3 by estimating the costs of spawning and synchronizing

- work-stealing complexity result.
- The span overhead is the smallest constant c_{∞} such that

$$T_p \leq T_1/p + c_\infty T_\infty.$$

- Recall that T_1/T_{∞} is the maximum possible speed-up that the application can obtain.
- We call parallel slackness assumption the following property

$$T_1/T_{\infty} >> c_{\infty} p \tag{11}$$

that is, $c_{\infty} p$ is much smaller than the average parallelism .

• Under this assumption it follows that $T_1/p>>c_\infty T_\infty$ holds, thus c_∞ has little effect on performance when sufficiently slackness exists. ₹.

(Moreno Maza)

Scheduling Theory and Implementation

Work overhead

- Let T_s be the running time of the C++ elision of a Cilk++ program.
- We denote by *c*₁ the work overhead

 $c_1 = T_1 / T_s$

• Recall the expected running time: $T_P \leq T_1/P + c_{\infty}T_{\infty}$. Thus with the parallel slackness assumption we get

$$T_P \le c_1 T_s / p + c_\infty T_\infty \simeq c_1 T_s / p.$$
(12)

• We can now state the work first principle precisely

Minimize c_1 , even at the expense of a larger c_{∞} . This is a key feature since it is conceptually easier to minimize c_1

rather than minimizing c_{∞} .

• Cilk++ estimates T_p as $T_p = T_1/p + 1.7$ burden_span, where burden_span is 15000 instructions times the number of continuation edges along the critical path.

The cactus stack

1 Parallelism Complexity Measures

4 Measuring Parallelism in Practice

3 Scheduling Theory and Implementation

2 cilk_for Loops

Announcements



- A cactus stack is used to implement C's rule for sharing of function-local variables.
- A stack frame can only see data stored in the current and in the previous stack frames.

	• □ ▶ • 6	나 수준에 수준에 문	$\mathcal{O}\mathcal{Q}$		∢ □ ▶	→ □ → → 注 → → 注 → □ 目	9 A CP
(Moreno Maza) Multithreaded	Parallelism and Performance N	CS 4435 - CS 9624	45 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	46 / 62
Scheduling Theory and Implement	tation			Me	asuring Parallelism in Practice		
Space bounds				Plan			
opuce bounds				1 Iun			



The space S_p of a parallel execution on p processors required by Cilk++'s work-stealing satisfies:

$$S_p \le p \cdot S_1 \tag{13}$$

where S_1 is the minimal serial space requirement.

	< [৩৫৫		< □	→ 《圖》 《言》 《言》	≣
(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	47 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	48 / 62

Cilkview



- Cilkview computes work and span to derive upper bounds on parallel performance
- Cilkview also estimates scheduling overhead to compute a burdened span for lower bounds.

```
Measuring Parallelism in Practice
```

The Fibonacci Cilk++ example

Code fragment
long fib(int n)
{
if (n < 2) return n;
long x, y;
$x = cilk_spawn fib(n-1);$
y = fib(n-2);
cilk_sync;
return x + y;
}



Fibonacci program timing

The environment for benchmarking:

- model name : Intel(R) Core(TM)2 Quad CPU Q6600 @ 2.40GHz
- L2 cache size : 4096 KB
- memory size : 3 GB

	#cores = 1	#cores = 2		#cores	5 = 4
n	timing(s)	timing(s)	speedup	timing(s)	speedup
30	0.086	0.046	1.870	0.025	3.440
35	0.776	0.436	1.780	0.206	3.767
40	8.931	4.842	1.844	2.399	3.723
45	105.263	54.017	1.949	27.200	3.870
50	1165.000	665.115	1.752	340.638	3.420

Quicksort

code in cilk/examples/qsort

```
void sample_qsort(int * begin, int * end)
{
   if (begin != end) {
         --end;
        int * middle = std::partition(begin, end,
            std::bind2nd(std::less<int>(), *end));
        using std::swap;
        swap(*end, *middle);
        cilk_spawn sample_qsort(begin, middle);
        sample_qsort(++middle, ++end);
        cilk_sync;
    }
}
```

51 / 62

52 / 62

#cores = 2

speedup

1.927

1.923

1.936

1.971

timing(s)

1.016

5.469

11.096

57.996

#cores = 4

speedup

3.619

3.694

3.608

3.677

timing(s)

0.541

2.847

5.954

31.086

of int

 $10 imes10^{6}$

 $50 imes 10^6$

 $100 imes 10^{6}$

 $500 imes 10^{6}$

Matrix multiplication

Code in cilk/examples/matrix

Timing of multiplying a 687 \times 837 matrix by a 837 \times 1107 matrix

	iterative			r	ecursiv	e
threshold	st(s)	pt(s)	su	st(s)	pt (s)	su
10	1.273	1.165	0.721	1.674	0.399	4.195
16	1.270	1.787	0.711	1.408	0.349	4.034
32	1.280	1.757	0.729	1.223	0.308	3.971
48	1.258	1.760	0.715	1.164	0.293	3.973
64	1.258	1.798	0.700	1.159	0.291	3.983
80	1.252	1.773	0.706	1.267	0.320	3.959

	4) Q (4		 ٢ 	▶ ▲@ ▶ ▲ 코 ▶ ▲ 코 ▶ 콜	: nac
(Moreno Maza) Multithreaded Parallelism and Performance M CS 4435 - CS 9624 Measuring Parallelism in Practice	53 / 62	(Moreno Maza) Measuring F	Multithreaded Parallelism and Performance M Parallelism in Practice	CS 4435 - CS 9624	54 / 62
<pre>(Moreno Maza) Multithreaded Parallelism and Performance M Measuring Parallelism in Practice The cilkview example from the documentation Using cilk_for to perform operations over an array in parallel: static const int COUNT = 4; static const int ITERATION = 1000000; long arr[COUNT]; long do_work(long k){ long x = 15; static const int nn = 87; for (long i = 1; i < nn; ++i)</pre>	53 / 62	(Moreno Maza) Measuring F 1) Parallelism Prof Work : Span : Burdened span : Parallelism : Burdened paralle Number of spawns Average instruct Strands along sp Average instruct	Multithreaded Parallelism and Performance M Parallelism in Practice ile 6,480,8 2,116,8 31,920 lism : /syncs: ions / strand : an : ions / strand on span	CS 4435 - CS 9624 301,250 ins 301,250 ins ,801,250 ins 3.06 0.20 3,000,000 720 4,000,001 : 529	54 / 62
<pre>x = x / i + k % i; return x; } int cilk_main(){ for (int j = 0; j < ITERATION; j++) cilk_for (int i = 0; i < COUNT; i++) arr[i] += do_work(j * i + i + j);</pre>		<pre>2) Speedup Estimate 2 processors: 4 processors: 8 processors: 16 processors: 32 processors:</pre>	$\begin{array}{r} 0.21 - 2.00 \\ 0.15 - 3.06 \\ 0.13 - 3.06 \\ 0.13 - 3.06 \\ 0.12 - 3.06 \end{array}$		
} (Moreno Maza) Multithreaded Parallelism and Performance N CS 4435 - CS 9624	シュペ 55 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	▶ ◀ @ ▶ ◀ ≧ ▶ ◀ ≧ ▶ ■ CS 4435 - CS 9624	シ シ シ へ で 56 / 62

Timing for sorting an array of integers:

#cores = 1

timing(s)

1.958

10.518

21.481

114.300

Measuring Parallelism in Practice	Measuring Parallelism in Practice	
A simple fix	1) Parallelism Profile	
<pre>Inverting the two for loops int cilk_main() { cilk_for (int i = 0; i < COUNT; i++) for (int j = 0; j < ITERATION; j++) arr[i] += do_work(j * i + i + j); }</pre>	<pre>Work : 5,2 Span : 1,3 Burdened span : 1,3 Parallelism : Burdened parallelism : Number of spawns/syncs: Average instructions / strand : Strands along span : Average instructions / strand on sp 2) Speedup Estimate 2 processors: 1.40 - 2.00 4 processors: 1.76 - 3.99 8 processors: 2.01 - 3.99 16 processors: 2.17 - 3.99 32 processors: 2.25 - 3.99</pre>	95,801,529 ins 26,801,107 ins 26,830,911 ins 3.99 3.99 3 529,580,152 5 an: 265,360,221

· · · · · · · · · · · · · · · · · · ·				· · · · · · · · · · · · · · · · · · ·			
(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	57 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	58 / 62
Measuring Parallelism in Practice			Announcements				
<u> </u>							
Liming				Plan			
<u> </u>							

	#cores = 1	#core:	s = 2	#cores = 4		
version	timing(s)	timing(s)	speedup	timing(s)	speedup	
original	7.719	9.611	0.803	10.758	0.718	
improved	7.471	3.724	2.006	1.888	3.957	

1) Parallelis	m Compl	lexity N	leasures
---------------	---------	----------	----------

- 2 cilk_for Loops
- 3 Scheduling Theory and Implementation
- 4 Measuring Parallelism in Practice



Acknowledgements

References

- Charles E. Leiserson (MIT) for providing me with the sources of its lecture notes.
- Matteo Frigo (Intel) for supporting the work of my team with Cilk++.

Announcements

- Yuzhen Xie (UWO) for helping me with the images used in these slides.
- Liyun Li (UWO) for generating the experimental data.

- Matteo Frigo, Charles E. Leiserson, and Keith H. Randall. The Implementation of the Cilk-5 Multithreaded Language. Proceedings of the ACM SIGPLAN '98 Conference on Programming Language Design and Implementation, Pages: 212-223. June, 1998.
- Robert D. Blumofe, Christopher F. Joerg, Bradley C. Kuszmaul, Charles E. Leiserson, Keith H. Randall, and Yuli Zhou. Cilk: An Efficient Multithreaded Runtime System. Journal of Parallel and Distributed Computing, 55-69, August 25, 1996.
- Robert D. Blumofe and Charles E. Leiserson. Scheduling Multithreaded Computations by Work Stealing. Journal of the ACM, Vol. 46, No. 5, pp. 720-748. September 1999.

(日)					シイク 前一 本田 マ 本田 マ キャット 日 マント			
(Moreno Maza)	Multithreaded Parallelism and Performance N	CS 4435 - CS 9624	61 / 62	(Moreno Maza)	Multithreaded Parallelism and Performance M	CS 4435 - CS 9624	62 / 62	