

15-853: Algorithms in the Real World

Parallelism: Lecture 1

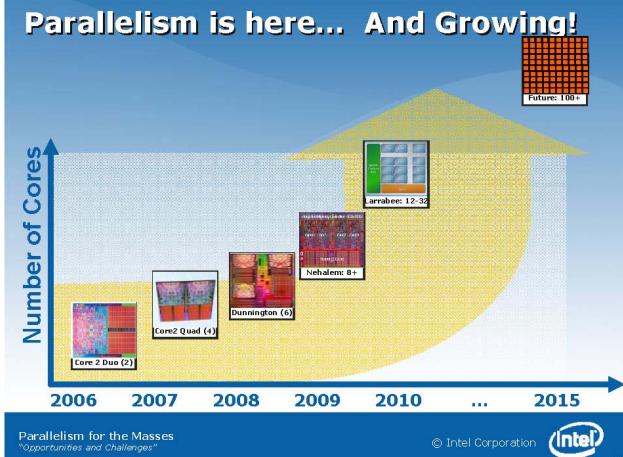
Nested parallelism

Cost model

Parallel techniques and algorithms

15-853

Page 1



15-853

Outline

Concurrency vs. Parallelism

Quicksort example

Nested Parallelism

- fork-join and parallel loops

Cost model: work and span

Techniques:

- Using collections: inverted index
- Divide-and-conquer: merging, mergesort, kd-trees, matrix multiply, matrix inversion, fft
- Contraction : quickselect, list ranking, graph connectivity, suffix arrays

15-853

Page 4

Parallelism in "Real world" Problems

Optimization
N-body problems
Finite element analysis
Graphics
JPEG/MPEG compression
Sequence alignment
Rijndael encryption
Signal processing
Machine learning
Data mining

15-853

Page5

Parallelism vs. Concurrency

- Parallelism: using multiple processors/cores running at the same time. Property of the machine
- Concurrency: non-determinacy due to interleaving threads. Property of the application.

		Concurrency	
		sequential	concurrent
Parallelism	serial	Traditional programming	Traditional OS
	parallel	Deterministic parallelism	General parallelism

15-853

6

Nested Parallelism

nested parallelism =
arbitrary nesting of parallel loops + fork-join

- Assumes no synchronization among parallel tasks except at joint points.
- Deterministic if no race conditions

Advantages:

- Good schedulers are known
- Easy to understand, debug, and analyze

15-853

Page7

Nested Parallelism: parallel loops

```
cilk_for (i=0; i < n; i++)      Cilk
    B[i] = A[i]+1;

Parallel.ForEach(A, x => x+1);  Microsoft TPL
                                (C#,F#)
```

```
B = {x + 1 : x in A}          Nesl, Parallel Haskell
```

```
#pragma omp for
for (i=0; i < n; i++)
    B[i] = A[i] + 1;          OpenMP
```

15-853

Page8

Nested Parallelism: fork-join

```
cobegin {  
  S1;  
  S2;}  
  
Dates back to the 70s or  
possibly 60s. Used in  
dialects of Pascal  
  
coinvoke(f1,f2) Java fork-join framework  
Parallel.invoke(f1,f2) Microsoft TPL (C#,F#)  
  
#pragma omp sections  
{  
  #pragma omp section OpenMP (C++, C, Fortran, ...)  
  S1;  
  #pragma omp section  
  S2;  
}  
15-853
```

Page9

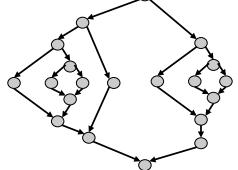
Nested Parallelism: fork-join

```
spawn S1;  
S2; cilk, cilk+  
sync;  
  
(exp1 || exp2) Various functional  
languages  
  
plet  
  x = exp1 Various dialects of  
  y = exp2 ML and Lisp  
in  
  exp3  
15-853
```

Page10

Serial Parallel DAGs

Dependence graphs of nested parallel computations are series parallel



Two tasks are parallel if not reachable from each other.
A data race occurs if two parallel tasks are involved in a race if they access the same location and at least one is a write.

15-853

Page11

Cost Model

Compositional:

Work : total number of operations

- costs are added across parallel calls

Span : depth/critical path of the computation

- Maximum span is taken across forked calls

Parallelism = Work/Span

- Approximately # of processors that can be effectively used.

15-853

Page12

Combining costs

Combining for parallel for:

```
pfor (i=0; i<n; i++)
  f(i);
```

$$W_{\text{pexp}}(\text{pfor } \dots) = \sum_{i=0}^{n-1} W_{\text{exp}}(f(i)) \quad \text{work}$$

$$D_{\text{pexp}}(\text{pfor } \dots) = \max_{i=0}^{n-1} D_{\text{exp}}(f(i)) \quad \text{span}$$

15-853

13

Why Work and Span

Simple measures that give us a good sense of efficiency (work) and scalability (span).

Can schedule in $O(W/P + D)$ time on P processors.

This is within a constant factor of optimal.

Goals in designing an algorithm

1. Work should be about the same as the sequential running time. When it matches asymptotically we say it is **work efficient**.
2. Parallelism (W/D) should be polynomial $O(n^{1/2})$ is probably good enough

15-853

14

Example: Quicksort

```
function quicksort(S) =
  if (#S <= 1) then S
  else let
    a = S[rand(#S)];
    S1 = {e in S | e < a};
    S2 = {e in S | e = a};
    S3 = {e in S | e > a};
    R = {quicksort(v) : v in [S1, S3]};
    in R[0] ++ S2 ++ R[1];
  in
```

Partition

Recursive calls

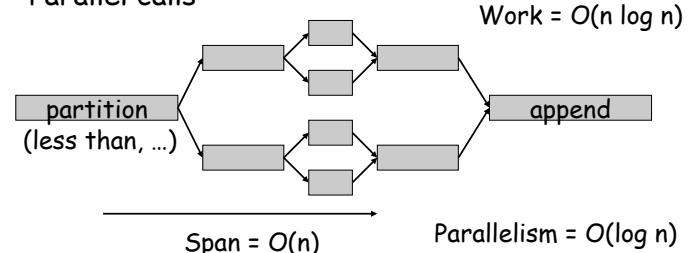
How much parallelism?

15-853

15

Quicksort Complexity

Sequential Partition and appending
Parallel calls



Not a very good parallel algorithm

15-853

*All randomized with high probability¹⁶

Quicksort Complexity

Now lets assume the partitioning and appending can be done with:

Work = $O(n)$

Span = $O(\log n)$

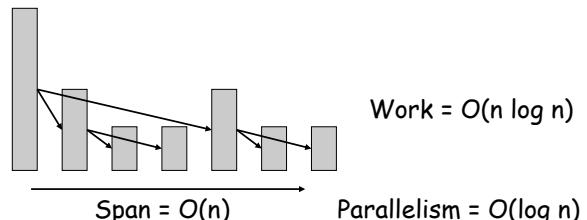
but recursive calls are made sequentially.

15-853

Page17

Quicksort Complexity

Parallel partition
Sequential calls



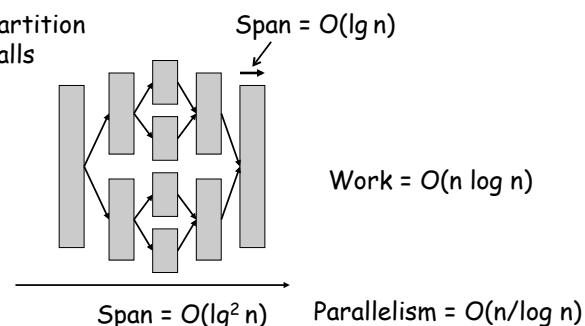
Not a very good parallel algorithm

15-853

*All randomized with high probability¹⁸

Quicksort Complexity

Parallel partition
Parallel calls



A good parallel algorithm

15-853

*All randomized with high probability¹⁹

Quicksort Complexity

Caveat: need to show that depth of recursion is $O(\log n)$ with high probability

15-853

Page20

Parallel selection

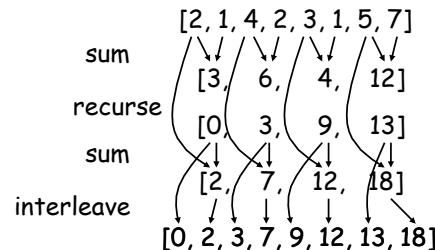
```
{e in S | e < a};  
  
S      = [2, 1, 4, 0, 3, 1, 5, 7]  
F = S < 4 = [1, 1, 0, 1, 1, 1, 0, 0]  
I = addscan(F) = [0, 1, 2, 2, 3, 4, 5, 5]  
  
where F  
R[I] = S = [2, 1, 0, 3, 1]
```

Each element gets sum of previous elements.
Seems sequential?

15-853

21

Scan



15-853

22

Scan code

```
function addscan(A) =  
if (#A <= 1) then [0]  
else let  
    sums = {A[2*i] + A[2*i+1] : i in [0:#a/2]};  
    evens = addscan(sums);  
    odds = {evens[i] + A[2*i] : i in [0:#a/2]};  
in interleave(evens, odds);  
  
W(n) = W(n/2) + O(n) = O(n)  
D(n) = D(n/2) + O(1) = O(log n)
```

15-853

23

Parallel Techniques

Some common themes in "Thinking Parallel"

1. Working with collections.
 - map, selection, reduce, scan, collect
2. Divide-and-conquer
 - Even more important than sequentially
 - Merging, matrix multiply, FFT, ...
3. Contraction
 - Solve single smaller problem
 - List ranking, graph contraction
4. Randomization
 - Symmetry breaking and random sampling

15-853

Page24

Working with Collections

reduce $\odot [a, b, c, d, \dots]$
= $a \odot b \odot c \odot d + \dots$

scan \odot ident $[a, b, c, d, \dots]$
= [ident, a, $a \odot b, a \odot b \odot c, \dots$

sort compF A

collect $[(2,a), (0,b), (2,c), (3,d), (0,e), (2,f)]$
= [(0, [b,e]), (2,[a,c,f]), (3,[d])]