

Parallelism Outline

Done:

- Intro to Parallelism vs. Concurrency
- Java threads
- How to use fork and join to write a parallel algorithm in Java
- Why using divide-and-conquer with lots of small tasks is best
 - Combines results in parallel

Now:

- More examples of simple parallel programs (map & reduce)
- Asymptotic analysis for fork-join parallelism
- Amdahl's Law
- Parallel Sorting

We looked at summing an array

 Summing an array went from O(n) sequential to O(log n) parallel (assuming a lot of processors and very large n)
 An exponential speed-up in theory



Extending Parallel Sum

- We can tweak the 'parallel sum' algorithm to do all kinds of things; just specify 2 parts (usually)
 - Describe how to compute the result for base problem size (Sum: Iterate through values sequentially and add them up)
 - Describe how to merge results
 - Describe how to merge results (Sum: Just add 'left' and 'right' results)





8

Reductions

- This class of computations are called reductions - We 'reduce' a large array of data to a single item
- Note: Recursive results don't have to be single numbers or strings. They can be arrays or objects with multiple fields. - Example: create a Histogram of test results from a much larger array of actual test results
- While many can be parallelized due to nice properties like associativity of addition, some things are inherently sequential How we process arr[i] may depend entirely on the result of processing arr[i-1]

Even easier: Data Parallel (Maps)

- · While reductions are a simple pattern of parallel programming, maps are even simpler
 - Operate on set of elements to produce a new set of elements (no combining results); generally input and output are of the same length
 - All operations can be done in parallel
 - E.g. Multiply each element of an array by 2.
- · Example:

int[] mult_by_two(int[] arr){ result = new int[arr.length]; FORALL(i=0; i < arr.length; i++) { result[i] = arr[i] * 2; }</pre> } return result; 3

Another Map Example:

Maps operate on a set of elements to produce a new set of elements (no combining results); generally input and output are of the same length

- All operations can be done in parallel

Another Example: Vector addition

int[] vector_add(int[] arr1, int[] arr2){
 assert (arr1.length == arr2.length);
 result = new int[arr1.length];
 FORALL(i=0; i < arr1.length; i++) {
 result[i] = arr1[i] + arr2[i];
 }
}</pre> return result;



(Java Details) Map vs reduce in ForkJoin framework

- In our examples:
- Reduce:
 - Parallel-sum extended RecursiveTask - Result was returned from compute()
- Map:

 - Class extended was RecursiveAction
 - Nothing returned from compute()
 - In the above code, the 'answer' array was passed in as a parameter
- · Doesn't have to be this way
 - Map can use RecursiveTask to, say, return an array
 - Reduce could use RecursiveAction; depending on what you're passing back via RecursiveTask, could store it as a class variable and access it via 'left' or 'right' when done

11



Analyzing algorithms

- Like all algorithms, parallel algorithms should be:
 Correct
 - Efficient
- For our algorithms so far, correctness is "obvious" so we'll focus on efficiency:
 - We still want asymptotic bounds
 - Want to analyze the algorithm without regard to a specific number of processors





13













- 1. Assume 256 processors
 - x-axis: sequential portion **S**, ranging from .01 to .25
 - y-axis: speedup T₁ / T_P (will go down as S increases)
- 2. Assume **S** = .01 or .1 or .25 (three separate lines)
 - x-axis: number of processors P, ranging from 2 to 32
 - y-axis: speedup T_1 / T_P (will go up as P increases)

I encourage you to try this out!

- Chance to use a spreadsheet or other graphing program
- Compare against your intuition
- A picture is worth 1000 words, especially if you made it

22





26

Sequential Quicksort review

Recall quicksort was sequential, in-place, expected time O(n log n)

	Best / expected case work
1. Pick a pivot element	O(1)
2. Partition all the data into:	O(n)
A. The elements less than the pivot	
B. The pivot	
C. The elements greater than the pivot	
3. Recursively sort A and C	2T(n/2)
Recurrence (assuming a good pivot): T(0)=T(1)=1 T(n)=n + 2T(n/2) = O(nlogn)	
Run-time: O(nlogn)	
How should we parallelize this?	
	25

Review: Common recurrences

T(n) = O(1) + T(n-1)	linear (e.g. recursive linear search)
T(n) = O(1) + T(n/2)	logarithmic (e.g. binary search)
T(n) = O(n) + T(n-1)	quadratic (e.g. Quicksort worst case)
T(n) = O(n) + T(n/2)	linear (see next few slides)
T(n) = O(n) + 2T(n/2)	O(n log n) (e.g. Mergesort)

