

# CSE 332: Parallel Sorting

Richard Anderson, Steve Seitz  
Winter 2014

# Announcements

- Project 3 PartA due Thursday night

# Recap

## Last week

- simple parallel programs
- common patterns: map, reduce
- analysis tools (work, span, parallelism)
- Amdahl's Law

## Now

- parallel quicksort, merge sort
- useful building blocks: prefix, pack

# Parallelizable?

Fibonacci (N)

$$\begin{aligned} \text{Fib}(N) &= \text{fib}(N-1) + \text{fib}(N-2) \\ &= 1 \quad \text{if } N \in \{1, 2\} \end{aligned}$$

Not really\* because of sequential dependencies

naive approach spawns  $2^n$  threads  
→ exponential time

# Parallelizable?

Prefix-sum:

input	6	3	11	10	8	2	7	8
output	6	9	20	30	38	40	47	55

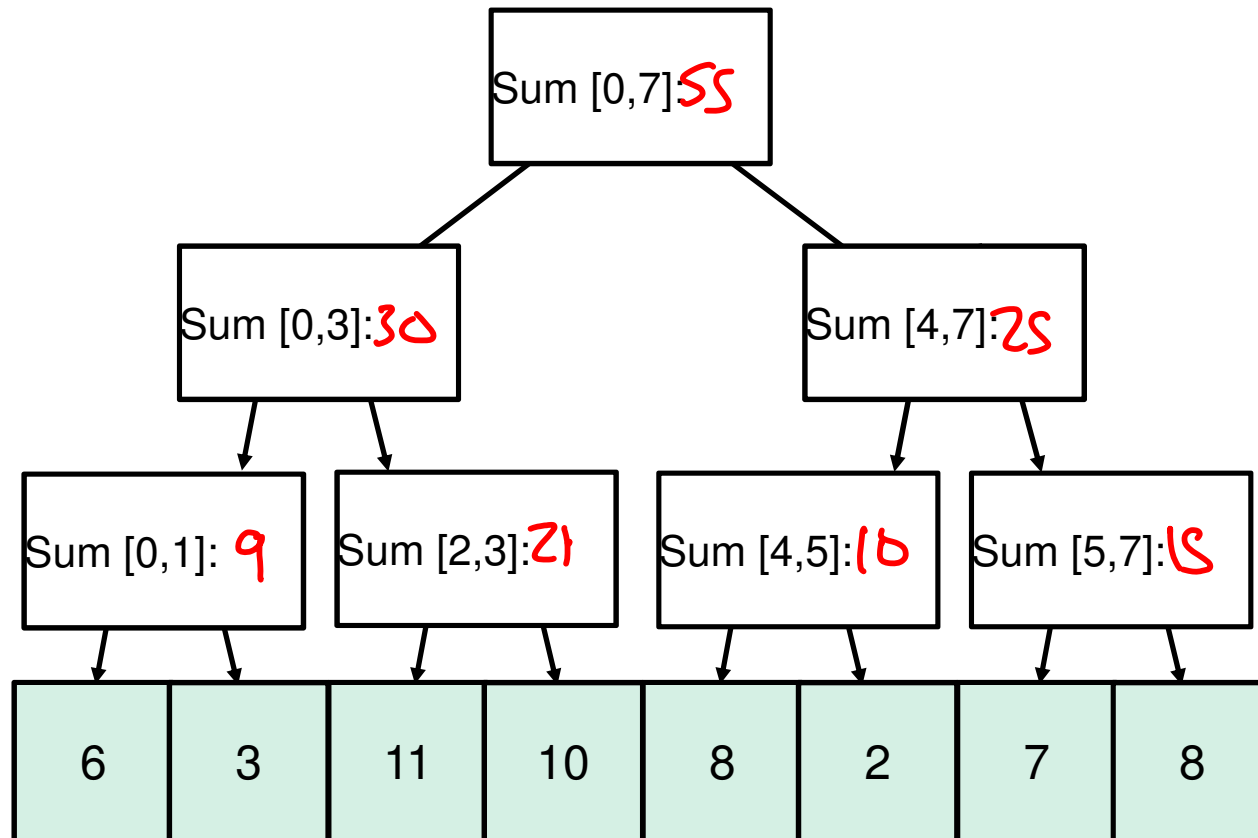
$$output[i] = \sum_0^{i-1} input[i]$$

# First Pass: Sum

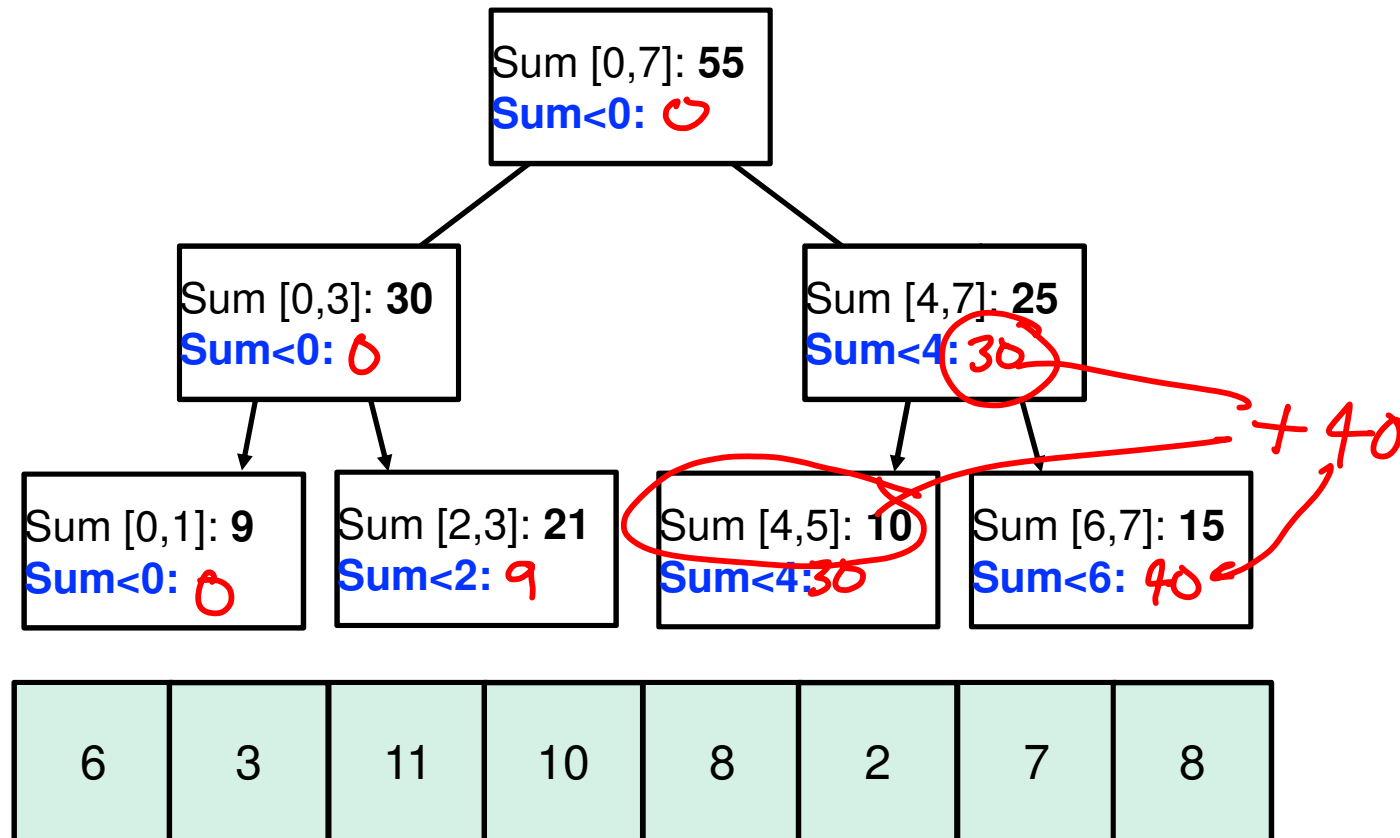
Sum [0,7]: 55

6	3	11	10	8	2	7	8
---	---	----	----	---	---	---	---

# First Pass: Sum

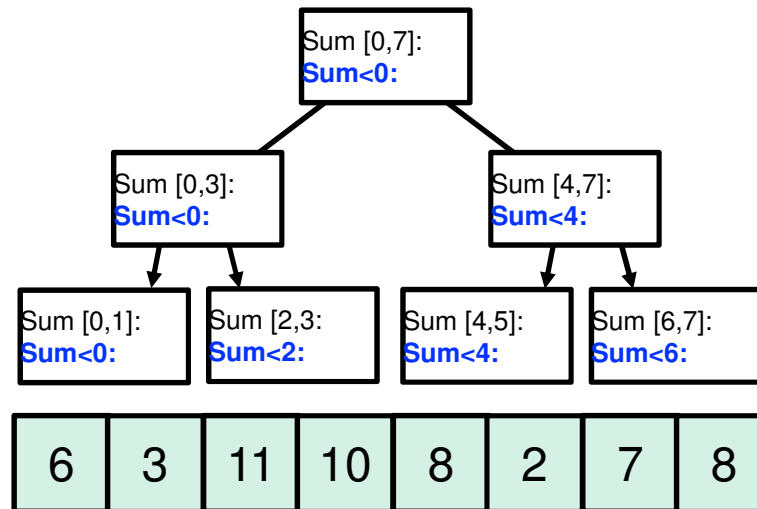


# 2nd Pass: Use Sum for Prefix-Sum





# 2nd Pass: Use Sum for Prefix-Sum



Go from root down to leaves

Root

–  $\text{sum} < 0 = 0$

Left-child

–  $\text{sum} < K = \text{parent's sum} < K$

Right-child

–  $\text{sum} < K = \text{parent's sum} < K + \text{siblings sum}[J, K]$

# Prefix-Sum Analysis

- First Pass (Sum):
  - span =  $O(\log n)$
- Second Pass:
  - single pass from root down to leaves
    - update children's sum <K value based on parent and sibling
  - span =  $O(\log n)$
- Total
  - span =  $O(\log n)$

# Parallel Prefix, Generalized

Prefix-sum is another common pattern (prefix problems)

- maximum element to the left of  $i$
- is there an element to the left of  $i$  satisfying some property?
- count of elements to the left of  $i$  satisfying some property
- ...

We can solve all of these problems in the same way

# Pack

Pack:

input	6	3	11	10	8	2	7	8	test: $x < 8?$
output	6	3	2	7					

Output array of elements satisfying **test**, in original order

# Parallel Pack?

Pack

input	6	3	11	10	8	2	7	8	test: $x < 8?$
output	6	3	2	7					

- Determining **which** elements to include is **easy**
- Determining **where** each element goes in output is **hard**
  - seems to depend on previous results

# Parallel Pack

1. map test input, output [0,1] bit vector

<b>input</b>	6	3	11	10	8	2	7	8	<b>test: <math>x &lt; 8?</math></b>
<b>test</b>	1	1	0	0	0	1	1	0	

# Parallel Pack

1. map test input, output [0,1] bit vector

<b>input</b>	6	3	11	10	8	2	7	8	<b>test: x &lt; 8?</b>
<b>test</b>	1	1	0	0	0	1	1	0	

2. transform bit vector into array of indices into result array

<b>pos</b>	1	2	2	2	2	3	4	4	
------------	---	---	---	---	---	---	---	---	--

# Parallel Pack

1. map test input, output [0,1] bit vector

<b>input</b>	6	3	11	10	8	2	7	8	<b>test: <math>x &lt; 8?</math></b>
<b>test</b>	1	1	0	0	0	1	1	0	

2. prefix-sum on bit vector

<b>pos</b>	1	2	2	2	2	3	4	4
------------	---	---	---	---	---	---	---	---

3. map input to corresponding positions in output

<b>output</b>	6	3	2	7				
---------------	---	---	---	---	--	--	--	--

- `if (test[i] == 1) output[pos[i]] = input[i]`



# Parallel Pack Analysis

- Parallel Pack
  1. map:  $O(\quad)$  span
  2. sum-prefix:  $O(\quad)$  span
  3. map:  $O(\quad)$  span
- Total:  $O(\quad)$  span

# Sequential Quicksort

## Quicksort (review):

1. Pick a pivot  $O(1)$
2. Partition into two sub-arrays  $O(n)$ 
  - A. values less than pivot
  - B. values greater than pivot
3. Recursively sort A and B  $2T(n/2)$ , avg

## Complexity (avg case)

- $T(n) = n + 2T(n/2)$        $T(0) = T(1) = 1$
- $O(n \log n)$

How to parallelize?

# Parallel Quicksort

## Quicksort

1. Pick a pivot O(1)
2. Partition into two sub-arrays O(n)
  - A. values less than pivot
  - B. values greater than pivot
3. Recursively sort A and B **in parallel** **T(n/2), avg**

## Complexity (avg case)

- $T(n) = n + T(n/2)$        $T(0) = T(1) = 1$
- **Span:  $O(\log n)$**
- **Parallelism (work/span) =  $O(n)$**

# Taking it to the next level...

- $O(\log n)$  speed-up with infinite processors is okay, but a bit underwhelming
  - Sort  $10^9$  elements 30x faster
- Bottleneck:

# Parallel Partition

Partition into sub-arrays

- A. values less than pivot
- B. values greater than pivot

What parallel operation can we use for this?

# Parallel Partition

- Pick pivot

8	1	4	9	0	3	5	2	7	6
---	---	---	---	---	---	---	---	---	---

- Pack (test:  $<6$ )

1	4	0	3	5	2				
---	---	---	---	---	---	--	--	--	--

- Right pack (test:  $\geq 6$ )

1	4	0	3	5	2	6	8	9	7
---	---	---	---	---	---	---	---	---	---

└──────────┘

# Parallel Quicksort

## Quicksort

1. Pick a pivot  $O(1)$
2. Partition into two sub-arrays  $O(\quad)$  span
  - A. values less than pivot
  - B. values greater than pivot
3. Recursively sort A and B in parallel  $T(n/2)$ , avg

## Complexity (avg case)

- $T(n) = O(\quad) + T(n/2)$        $T(0) = T(1) = 1$
- Span:  $O(\quad)$
- Parallelism (work/span) =  $O(\quad)$

# Sequential Mergesort

Mergesort (review):

- |                               |           |
|-------------------------------|-----------|
| 1. Sort left and right halves | $2T(n/2)$ |
| 2. Merge results              | $O(n)$    |

Complexity (worst case)

- $T(n) = n + 2T(n/2)$        $T(0) = T(1) = 1$
- $O(n \log n)$

How to parallelize?

- Do left + right in parallel, improves to  $O(n)$
- To do better, we need to...




# Parallel Merge




How to merge two sorted lists in parallel?

# Parallel Merge



1. Choose median  $M$  of left half  $O(\quad)$
2. Split both arrays into  $< M, \geq M$   $O(\quad)$ 
  - how?


# Parallel Merge




1. Choose median  $M$  of left half
2. Split both arrays into  $< M, \geq M$ 
  - how?
3. Do two submerges in parallel



*merge*



*merge*




*merge*



*merge*





*merge*



*merge*



*merge*






*merge*




*merge*




When we do each merge in parallel:

- ✦ we split the bigger array in half
- ✦ use binary search to split the smaller array
- ✦ And in base case we copy to the output array

*merge*



*merge*



# Parallel Mergesort Pseudocode

```
Merge(arr[], left1, left2, right1, right2, out[], out1, out2 )
    int leftSize = left2 - left1
    int rightSize = right2 - right1
    // Assert: out2 - out1 = leftSize + rightSize
    // We will assume leftSize > rightSize without loss of generality

    if (leftSize + rightSize < CUTOFF)
        sequential merge and copy into out[out1..out2]

    int mid = (left2 - left1)/2
    binarySearch arr[right1..right2] to find j such that
        arr[j] ≤ arr[mid] ≤ arr[j+1]

    Merge(arr[], left1, mid, right1, j, out[], out1, out1+mid+j)
    Merge(arr[], mid+1, left2, j+1, right2, out[], out1+mid+j+1, out2)
```

# Analysis

## Parallel Merge (worst case)

- Height of partition call tree with  $n$  elements:  $O(\log n)$
- Complexity of each thread (ignoring recursive call):  $O(n)$
- Span:  $O(n)$

## Parallel Mergesort (worst case)

- Span:  $O(n)$
- Parallelism (work / span):  $O(\log n)$

## Subtlety: uneven splits



- but even in worst case, get a 3/4 to 1/4 split
  - still gives  $O(\log n)$  height

# Parallel Quicksort vs. Mergesort

## Parallelism (work / span)

- quicksort:  $O(n / \log n)$       avg case
- mergesort:  $O(n / \log^2 n)$       worst case