CSE 332 Data Structures & Parallelism

Parallel Prefix Sum

Melissa Winstanley Spring 2024

The prefix-sum problem

Given int[] input, produce int[] output where:



- Work: O(n), Span: O(n)
- This algorithm is sequential, but a different algorithm has Work: O(n), Span:
 O(log n

The prefix-sum problem

Sequential can be an intro to CS exam problem



int[] prefix_sum(int[] input){
 int[] output = new int[input.length];
 output[0] = input[0];
 for(int i=1; i < input.length; i++)
 output[i] = output[i-1]+input[i];
 return output;</pre>

Parallel prefix-sum

- The parallel-prefix algorithm does two passes
 - Each pass has O(n) work and O(log n) span
 - So in total there is O(n) work and O(log n) span
 - So like with array summing, parallelism is n/log n
 - An exponential speedup

- First pass builds a tree bottom-up: the "up" pass
- Second pass traverses the tree top-down: the "down" pass



Parallel Prefix: The Up Pass

We build want to build a binary tree where

- Root has sum of the range [x,y)
- If a node has sum of [lo,hi) and hi>lo,
 - Left child has sum of [lo,middle)
 - Right child has sum of [middle,hi)
 - A leaf has sum of [i,i+1), which is simply input[i]

It is critical that we actually <u>create the tree</u> as we will need it for the down pass

- We do not need an actual linked structure
- We could use an array as we did with heaps

Analysis of first step: Work = Span =







Parallel prefix, generalized

Just as sum-array was the simplest example of a common pattern, prefix-sum illustrates a pattern that arises in many, many problems

- Minimum, maximum of all elements 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
 This last one is perfect for an efficient parallel pack...
 - Perfect for building on top of the "parallel prefix trick"

Pack (think "Filter")

[Non-standard terminology]

Given an array input, produce an array output containing only elements such that **f(element)** is true

Example: input [17, 4, 6, 8, 11, 5, 13, 19, 0, 24] f: "is element 10" output [17, 11, 13, 19, 24]

Parallelizable?

- Determining *whether* an element belongs in the output is easy
- But determining <u>where</u> an element belongs in the output is hard; seems to depend on previous results....

Parallel Pack = (Soln) parallel map + parallel prefix + parallel map

1. **Parallel map** to compute a **bit-vector** for true elements:

input [17, 4, 6, 8, 11, 5, 3 9 0, 24]
bits 1, 0, 0, 0, 1 0, 1 0, 1 0, 1
2. Parallel-prefix sum on the bit vector:
 bitsum [1, 1, 1, 1, 2, 2, 3, 4, 4, 5]
3. Parallel map to produce the output:
 output [17, 11, 13, 19, 24]

In this example, Filter = element > 10

Parallel Pack = (Soln) parallel map + parallel prefix + parallel map

1. **Parallel map** to compute a **bit-vector** for true elements:

input [17, 4, 6, 8, 11, 5, 13, 19, 0, 24]
bits [1, 0, 0, 0, 1, 0, 1, 1, 0, 1]

2. Parallel-prefix sum <u>on the bit-vector</u>:

bitsum [1, 1, 1, 1, 2, 2, 3, 4, 4, 5]

3. **Parallel map** to produce the output:

output [17, 11, 13, 19, 24]

output = new array of size bitsum[n-1]
FORALL(i=0; i < input.length; i++) {
 if(bits[i]==')
 OUtput[bitsum[i]-1]=input[i]
}</pre>

In this example, Filter = element > 10

Parallel Pack = (Soln) parallel map + parallel prefix + parallel map

1. **Parallel map** to compute a **bit-vector** for true elements:

In this example, Filter = element > 10



Sequential Quicksort review

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

Best / expected case work

- 1. Pick a pivot element
- 2. Partition all the data into:
 - 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 $T(n)=\frac{2T(n/2)+1}{2T(n/2)+1}$

Run-time: O(nlogn) How should we parallelize this?



Parallel Quicksort (version 1)

Best / expected case work

- Pick a pivot element 1.
- 2. Partition all the data into:
 - A. The elements less than the pivot
 - The pivot Β.
 - C. The elements greater than the pivot
- 3. **Recursively sort A and C**

First: Do the two recursive calls in parallel

- Work: O(nlogn)Span: $T(n) = O(n) + T(n/2) \longrightarrow O(n)$

O(1) O(n)



Review: Really common recurrences

Should know how to solve recurrences but also recognize some really common ones:



Note big-Oh can also use more than one variable

• Example: can sum all elements of an n-by-m matrix in O(nm)

Pack (think "Filter")

[Non-standard terminology]

Given an array input, produce an array output containing only elements such that f(element) is true

Example: input [17, 4, 6, 8, 11, 5, 13, 19, 0, 24]

f: "is element > 10"

output [17, 11, 13, 19, 24]

Parallelizable?

- Determining *whether* an element belongs in the output is easy
- But determining <u>where</u> an element belongs in the output is hard; seems to depend on previous results....

Parallel partition (not in place)

Partition all the data into:

- A. The elements less than the pivot
- B. The pivot
- C. The elements greater than the pivot

This is just two packs!

- We know a pack is O(n) work, O(log n) span
- Pack elements less than pivot into left side of aux array
- Pack elements greater than pivot into right size of aux array
- Put pivot between them and recursively sort
- With a little cleverness, can do both packs at once (!) but no big-O change

With O(10Gn) span for partition, the total span for quicksort is T(n) = O(10Gn) + T(n/2)

Parallel Quicksort Example (version 2)

Step 1: pick pivot as median of three

Steps 2a and 2c (combinable): pack less than pivot, then pack greater than pivot into a second array

 Fancy parallel prefix to pull this off (not shown)

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

Step 3: Two recursive sorts in parallel

- Can sort back into original array (like in mergesort)

Parallelize Mergesort?

Recall mergesort: sequential, **not**-in-place, worst-case O(n log n)

- 1. Sort left half and right half
- 2. Merge results



Just like quicksort, doing the two recursive sorts in parallel changes the recurrence for the **Span** to T(n) = O(n) + 1T(n/2) = O(n)

- Again, **Work** is O(nlogn), and
- parallelism is **work/span** = O(log n)
- To do better, need to parallelize the merge
 - The trick won't use parallel prefix this time...



Idea: Suppose the larger subarray has m elements. In parallel;

- Merge the first m/2 elements of the larger half with the "appropriate" elements of the smaller half
- Merge the second m/2 elements of the larger half with the rest of the smaller half

Parallelizing the merge (in more detail)

Need to merge two **sorted** subarrays (may not have the same size) **Idea**: Recursively divide subarrays in half, merge halves in parallel

Suppose the larger subarray has *m* elements. In parallel:

- Pick the median element of the larger array (here 6) in constant time
- In the other array, use binary search to find the first element greater than or equal to that median (here 7)

Then, in parallel:

- Merge half the larger array (from the median onward) with the upper part of the shorter array
- Merge the lower part of the larger array with the lower part of the shorter array

Example: Parallelizing the merge

Example: Parallelizing the merge 0 4 6 8 9 1 2 3 5 7

1. Get median of bigger half: O(1) to compute middle index

Example: Parallelizing the merge



- 1. Get median of bigger half O(1) to compute middle index
- Find how to split the smaller half at the same value:
 O(log n) to do binary search on the sorted small half



- 1. Get median of bigger half: O(1) to compute middle index
- Find how to split the smaller half at the same value:
 O(log n) to do binary search on the sorted small half
- 3. Size of two sub-merges conceptually splits output array: O(1)

Example: Parallelizing the merge



- 1. Get median of bigger half: O(1) to compute middle index
- Find how to split the smaller half at the same value:
 O(log n) to do binary search on the sorted small half
- 3. Size of two sub-merges conceptually splits output array: O(1)
- 4. Do two submerges in parallel

Parallel Merge Pseudocode



Merge(arr[), left, left, right, right, out], out, out) int leftSize = $left_2 - left_1$ 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 = $(left_2 - left_1)/2$ binarySearch arr[right,..right,] to find j such that $arr[j] \le arr[mid] \le arr[j+1]$ Merge(arr[], left₁, mid, right₁, j, out₁, out₁, out₁+mid+j) Merge(arr[], mid+1, left₂, j+1, right2, out[], out₁+mid+j+1, out₂)

left1

Analysis

• Doing the two recursive calls in parallel but a sequential merge:

Work: same as sequential $O(n \log n)$ Span: $T(n) \neq 1$ (n/2) + O(n) which is O(n)

- Parallel merge makes work and span harder to compute...
 - Each merge step does an extra O(log n) binary search to find how to split the smaller subarray
 - To merge *n* elements total, do two smaller merges of possibly different sizes
 - But worst-case split is (3/4)*n* and (1/4)*n*
 - Happens when the two subarrays are of the same size (n/2) and the "smaller" subarray splits into two pieces of the most uneven sizes possible: one of size n/2, one of size 0

Analysis continued

For just a parallel merge of n elements:

- which is O(n) **Work** is T(n) = (3n/4) + T(n/4) $O(\log n)$ which is $O(\log^2 n)$
- **Span** is $T(n) = T(3n/4) + O(\log n)$,
- (neither bound is immediately obvious, but "trust me")

So for **mergesort** with parallel merge overall:

- **Work** is T(n) = 2T(n/2) + O(n), which is $O(n \log n)$
- **Span** is $T(n) = (T(n/2) + O(\log^2 n))$, which is $O(\log^3 n)$

So parallelism (work / span) is $O(n / \log^2 n)$

- Not quite as good as quicksort's $O(n / \log n)$
 - But (unlike Quicksort) this is a worst-case guarantee
- And as always this is just the asymptotic result

Summary

Quicksort (best case)

 $O(n \log n)$ - sequential n^2 worst case

--> O(n) span - parallel calls to quicksort

--> O(log² n) span - parallel partition

Mergesort (worst case)

O(n log n) - sequential

--> O(n) span - parallel calls to mergesort

--> O(log³ n) span - parallel merge

Toward sharing resources (memory)

So far, we have been studying parallel algorithms using the fork-join model

- Reduce span via parallel tasks

Fork-Join algorithms all had a very simple *structure* to avoid race conditions

- Each thread had memory "only it accessed"
 - Example: each array sub-range accessed by only one thread
- Result of forked process not accessed until after join() called
- So the structure (mostly) ensured that bad simultaneous access wouldn't occur

Strategy won't work well when:

- Memory accessed by threads is overlapping or unpredictable
- Threads are doing **independent tasks** needing **access to same resources** (rather than implementing the same algorithm)

