

Algorithms Are Algorithms, Some With More Parallelism than Others

Guy Blelloch
Carnegie Mellon University

Context:

Intro algorithms + data structure course at Carnegie Mellon (15-210)

- Teach parallelism throughout
- Sophomore level
- Started 12 years ago
- Taught around 5000 students, about half are CS majors
- Consists of both theory and programming.
- All but one assignment has parallelism
- Jointly developed with Umut Acar and others. Eight faculty have taught it.

Context: our curriculum

Students accepted directly into CS. 75% have had significant programming

0th year

- Intro to programming

1st year

- Intro to imperative programming (assumes AP CS)
- Discrete math for CS
- Intro to functional programming (introduce parallelism)

Context: deterministic parallelism

We teach “deterministic” parallelism, i.e.:

- no race conditions
- no locks
- no concurrency issues

Hence we use a mostly “functional”, i.e., side effect free, style for describing and programming algorithms.

Philosophy

There is no fundamental difference between parallel and sequential algorithms

- Need to recognize when parallelism is already present
- Need to understand how to make solutions more parallel
- Difference is only in costs and some added techniques

Emphasis on “**Parallel Thinking**”: i.e. start by thinking about solutions as parallel, or techniques that allow for parallelism

Pros and Cons

Pros:

- Do not have to diverge much from standard algorithms course structure
- Emphasizes that parallelism is not esoteric
- Learn to think about parallelism abstractly and naturally

Cons:

- Requires revising an algorithms course throughout (significant work for faculty). Many faculty not fully comfortable with the topic.

Why are sequential algorithms so successful?

- abstract away from detail
- broadly applicable
- general techniques
- simple cost model
- easy to program
- elegant and often simple
- help understand scaling, i.e., big-O
- lead to good efficiency in practice
- interesting theory


Lead to good efficiency?

Natural question: do ideas taught in course help in practice

speedup on 32 cores

Mean = 29.8

32 core/64 hyperthread single chip
intel processor (Ice Lake)



Outline of the talk:

Go through a sample of the topics covered in the course

For each point out how different from standard treatment (i.e. how does parallel thinking show up)

Would love feedback, especially on how to make it more accessible to other Universities.

Course Outline

- Cost models and analysis
- Techniques: D&C, brute force, ...
- Algorithms on sequences: merging, mergesort, scan, ..
- Randomization: quicksort, ...
- BSTs and Priority Queues
- Graph algorithms (search, shortest paths, MSTs, ...)
- Dynamic programming

Relatively traditional at this level of detail

Parallel Model

$(e1 \parallel e2)$: evaluate $e1$ and $e2$ in parallel, returning a pair when both done

Sequence operations: e.g.

- Map : $[f(x) : x \text{ in } S]$
- Tabulate : $[f(i) : 0 \leq i < n]$
- Filter : $[x \text{ in } S \mid f(x)]$

All calls to f are in parallel, returning a sequence

Sidebar: what could be added in cs1

Python:

Replace

for x in S: f(x) \Rightarrow [f(x) for x in S]

r = 0

For x in S: \Rightarrow reduce(add,S)

 r = r + x

Reduce available in functools

Point out that second can be parallel if f(x) does not have a side effect.

Intellectual load:

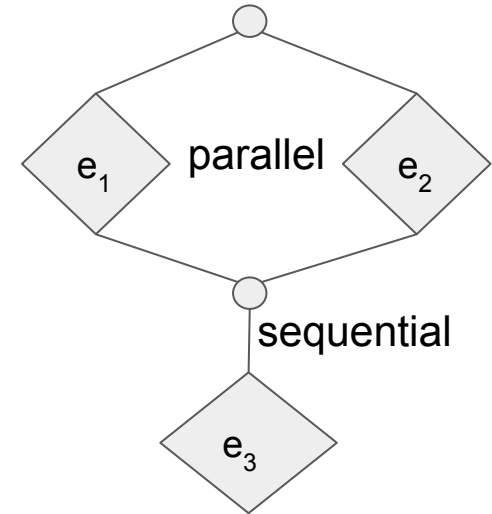
- What is a side effect
- How it is safe to run in parallel if no effect
- Passing functions

Cost Model : Work and Span

Work: Total number of operations (like sequential time)

Span: Critical path (assuming unbounded number of processors)

	Sequential composition	Parallel composition
Work	add	add
Span	add	max



No PRAM

Cost Model : work and span

$$e = (x,y) = (e_1 \parallel e_2); \\ e_3;$$

$$W(e) = W(e_1) + W(e_2) + W(e_3) + 2$$

$$S(e) = \max(S(e_1), S(e_2)) + S(e_3) + 2$$

	Sequential composition	Parallel composition
Work	add	add
Span	add	max

Cost model

Priorities:

- Most important is to reduce the work
- Then to reduce the span

Parallelism = Work/Span (roughly # of processors can make use of)

Some discussion of scheduling, leading to:

$\text{Time} = \text{Work}/\text{Processors} + \text{Span}$

Recurrences and Big-O

Same as sequential algorithms but for work and span instead of time.

E.g. for mergesort:

- $W(n) = 2W(n/2) + O(n) = O(n \log n)$ using sum over recursive calls
- $S(n) = S(n/2) + O(\log n) = O(\log^2 n)$ using max over recursive calls

Techniques

- Brute force
- Divide-and-conquer
- Contraction
- Graph search
- Dynamic programming

Techniques

- Brute force (note that naturally parallel)
- Divide-and-conquer (naturally parallel, and use more aggressively)
- Contraction (not usually covered in sequential algorithms)
- Graph search (BFS is parallel but not DFS)
- Dynamic programming (view more abstractly as a DAG)

Divide and Conquer

Same as with sequential, but:

- Recurrences for span
- More emphasis (use where iteration is used sequentially), e.g., reduce, parenthesis matching, MCSS
- Often need to strengthen induction

Divide and Conquer: reduce

Taking the sum with respect to a binary associative function f

reduce f identity $S =$

Case splitMid(x) of

Empty => return identity

| Single(x) => return x

| Pair(L, R) => return $f(\text{reduce } f \text{ identity } L \parallel \text{reduce } f \text{ identity } R)$

$$W(n) = 2W(n/2) + W_f = O(n) \quad \text{assuming } W_f = S_f = O(1)$$

$$S(n) = S(n/2) + S_f = O(\log n)$$

Would normally do iteratively in sequential setting.

Divide and Conquer: mergeSort

Taking the sum with respect to a binary associative function f

mergesort $S =$

 case splitMid(x) of

 Empty => return []

 | Single(x) => S

 | Pair(L, R) => return merge(mergesort L || mergesort R)

$$W(n) = 2W(n/2) + O(n) = O(n \log n)$$

$$S(n) = S(n/2) + \log n = O(\log^2 n)$$


Assuming for merge: $W(n) = O(n)$, $S(n) = O(\log n)$

Divide and Conquer: mergeSort

Or:

Mergesort $S = \text{reduce}(\text{merge}, [], [[x] : x \text{ in } S])$

Divide and Conquer: Merging



$\text{merge}(A_L, B_L) \parallel \text{merge}(A_R, B_R)$

At leaf or recursion place values in correct position

$$W(n) = 2W(n/2) + O(\log n) = O(n)$$

$$S(n) = S(n/2) + O(\log n) = O(\log^2 n)$$

Very different from sequential algorithm.

Can get $S(n) = O(\log n)$ with more advanced approach.

Divide and Conquer: Parenthesis Matching

Do these parentheses match: $((()()()()))((()()))$

Strengthen the induction to return number of unmatched parentheses on left and right:

parenMatch' S =

Case split_mid(x) of

Empty => return (0,0)

| Single(x) => if x = '(' then return (0,1) else return (1,0)

| Pair(L,R) =>

((la,lb), (ra,rb)) = (parenMatch' L, parenMatch' R)

if (lb > ra) then return (la, lb-ra + rb)

else return (la + ra - lb, rb)

parenMatch S =

(a,b) = ParenMatch' S

return (a = 0) and (b = 0)

Technique: Contraction

Basic structure:

- Reduce problem to one that is a constant factor smaller
- Recurse on the problem
- Use result to solve the problem

Some applications: Scan, kthSmallest, Graph Contraction

Not usually covered in intro algorithms/data structure course, but important for parallelism.

Contraction Example: Scan

For binary associative operator f , return “running sum” of previous elements

e.g.: Scan + 0 [2, 1, 3, 1, 2] \rightarrow [0, 2, 3, 6, 7], 9

Seems inherently sequential

Parallel version has many applications:

E.g.: filter, flatten, carry propagation, skyline, ...

Scan: using contraction

scan(f, identity, S) =

 If $|A| = 0$ then return ($[],$ identity)

 else if $|A| = 1$ then return ($[$ identity $], S[0]$)

 else

 pairSums = $[f(A[2i], A[2i+1]) : 0 \leq i < |S|/2]$

 // pairwise sum the elements

 (r,s) = scan(f, identity, pairSums)

 // recurse

 return $[$ if $(i\%2 = 0)$ then $r[2i]$ else $r[2i] + a[i]$ $]$

 // put together results

$$W(n) = W(n/2) + O(n) = O(n)$$

$$S(n) = S(n/2) + O(1) = O(\log n)$$

Assuming f is constant work:

Randomized algorithms

Analyze:

- expectation for work (e.g. $E[W_{\text{quicksort}}(n)] = O(n \log n)$)
- high probability bounds for span (e.g. $S[W_{\text{quicksort}}(n)] = O(\log^2 n)$ w.h.p.)

Problem is there is no equivalent of linearity of expectation for maximum.

Example of students taking an exam.

Randomized algorithms: kthSmallest

kthSmallest(A, k) =

 p = pick an element of A uniformly at random

 L = [x in A | x < p]

 R = [x in A | x > p]

 if (k < |L|) then return kthSmallest(L, k)

 else if (k < |A| - |R|) then return p

 else return kthSmallest(R, k - (|A| - |R|))

Show that number of recursive calls is $O(\log n)$ with high probability

Also work is $O(n)$ in expectation.

Randomized algorithms: Quicksort

```
quicksort(A) =  
  p = pick an element of A uniformly at random  
  L = [x in A | x < p]  
  R = [x in A | x > p]  
  (L', R') = (quicksort(L) || quicksort(R))  
  return L ++ [p] ++ R
```

Work = $O(n \log n)$ in expectation

Span = $O(\log^2 n)$ with high probability

Binary search trees:

Traditionally teach insertion, deletion, find on sets or dictionaries

None of these have parallelism.

Extend to support:

- Filter, reduce, ..
- Union, intersection, ..

Allows for bulk operations on sets and dictionaries (Parallel Thinking)

Can all be built up from “join”, and are highly parallel

Binary search trees: Filter example

filter f A =

case A of


Leaf => A

| Node(L,k,R) =>

(L', R') = (filter f L, filter f R)

if (f(k)) return join(L',k,R')

else return joinPair(L', R')



Work = $O(n)$

Span = $O(\log^2 n)$

Join handles the rebalancing.

We use treaps for balancing, but could use AVL trees or Red Black trees.




Graphs:

30% of course is on graphs

Some key points:

- BFS is parallel, span is $O(d \log n)$, where d is the diameter of the graph
- DFS and Dijkstra are sequential
- Bellman Ford is parallel: $O(mn)$ work and $O(n \log n)$ span.
- Johnson is parallel
- Use graph contraction for connectivity and Min Spanning Tree (MST)
MST uses Boruvka's algorithm, but we also mention Kruskal and Prim

Graphs: Example of star contraction



Graphs: Star contraction

In one round remove at least $\frac{1}{4}$ of the vertices in expectation.

Therefore finishes in $O(\log n)$ rounds with high probability.

Total work is $O(m \log n)$

Used for both Connectivity and Boruvkas MST algorithm

Dynamic Programming:

Some examples we cover:

- Subset sum: $W(n,k) = O(nk)$, $S(n, k) = O(n)$
- Minimum edit distance: $W(n,m) = O(nm)$, $S(n,m) = O(n + m)$
- Optimal BST: $W(n) = O(n^3)$, $S(n) = O(n \log n)$
- Matrix chain Product: $W(n) = O(n^3)$, $S(n) = O(n \log n)$

Discuss bottom up and top down approaches.

Lessons Learned

- **Deterministic** parallelism not difficult for the students to understand.
 - They get the hang quickly.
- List of “new” ideas that need to be covered, and which are easier/harder
- resistance by faculty to changing existing courses
 - faculty often not comfortable with topic
- best to do along with considering whole curriculum

Have not yet done a careful study

Lessons Learned: effect on whole curriculum

- ideally requires changes in prerequisites
 - parallel thinking earlier
 - programming with comprehensions and functional style earlier
 - tail bounds in probability
- changes to follow up courses
 - how to cover parallelism in Senior level algorithms class?
 - where to cover concurrent and distributed algorithms?
- important to have complementary course on systems issues
 - gpus, distributed computing, locks, ...

Conclusions

Parallel and sequential algorithms can be integrated such that students

- consider parallelism throughout
- learn “standard” algorithmic topics, e.g. D&C, DP, mergesort, quicksort, BFS, DFS, ...
- learn to think parallel, e.g., ask if an algorithm has parallelism, or if and how it can be added

Including in an existing UG curriculum has challenges