

15-150 Fall 2025

Lecture 19

Parallelism

Cost Semantics and Sequences

today

parallel programming

- parallelism and functional style
- cost semantics
- Brent's Theorem and speed-ups
- sequences: an abstract type with *efficient parallel operations*

parallelism

exploiting *multiple processors*

evaluating *independent code simultaneously*

- low-level implementation
 - *scheduling* work onto processors, tell each processor to do at each time step
- high-level planning
 - *designing code abstractly*
 - *without baking in a schedule*

our approach

Deal with scheduling implicitly

- Programmer specifies what to do
- Compiler determines how to schedule the work
- **Parallelism is deterministic**

Our thesis: this approach to parallelism will *prevail*..

(and 15-210 builds on these ideas...)

functional benefits

- No side effects, so...
evaluation order doesn't affect *correctness*
- Can build *abstract types* that support efficient *parallel-friendly* operations
- Can use *work* and *span* to predict potential for *parallel speed-up*
 - Work and span are *independent* of scheduling details

caveat

- In practice, it's hard to achieve speed-up
- Current language implementations don't make it easy
- Problems include:
 - scheduling overhead
 - locality of data (cache problems)
 - runtime sensitive to scheduling choices

what can programmers do?

- Lists bake in sequential evaluation. Trees don't.
- Today, we introduce sequences that have a linear structure like lists but offer parallelism of trees.
- Reason about time complexity using work and span

cost semantics

We already introduced *work* and *span*

- *Work* estimates the *sequential* running time on a *single* processor
- *Span* takes account of data dependency, estimates the *parallel* running time with *unlimited* processors

cost semantics

- We showed how to calculate *work* and *span* for *recursive functions* with **recurrence relations**
- Now we introduce **cost graphs**, another tool to deal with work and span
- Cost graphs also allow us to talk about *schedules*...
- ... and the potential for *speed-up*

cost graphs


A *cost graph* is a ***series-parallel graph***

- a *directed* acyclic graph, with *source* and *sink* (constant time)
- branching indicates *potential parallelism*

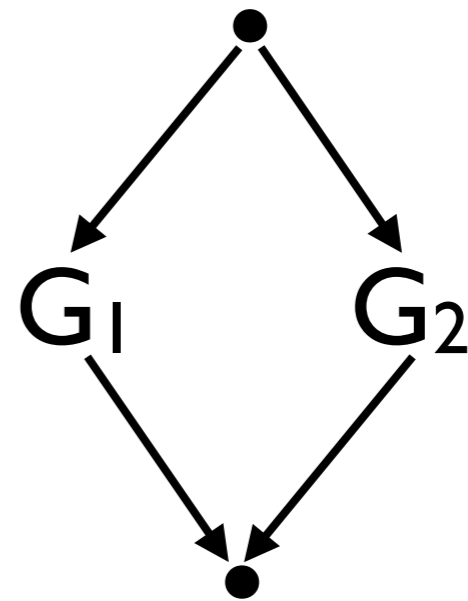
series-parallel graphs



a single node



sequential
composition




parallel
composition

(n-ary parallelism allowed)

example

$$(1+2) * 3$$



(Edges are implicitly directed downward)


work and span

of a cost graph

- The **work** is the *number of nodes*
- The **span** is the *length of the longest path from source to sink*


$$\mathbf{span}(G) \leq \mathbf{work}(G)$$

span



$$= \text{span } G_1 + \text{span } G_2 + c$$

sequential code ... add the span



$$= \max(\text{span } G_1, \text{span } G_2) + c$$


parallel code ... max the span

sources and sinks

- Sometimes we omit them from pictures
- No loss of generality
 - easy to put them in
- No difference, asymptotically
 - a single node represents an additive constant amount of work and span
- Allows easier explanation of *execution*

example

$$(1+2) * 3$$



work = 7

span = 5



Brent's Theorem

An expression with work w and span s can be evaluated on a p -processor machine in time $\Omega(\max(w/p, s))$.


scheduling

- p pebbles, with p the number of processors
- Start with one pebble on cost graph G 's source
- Putting a pebble on a node visits the node
- At each time step, pick up all pebbles and put at most p on the graph, no more than one per node. Can only put a pebble on an unvisited node all of whose ancestors have been visited.

This could be a cost graph for $(1+2) * (3+4)$




This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2		
	3		
	4		
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2		
	3		
	4		
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3		
	4		
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3		
	4		
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4		
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4		
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4	h	i
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4	h	i
	5		
	6		

This could be a cost graph for $(1+2) * (3+4)$




		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4	h	i
	5	e	j
	6		

This could be a cost graph for $(1+2) * (3+4)$



		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4	h	i
	5	e	j
	6		


This could be a cost graph for $(1+2) * (3+4)$



		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4	h	i
	5	e	j
	6	(idle)	f

work = 10

span = 5



		processors	
		1	2
time	1	a	(idle)
	2	b	g
	3	c	d
	4	h	i
	5	e	j
	6	(idle)	f

next

- Exploiting parallelism in ML
- A signature for *parallel collections*
- *Cost analysis* of implementations
- *Cost benefits* of parallel algorithm design

sequences

signature SEQ =

```
sig  
  type 'a seq (* abstract *)  
  exception Range of string  
  val empty : unit ->'a seq  
  val tabulate : (int -> 'a) -> int -> 'a seq  
  val length : 'a seq -> int  
  val nth : 'a seq -> int -> 'a  
  val map : ('a -> 'b) -> 'a seq ->'b seq  
  val reduce : ('a * 'a -> 'a) -> 'a -> 'a seq -> 'a  
  val mapreduce : ('a -> 'b) -> 'b -> ('b * 'b -> 'b) -> 'a seq -> 'b  
  val filter: ('a -> bool) -> 'a seq -> 'a seq  
end
```

implementations

- Many ways to implement the signature
 - lists, balanced trees, arrays, ...
- For each one, can give a *cost analysis*
- There may be implementation *trade-offs*
 - *arrays: item access is $O(1)$*
 - *trees: item access is $O(\log n)$*

Seq :SEQ

- An abstract parameterized type of *sequences*
- Think of a sequence as a *parallel collection*
- With *parallel-friendly* operations
 - *constant-time* access to items
 - *efficient* map and reduce

sequence values

A value of type `t seq`
is a sequence of values of type `t`

- We use math notation like

$\langle v_0, \dots, v_{n-1} \rangle$

$\langle \rangle$

for sequence values

$\langle 1, 2, 4, 8 \rangle$ is a value of type `int seq`

Reminder:
A client would
write `t Seq.seq`

equivalence

- Two sequence values are *extensionally equivalent* iff they have the same length and have extensionally equivalent items at all

$$\langle v_0, \dots, v_{n-1} \rangle \cong \langle u_0, \dots, u_{m-1} \rangle$$

if and only if

$$n \cong m \text{ and for all } i, v_i \cong u_i$$

operations

For our given structure $\text{Seq} : \text{SEQ}$, we specify

- the (extensional) *behavior*
- the *cost semantics*

of each operation

Other implementations of SEQ may achieve *different* work and span profiles

Learn to choose wisely!

`empty ()` returns $\langle \rangle$


- Type can be `t seq` for any type `t`
- Cost graph



work and span $O(1)$

$\text{tabulate } f \ n \cong \langle f \ 0, \dots, f(n-1) \rangle$

- If G_i is cost graph for $f(i)$,
the cost graph for $\text{tabulate } f \ n$ is



If f is $O(1)$, the work for $\text{tabulate } f \ n$ is $O(n)$

If f is $O(1)$, the span for $\text{tabulate } f \ n$ is $O(1)$

`tabulate f n` \cong $\langle f\ 0, \dots, f(n-1) \rangle$

examples

- `tabulate (fn x:int => x) 6` $\langle 0, 1, 2, 3, 4, 5 \rangle$
- `tabulate (fn x:int => x*x) 6` $\langle 0, 1, 4, 9, 16, 25 \rangle$

`nth` $\langle v_0, \dots, v_{n-1} \rangle$ $i \cong v_i$ if $0 \leq i < n$
 \cong **raise** Range otherwise

- Work is $O(1)$
- Span is $O(1)$
- Cost graph is



Contrast: `List.nth`
work, span $O(n)$

$$\text{length } \langle v_0, \dots, v_{n-1} \rangle \cong n$$

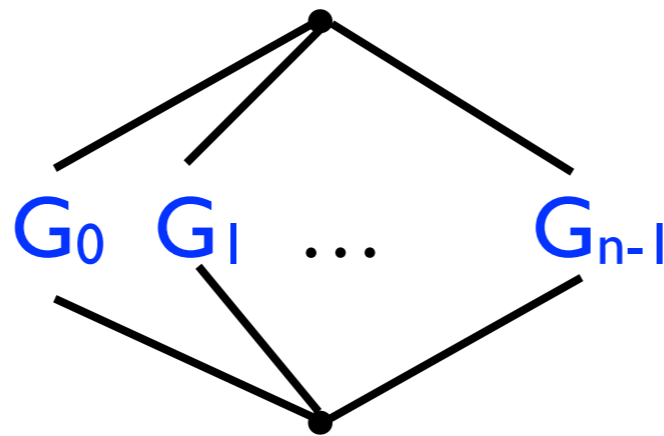
- Work is $O(1)$
- Span is $O(1)$
- Cost graph is



Contrast: $\text{List.length } [v_0, \dots, v_{n-1}] \cong n$
work, span $O(n)$

$$\text{map } f \langle v_0, \dots, v_{n-1} \rangle \cong \langle f v_0, \dots, f v_{n-1} \rangle$$

$\text{map } f \langle v_0, \dots, v_{n-1} \rangle$ has cost graph



where each G_i
is cost graph for $f v_i$

- If f is constant time, $\text{map } f \langle v_0, \dots, v_{n-1} \rangle$ has *work* $O(n)$, *span* $O(1)$

(contrast with `List.map`)

reduce

reduce is used to *combine a sequence*

reduce : ('a * 'a -> 'a) -> 'a ->'a seq -> 'a

Compare it with

foldr: ('a * 'b -> 'a) -> 'b -> 'a list -> 'b

reduce

$$('a * 'a -> 'a) -> 'a -> 'a \text{ seq } -> 'a$$

$$\text{reduce } g \ z \ \langle v_0, \dots, v_{n-1} \rangle \cong v_0 \odot v_1 \dots \odot v_{n-1} \odot z$$

where g is an *associative* function with a base value z
where we represent g with the infix operator \odot

- $g : t * t -> t$ is **associative** iff for all $x_1, x_2, x_3 : t$


$$g(x_1, g(x_2, x_3)) = g(g(x_1, x_2), x_3)$$

- Sometimes we will assume that z is an *identity element* for g , i.e. for all $x : t$, $g(x, z) = x$

$$\text{reduce } g \ z \ \langle v_0, \dots, v_{n-1} \rangle \cong v_0 \odot v_1 \dots \odot v_{n-1}$$

$$\text{reduce } g \ z \ \langle \rangle \cong z$$

$$\text{reduce } g \ z \ \langle v_0, \dots, v_{n-1} \rangle \cong v_0 \odot v_1 \dots \odot v_{n-1} \odot z$$



assuming g is $O(1)$

work is $O(n)$

span is $O(\log n)$

mapreduce $f \ z \ g \ \langle v_0, \dots, v_n \rangle \cong (f \ v_0) \odot \dots \odot (f \ v_{n-1}) \odot z$

assuming f and g are $O(1)$

has *work* $O(n)$

and *span* $O(\log n)$

filter p $s \cong s'$

with S' a sequence consisting of all x_i in S such that $p(x)$ true for all x_i in S . The order of retained elements in S' is the same as in S

Assuming p is $O(1)$, has *work* $O(n)$

and *span* $O(\log n)$

```
mapreduce f z g ⟨v1, ..., vn⟩ = (f v1) g ... g (f vn) g z
```

Example: filter

```
val singleton : 'a -> 'a seq (* gives a single element  
                             sequence *)
```

```
val append : 'a seq * 'a seq -> 'a seq
```

```
fun filter (p: 'a -> bool) : 'a seq -> 'a seq =  
  let val nothing = empty ()  
      fun keep x = if p (x) then singleton x  
                  else nothing  
  in  
    mapreduce keep nothing append  
  end
```

$S(n) = O(\log n)$, $W(n) = O(n \log n)$ assuming append has span $O(1)$

Example: count

using map

```
fun sum (s : int Seq.seq) : int = _____
```

```
type row = int Seq.seq
```

```
type room = row Seq.seq
```

```
fun count (class: room) : int = sum _____
```

Example: count

using map

```
fun sum (s : int Seq.seq) : int = Seq.reduce (op +) 0 s
```

```
type row = int Seq.seq
```

```
type room = row Seq.seq
```

```
fun count (class: room) : int = sum _____
```

Example: count

using map

```
fun sum (s : int Seq.seq) : int = Seq.reduce (op +) 0 s
```

```
type row = int Seq.seq
```

```
type room = row Seq.seq
```

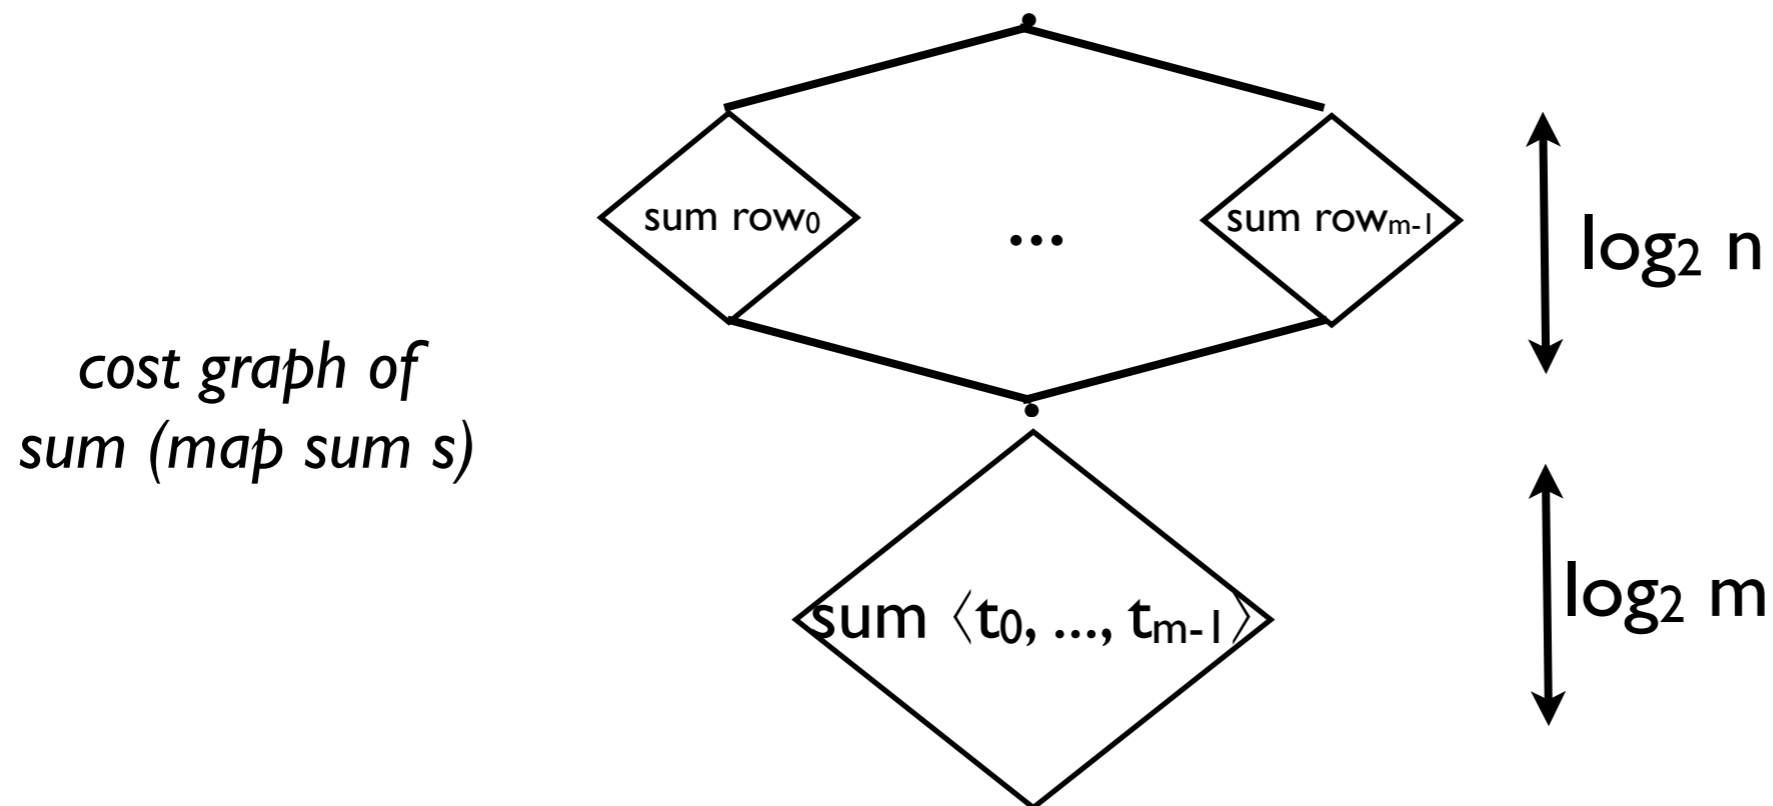
```
fun count (class: room) : int = sum (Seq.map sum class)
```

analysis

Let $t_i = \text{sum row}_i$

m rows of length n each

count $s = \text{sum} \langle t_0, \dots, t_{m-1} \rangle$



work is $O(mn)$
span is $O(\log n + \log m)$

```
mapreduce f z g ⟨v1, ..., vn⟩ = (f v1) g ... g (f vn) g z
```

Alternatively

using mapreduce

```
fun sum (s : int Seq.seq) : int = Seq.reduce (op +) 0 s
```

```
type row = int Seq.seq
```

```
type room = row Seq.seq
```

```
fun count (class: room) : int =
```

```
    Seq.mapreduce sum 0 (op +) class
```