

1/31/2023 Model

- PRAM
- Vector RAM (VRAM)
- Circuit Models (NC, P)
- M_P-RAM : Multiprocessor RAM

Goals for a parallel model:

- simple
 - guide you (alg. designer) toward efficient algorithms
 - understand how perf / alg. scales as we vary the input size
 - robust across a variety of machines
 - useful for understanding alg. design techniques
- ☐ can naturally express algs. in pseudocode and by extension real code.

What about the PRAM?

- perhaps not robust / implementable
- but alg. design techniques can be illustrated in the PRAM




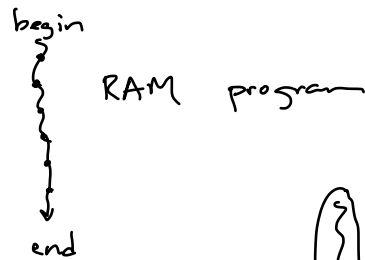
- PRAM assumes - fixed (P) processors, making it rather unwieldy for coding


Work-Depth / Work-Span

- shared random access memory
 - dynamic task creation / closure.
- work = # operations
 depth/span = longest set of sequential dependencies



MP-RAM : Multiprocessor-RAM

- set of dynamically evolving processes
- unbounded memory






- augment with FORK (k)
 - k children processes created
 - child i gets "id" i
- nested forks OK!




Binary-forking vs k-ary forking

work-stealing scheduler (~~Bin~~, Blumofe-Leiserson)



work = total number of basic operations

depth = longest chain of sequential dependencies



$$W = 10$$

$$D = 5$$

Work = time using 1 processor

Depth = time using ∞ processors.

$$T \geq \frac{W}{P} \Rightarrow T \geq \max\left(\frac{W}{P}, D\right)$$

$$T \geq D$$

Brent's Thm: Given a DAG with work W , Depth D , using P processors it can be scheduled in $O\left(\frac{W}{P} + D\right)$ steps

$$\max\left(\frac{W}{P}, D\right) \leq T \leq \underbrace{O\left(\frac{W}{P} + D\right)}_{\text{work-stealing}}$$

$\frac{W}{P} + D$ time

$P = \frac{W}{D}$

parallelism

$D + D$

- Parallelism = $\frac{W}{D}$

- $W = O(n \log n)$


$\frac{W}{D} = O(n)$

- $D = O(\log n)$

parallel for loop: $[0, 1, 2, \dots, |A|-1]$

perform i in $[0 : |A|]$

$B[i] \leftarrow f(A[i])$



par do, parallel do, fork: ||

$sum(A) =$ $+$ $A[0 : |A|/2]$

if ($|A| = 1$) return $A[0]$

$l \leftarrow sum(A[0, |A|/2])$ ||

$r \leftarrow sum(A[|A|/2, |A|])$

return $l + r$

$$C = A \parallel B$$

$$w(C) = w(A) + w(B)$$

$$D(C) = \max(D(A), D(B)) + 1$$

Sum:

$$w(n) = 2w(n/2) + O(1) = O(n)$$

$$D(n) = \max(D(n/2), D(n/2)) + O(1) \in O(\log_2 n)$$


Filter (A, f):

[1, 2, 3, ..., 10]

$$w = O(n)$$

↓
[2, 4, 6, 8, 10]

$$D = O(\log n)$$



$$D = \text{alloc}(|L| + |R|)$$

Filter (A, f):

if (|A| = 1):

 if f(A[0]) return A

 else return []

endif

(L, R) = filter(A₁^f) || filter(A₂^f)

o = alloc(|L| + |R|)


{ copy L, R into o }

return o

endfun

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

$$\in O(\log^2 n)$$



Nick's class:

- poly work
- poly-logarithmic depth

SSSP: $\frac{n^3}{P}$ work $\log^2 n$ depth
 \downarrow
 $n \log(n)$ work/depth

work-efficiency: parallel alg uses the same amount of work (asymptotically) as that of the best seq. algorithm.

$$\boxed{\frac{W}{P}} + \underline{\underline{O(D)}}$$

