

8E380

Exam 2
today
@
SPM

Mont
104

Input:

$$X = k_1, k_2, \dots, k_n; \oplus$$

Output:

$$k_1, k_1 \oplus k_2, k_1 \oplus k_2 \oplus k_3, \dots, k_1 \oplus k_2 \oplus \dots \oplus k_n$$

$$S = n$$

LEMMA: We can solve the above

problem in $O(\log n)$ Time

Using n CREW PRAM Processors.

PROOF:

$$\left. \begin{array}{l} \text{let } P = n \\ X_1 = k_1, k_2, \dots, k_{\frac{n}{2}} \text{ and} \\ X_2 = k_{\frac{n}{2}+1}, k_{\frac{n}{2}+2}, \dots, k_n \end{array} \right\}$$

Step 0: Assign k_i to processor i ,
 $1 \leq i \leq n$;

Step 1: Using $\frac{n}{2}$ proc.
perform a prefix comp.

Recursively on X_1 to
get $k'_1, k'_2, \dots, k'_{\frac{n}{2}}$;

Using the other $\frac{n}{2}$ proc.
perform a prefix comp.

Recursively on X_2 to get

$k'_{\frac{n}{2}+1}, k'_{\frac{n}{2}+2}, \dots, k'_n$;

Step 2: Output the FIRST half without Change.
Using $\frac{n}{2}$ processors, 'preadd' $k'_{\frac{n}{2}}$ to
each element of $k'_{\frac{n}{2}+1}, k'_{\frac{n}{2}+2}, \dots, k'_n$
and output these.

Let $T(n)$ be the RUN TIME of this ALG.
ON ANY INPUT OF SIZE n , USING
 n PROCESSORS.

Then,

$$T(n) = T\left(\frac{n}{2}\right) + O(1)$$
$$= O(\log n).$$

$$a=1, b=2,$$

$$f(n) = \Theta(1).$$
$$n^{\log_b a}$$

$$= n^{\log_2 1} = 1.$$

Case 2 applies.

$$\Rightarrow T(n) = \Theta(\log n).$$

$$\text{WORK DONE} = PT$$
$$= O(n \log n).$$

LEMMA: We can solve the prefix comp. problem in $O(\log n)$ time using $\frac{n}{\log n}$ CREW PRAM PROC.

BRENT'S TRICK.

Reduce the input size;
Run a Non optimal alg.
on the reduced input;
Use these answers to
derive answers for the
original input.

$$P = \left\lceil \frac{n}{\log n} \right\rceil$$

o) Assign $K_{(i-1)\log n+1}, K_{(i-1)\log n+2}, \dots, K_{i\log n}$ to processor i , for $1 \leq i \leq P$.

① for $1 \leq i \leq P$ in parallel do

Processor i performs a prefix Comp on its $\log n$ elements to get:

$$K'_{(i-1)\log n+1}, K'_{(i-1)\log n+2}, \dots, K'_{i\log n}$$

(2) All the P processors perform a prefix

Comp. on

$K'_{1 \log n}, K'_{2 \log n}, \dots, K'_n$ to get

$K''_{1 \log n}, K''_{2 \log n}, \dots, K''_n$

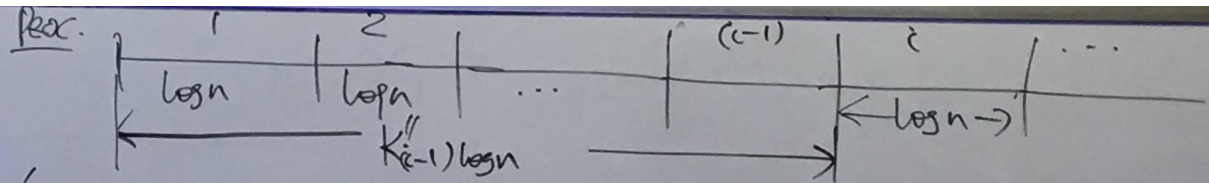
$X = \boxed{K_1, K_2, \dots, K_{1 \log n}}, \dots, K_{2 \log n}, \dots, K_{3 \log n}, \dots, K_n$

$K''_{i \log n} = \text{Sum of the FIRST}$

$i \log n$ elements of X , for $1 \leq i \leq P$.

(3) For $1 \leq i \leq p$ in Π^R do

Processor i outputs $K_{(i-1)\log n}'' \oplus K_{(i-1)\log n+1}', K_{(i-1)\log n}'' \oplus K_{(i-1)\log n+2}', \dots, K_{(i-1)\log n}'' \oplus K_{i\log n}'$



Step 1 takes $O(\log n)$ TIME.

Step 2 takes $O\left(\log\left(\frac{n}{\log n}\right)\right)$ TIME
 $= O(\log n)$.

Step 3 takes $O(\log n)$ TIME.

Total RUN TIME = $O(\log n)$. \square

PROBLEM:

INPUT: $X = k_1, k_2, \dots, k_n$
and another road # k .

Goal: Permute X s.t.
all the elements
less than k appear
before the rest.

$$K=17.$$

Ex. $X = 5, 11, 8, 12, 25, 32, 3, 6, 15;$

Output: $5, 11, 8, 12, 3, 6, 15, 25, 32$

Fact. We can solve this in $O(\log n)$ time using $\frac{n}{\log n}$ processors.

Algorithm:

(1) Generate an array $A[1:n]$ of BITS
 s.t. $A[i] = \begin{cases} 1 & \text{if } K_i < K \\ 0 & \text{otherwise} \end{cases}, 1 \leq i \leq n.$

A	1	1	1	1	0	0	1	1	1
---	---	---	---	---	---	---	---	---	---

(2) Perform a prefix SUMS Comp.
 on $A[1], A[2], \dots, A[n].$

1	2	3	4	4	4	5	6	7
---	---	---	---	---	---	---	---	---

These sums can be used as addresses for writing elements of X that are $< K$.

③ Use a linear procedure to write elements of $X \geq k$.

TOTAL RUN TIME
 $= O(\log n)$.

PROBLEM:

INPUT: $X = k_1, k_2, \dots, k_n; k$

Output: $\text{Rank}(k, X) = \left| \{x \in X: x < k\} \right| + 1;$

This can be done in $O(\log n)$ TIME
USING $\frac{n}{\log n}$ PROC.

$X = 3, 8, 21, 17, 6, 5, 12, 34, 16; k = 12$

① Generate an array $A[1..n]$ s.t.

$$A[i] = \begin{cases} 1 & \text{if } k_i < k \\ 0 & \text{otherwise} \end{cases}$$

1	1	0	0	1	1	0	0	0
---	---	---	---	---	---	---	---	---

② Perform a prefix Sums Comp. on $A[1], A[2], \dots, A[n]$; let the results be $B[1], B[2], \dots, B[n]$.

③ Output $B[n] + 1$

total TIME = $O(\log n)$.

SORTING. $X = k_1, k_2, \dots, k_n$

① For $1 \leq i \leq n$ in \parallel^k do

Using $\frac{n}{\log n}$ proc. Compute $y_i = \text{Rank}(k_i, X);$

(2) For $1 \leq i \leq n$ in \mathbb{R} do
Processor i outputs
 k_i in cell r_i ;

\Rightarrow We can sort n elements
in $O(\log n)$ TIME USING
 $\frac{n^2}{\log n}$ CREW PRAM PROC.
WORK DONE $= O(n^2)$.

CSE 3500 Algorithms and Complexity – Fall 2016

Lecture 23: November 15, 2016

Parallel Algorithms: Prefix Computation

- **Input:** A sequence $X = k_1, k_2, \dots, k_n$ of elements from a domain Σ . \oplus is a binary, associative, and unit operation defined on Σ . Recall that an operation \oplus on Σ is associative if for any three elements x, y, z in Σ , the following holds: $x \oplus (y \oplus z) = (x \oplus y) \oplus z = x \oplus y \oplus z$.
- **Output:** $k_1, k_1 \oplus k_2, k_1 \oplus k_2 \oplus k_3, \dots, k_1 \oplus k_2 \oplus \dots \oplus k_n$.

A Divide-and-conquer Algorithm for Prefix Computation

- We will first discuss a divide-and-conquer parallel algorithm that employs n CREW PRAM processors.
- Step 0) We partition the input into two equal halves: $X_1 = k_1, k_2, \dots, k_{n/2}$ and $X_2 = k_{(n/2)+1}, k_{(n/2)+2}, \dots, k_n$.
- Step 1) $\frac{n}{2}$ processors recursively perform a prefix computation on X_1 . Let the output be $k'_1, k'_2, \dots, k'_{n/2}$; At the same time the other $\frac{n}{2}$ processors recursively perform a prefix computation on X_2 . Let the output be $k'_{(n/2)+1}, k'_{(n/2)+2}, \dots, k'_n$.
- Note that $k'_1, k'_2, \dots, k'_{n/2}$ is indeed the first half of the prefix outputs for X . Thus we can output these without any modifications.
- Step 3) We can modify $k'_{(n/2)+1}, k'_{(n/2)+2}, \dots, k'_n$ by pre-adding $k'_{n/2}$ to every element. (Here the word ‘adding’ refers to the operator \oplus). The modified values will be the second half of the prefix outputs for X . This modification can be done in $O(1)$ time using $\frac{n}{2}$ CREW PRAM processors.

Run time analysis: Like for any recursive algorithm, we have to write a recurrence relation for the run time, and solve it. Let $T(n)$ be the run time of the above algorithm on any input of size n , where the number of processors used is n .

Then we get the following recurrence relation: $T(n) = T(n/2) + O(1)$ which solves to $T(n) = O(\log n)$. As a result, we get the following Lemma.

Lemma 1: *We can solve the prefix computation problem on any input of size n in $O(\log n)$ time using n CREW PRAM processors. \square*

An Optimal Prefix Computation Algorithm

- We can get an optimal prefix computation algorithm using a technique due to Richard Brent. The idea is to reduce the input size sufficiently, employ a nonoptimal algorithm to solve the problem on the reduced input, and use these results to obtain the results for the original input. Let $P = \frac{n}{\log n}$ and let $X = k_1, k_2, \dots, k_n$ be the input. A detailed description is given below.

- 0) Assign $\log n$ elements per processor. Specifically, assign the elements $k_{(i-1)\log n+1}, k_{(i-1)\log n+2}, \dots, k_{i\log n}$ to processor i , for $1 \leq i \leq P$;
- 1) **for** $i = 1$ **to** P **in parallel do**
- 2) Processor i performs a prefix computation on its $\log n$ elements $k_{(i-1)\log n+1}, k_{(i-1)\log n+2}, \dots, k_{i\log n}$ to get $k'_{(i-1)\log n+1}, k'_{(i-1)\log n+2}, \dots, k'_{i\log n}$;
- 3) P processors collectively perform a prefix computation on $k'_{\log n}, k'_{2\log n}, \dots, k'_n$ to get $k''_{\log n}, k''_{2\log n}, \dots, k''_n$;
- 4) **for** $i = 2$ **to** n **in parallel do**
- 5) Processor i outputs $k''_{(i-1)\log n} \oplus k'_{(i-1)\log n+1}, k''_{(i-1)\log n} \oplus k'_{(i-1)\log n+2}, \dots, k''_{(i-1)\log n} \oplus k'_{i\log n}$;

Run time analysis: Step 2 takes $O(\log n)$ time. In step 3 we have to perform a prefix computation on $\frac{n}{\log n}$ elements using $\frac{n}{\log n}$ processors. Using Lemma 1, we infer that Step 3 takes $O\left(\log\left(\frac{n}{\log n}\right)\right) = O(\log n)$ time. Step 5 takes $O(\log n)$ time. Thus the total run time of the algorithm is $O(\log n)$ resulting in the following Lemma.

Lemma 2: *We can solve the prefix computation problem on any input of size n in $O(\log n)$ time using $\frac{n}{\log n}$ CREW PRAM processors. \square*

Observation: The above algorithm is asymptotically optimal since $S = n$ and the work done by the parallel algorithm is $O(n)$.

Some Applications

- Numerous problems can be solved efficiently by reducing them to prefix computations. We'll see two examples.
- *Computing the rank of an element:* Here we are given a sequence $X = k_1, k_2, \dots, k_n$ of arbitrary real numbers and a number $k \in X$. The goal is to compute the rank of k in X . (Recall that $\text{rank}(k, X) = |\{q \in X : q < k\}| + 1$). This problem can be reduced to prefix sums computation as follows.

- 1) Create a bit array $a[1 : n]$ such that $a[i] = 1$ if $k_i < k$ and $a[i] = 0$ otherwise, for $1 \leq i \leq n$.
 - 2) Compute the prefix sums of $a[1], a[2], \dots, a[n]$ to get $b[1], b[2], \dots, b[n]$;
 - 3) Output $b[n] + 1$;
- Observe that step 1 can be done in $O(\log n)$ time using $\frac{n}{\log n}$ CREW PRAM processors (by assigning $\log n$ input elements per processor). Step 2 takes $O(\log n)$ time using $\frac{n}{\log n}$ CREW PRAM processors (c.f. Lemma 2). Step 3 takes 1 unit of time. Thus the entire algorithm runs in $O(\log n)$ time using $\frac{n}{\log n}$ CREW PRAM processors.
 - As a corollary to the above algorithm, we can sort n given elements in $O(\log n)$ time using $\frac{n^2}{\log n}$ CREW PRAM processors. The idea is to assign $\frac{n}{\log n}$ processors per key and compute its rank in parallel. Followed by this, the keys are output based on their ranks. Let $X = k_1, k_2, \dots, k_n$ be the input.

- 1) **for** $i = 1$ **to** n **in parallel do**
 - 2) Using $\frac{n}{\log n}$ processors compute the rank r_i of k_i ;
 - 3) **for** $i = 1$ **to** n **in parallel do**
 - 4) Processor i outputs k_i in memory cell r_i ;
- *Splitting an input sequence:* Here the input is a sequence $X = k_1, k_2, \dots, k_n$ of arbitrary real numbers and another real number k . The goal is to permute X such that all the elements of X that are smaller than k appear before the rest of the elements. We can also reduce this problem to prefix computations as follows.
 - 1) Create a bit array $a[1 : n]$ such that $a[i] = 1$ if $k_i < k$ and $a[i] = 0$ otherwise, for $1 \leq i \leq n$.
 - 2) Compute the prefix sums of $a[1], a[2], \dots, a[n]$ to get $b[1], b[2], \dots, b[n]$;
 - 3) **for** $i = 1$ **to** n **in parallel do**
 - 4) **if** $k_i < k$ **then** write k_i in cell $b[i]$;

/* Note that all the elements less than k have now
been placed in successive cells */

 - 5) Place the elements of X that are $> k$ in a similar manner;

Run time analysis: Step 1 can be done in $O(\log n)$ time using $\frac{n}{\log n}$ processors (by assigning $\log n$ elements per processor). Step 2 takes $O(\log n)$ time (c.f. Lemma 2). Step 3 also can be done in $O(\log n)$ time (if we assign $\log n$ elements per processor). Steps 1 through 4 thus take a total of $O(\log n)$ time. As a result, Step 5 will also take the same amount of time. Therefore, the total run time of the algorithm is $O(\log n)$.

Parallel Sorting

- A number of parallel algorithms have been proposed in the literature. The first deterministic optimal parallel algorithm proposed was by AKS in 1981. This was for a sorting network. Reischuk parallellized Frazer and McKellar's algorithm on the CRCW PRAM to get an asymptotically optimal randomized sorting algorithm (1981). In the next Lecture we will study Preparata's algorithm.