## COMP 322: Fundamentals of Parallel Programming

# Lecture 37: Algorithms based on Parallel Prefix (Scan) operations (contd) 

Instructors: Vivek Sarkar, Mack Joyner Department of Computer Science, Rice University \{vsarkar, mjoyner\}@rice.edu
Acknowledgements:

- Book chapter on "Prefix Sums and Their Applications", Guy E. Blelloch, CMU
- Slides on "Parallel prefix adders", Kostas Vitoroulis, Concordia University


## http://comp322.rice.edu/

## Recap of Parallel Prefix and Scan

- In Lecture 13, we learned how to compute the parallel prefix sum (set of partial sums), $X$, for an input array A of size $n$ with WORK $=0(n)$ and $C P L=O(\log n)$ - For $P$ processors, $T_{P}=O(n / P+\log (P))$
- In Lecture 36, we learned the scan operator which generalizes parallel prefix sum to parallel prefix for any binary associative operator, $\oplus$ (need not be restricted to xor).
- Specifically, scan takes a binary associative operator $\oplus$, and an array of $n$ elements $\left[A_{0}, A_{1}, \ldots, A_{n-1}\right]$ as input, and returns array $\left[A_{0},\left(A_{0} \oplus A_{1}\right), \ldots,\left(A_{0} \oplus A_{1} \oplus \ldots \oplus A_{n-1}\right)\right]$ as output. The output from scan can also be specified as a recurrence:

$$
x_{i}= \begin{cases}a_{0} & i=0 \\ x_{i-1} \oplus a_{i} & 0<i<n\end{cases}
$$

- From Lecture 13, we see that scan can also be performed with WORK $=0(n)$ and CPL = O(log n)
- Note that any constant sequence of scan operators can also be performed with WORK $=0(n)$ and CPL $=0(\log n)$


## Parallelizing Prefix Sum (Lecture 13)

Observation: each prefix sum can be decomposed into reusable terms of power-of-2-size e.g.

$$
\begin{aligned}
X[6] & =A[0]+A[1]+A[2]+A[3]+A[4]+A[5]+A[6] \\
& =(A[0]+A[1]+A[2]+A[3])+(A[4]+A[5])+A[6]
\end{aligned}
$$

Approach:

- Combine reduction tree idea from Parallel Array Sum with partial sum idea from Sequential Prefix Sum
- Use an "upward sweep" to perform parallel reduction, while storing partial sum terms in tree nodes
- Use a "downward sweep" to compute prefix sums while reusing partial sum terms stored in upward sweep


## Parallel Prefix Sum: Upward Sweep (while calling scan recursively)

Upward sweep is just like Parallel Reduction, except that partial sums are also stored along the way

1. Receive values from left and right children
2. Compute left+right and store in box
3. Send left+right value to parent

15


## Parallel Prefix Sum: Downward Sweep (while returning from recursive calls to scan)

1. Receive value from parent (root receives 0 )
2. Send parent's value to LEFT child (prefix sum for elements to left of left child's subtree)
3. Send parent's value+ left child's box value to RIGHT child (prefix sum for elements to left of right child's subtree)
4. Add $A[i]$ to get inclusive prefix sum


Inclusive prefix sums
346
$6 \quad 10$
1112
15

## Worksheet \#36 problem statement: Parallelizing the Split step in Radix Sort

The Radix Sort algorithm loops over the bits in the binary representation of the keys, starting at the lowest bit, and executes a split operation for each bit as shown below. The split operation packs the keys with a 0 in the corresponding bit to the bottom of a vector, and packs the keys with a 1 to the top of the same vector. It maintains the order within both groups. The sort works because each split operation sorts the keys with respect to the current bit and maintains the sorted order of all the lower bits. Your task is to show how the split operation can be performed in parallel using scan operations, and to explain your answer.


# Worksheet \#36 solution: Parallelizing the Split step in Radix Sort 



## FIGURE 1.9

The split operation packs the elements with a 0 in the corresponding flag position to the bottom of a vector, and packs the elements with a 1 to the top of the same vector. The permute writes each element of A to the index specified by the corresponding position in Index.

## Binary Addition

This is the pen and paper addition of two 4-bit binary numbers $\mathbf{x}$ and $\mathbf{y}$. c represents the generated carries. s represents the produced sum bits.


A stage of the addition is the set of $\mathbf{x}$ and $\mathbf{y}$ bits being used to produce the appropriate sum and carry bits. For example the highlighted bits $\mathrm{x}_{2}$, $y_{2}$ constitute stage 2 which generates carry $\mathrm{C}_{2}$ and sum $\mathrm{s}_{2}$.

Each stage $i$ adds bits $a_{i}, b_{i}, c_{i-1}$ and produces bits $s_{i}, c_{i}$ The following hold:

| $\mathrm{a}_{\mathrm{i}}$ | $\mathrm{b}_{\mathrm{i}}$ | $\mathrm{c}_{\mathrm{i}}$ | Comment: | Formal definition: |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 0 | 0 | The stage "kills" an incoming carry. | "Kill" bit: | $k_{i}=\overline{x_{i}+y_{i}}$ |
| 0 | 1 | $\mathrm{c}_{\mathrm{i}-1}$ | The stage "propagates" an incoming carry | "Propagate" bit: | $p_{i}=x_{i} \oplus y_{i}$ |
| 1 | 0 | $\mathrm{c}_{\mathrm{i}-1}$ | The stage "propagates" an incoming carry |  |  |
| 1 | 1 | 1 | The stage "generates" a carry out | "Generate" bit: | $g_{i}=x_{i} \bullet y_{i}$ |

## Binary Addition

| $\mathrm{a}_{\mathrm{i}}$ | $\mathrm{b}_{\mathrm{i}}$ | $\mathrm{c}_{\mathrm{i}}$ | Comment: | Formal definition: |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 0 | 0 | The stage "kills" an incoming carry. | "Kill" bit: | $k_{i}=\overline{x_{i}+y_{i}}$ |
| 0 | 1 | $\mathrm{c}_{\mathrm{i}-1}$ | The stage "propagates" an incoming carry | "Propagate" bit: | $p_{i}=x_{i} \oplus y_{i}$ |
| 1 | 0 | $\mathrm{c}_{\mathrm{i}-1}$ | The stage "propagates" an incoming carry |  |  |
| 1 | 1 | 1 | The stage "generates" a carry out | "Generate" bit: | $g_{i}=x_{i} \bullet y_{i}$ |

The carry $\mathrm{c}_{\mathrm{i}}$ generated by a stage $i$ is given by the equation:

$$
c_{i}=g_{i}+p_{i} \cdot c_{i-1}=x_{i} \cdot y_{i}+\left(x_{i} \oplus y_{i}\right) \cdot c_{i-1}
$$

This equation can be simplified to:

$$
c_{i}=x_{i} \cdot y_{i}+\left(x_{i}+y_{i}\right) \cdot c_{i-1}=g_{i}+a_{i} \cdot c_{i-1}
$$

The " $a_{i}$ " term in the equation being the "alive" bit.
The later form of the equation uses an OR gate instead of an XOR which is a more efficient gate when implemented in CMOS technology. Note that:

$$
a_{i}=\overline{k_{i}}
$$

Where $\mathrm{k}_{\mathrm{i}}$ is the "kill" bit defined in the table above.
COMP 322, Spring 2017 (V. Sarkar, M. Joyner)

## Binary addition as a prefix sum problem.

- We define a new operator:
- Input is a vector of pairs of 'propagate' and 'generate' bits:

$$
\left(g_{n}, p_{n}\right)\left(g_{n-1}, p_{n-1}\right) \ldots\left(g_{0}, p_{0}\right)
$$

- Output is a new vector of pairs:

$$
\left(G_{n}, P_{n}\right)\left(G_{n-1}, P_{n-1}\right) \ldots\left(G_{0}, P_{0}\right)
$$

- Each pair of the output vector is calculated by the following definition:


$$
\begin{aligned}
& \left(G_{0}, P_{0}\right)=\left(g_{0}, p_{0}\right) \\
& \left(g_{x}, p_{x}\right) \circ\left(g_{y}, p_{y}\right)=\left(g_{x}+p_{x} \cdot g_{y}, p_{x} \cdot p_{y}\right) \\
& \text { with +, being the OR, AND operations }
\end{aligned}
$$

## 1973: Kogge-Stone adder



- The Kogge-Stone adder has:
$\square$ Low depth
$\square$ High node count (implies more area).
$\square$ Minimal fan-out of 1 at each node (implies faster ${ }_{1}$ performance).


## Summary

- A parallel prefix adder can be seen as a 3-stage process:

- There exist various architectures for the carry calculation part.
- Trade-offs in these architectures involve the
$\square$ area of the adder
$\square$ its depth
$\square$ the fan-out of the nodes
$\square$ the overall wiring network.


## Parallel Algorithms, Computation Graphs, and Circuits

- Today's lecture shows that parallel algorithms, computation graphs, and circuits represent different approaches to parallel computational thinking
- A parallel algorithm unfolds into a computation graph when executing
- A circuit represents an "unrolled" computation graph in hardware e.g., see bitonic sorting network in https://en.wikipedia.org/wiki/Bitonic sorter



## And the same principles can be applied to distributed computing too, e.g., MPI_Scan

```
1.main(int argc, char **argv){
int mype, ierr;
    ierr = MPI_Init(&argc, &argv);
    ierr = MPI_Comm_rank(WCOMM, &mype);
    doub7e 1 = 0.5;
    double exp_sum = 0;
    double exp_pdf_i = 0.0;
    double exp_cdf_i = 0.0;
    double DIV_CONST = 2.0;
    int i;
    for(i = 0; i < NUMPTS; i++)
    {
        if (i == mype)
        {
            exp_pdf_i = 1*exp(-1 * ((double) i) / DIV_CONST);
        }
    }
    //calculate the cumulative frequency histogram (exp_cdf_i) from exp_pdf_i
    ierr = MPI_Scan( \
        &exp_pdf_i, &exp_cdf_i, 1, MPI_DOUBLE, MPI_SUM, WCOMM);
        for (i = 0; i < NUMPTS; i++)
        {
        if (i == mype)
        {
            printf("process %d: cumulative sum = %lf\n", \
            mype, exp_cdf_i);
        }
    }
```


## Announcements \& Reminders

- HW5 is now available
- Due April 21st, with automatic extension until May 1st
- We are making an optional quiz available for Unit 10
- Also due April 21st, with automatic extension until May 1st
- Optional, but we will pick the best 9 of your 10 quiz scores
- As always, please post any issues with quiz questions on Piazza
- Two more lectures + grand finale at 3pm on Friday
- April 19th, 1 pm: GPU Computing
- April 21st, 1pm: Course review (scope of Exam 2) in class
- April 21st, 3pm: COMP 322 grand finale in DH 3092 (group office hours time slot, but cake will be served!)

