


Programming Massively Parallel Processors


A Hands-on Approach

CHAPTER 11 > Prefix Sum (Scan)





- A **scan** operation:
 - Takes:
 - An input array $[x_0, x_1, \dots, x_{n-1}]$
 - An associative operator \oplus
 - e.g., sum, product, min, max
 - Returns:
 - An output array $[y_0, y_1, \dots, y_{n-1}]$ where
 - **Inclusive scan**: $y_i = x_0 \oplus x_1 \oplus \dots \oplus x_i$
 - **Exclusive scan**: $y_i = x_0 \oplus x_1 \oplus \dots \oplus x_{i-1}$

- Addition example:



- In general:



- Sequential scan for sum:

```
output[0] = input[0];  
for(i = 1; i < N; ++i) {  
    output[i] = output[i-1] + input[i];  
}
```

Inclusive Scan

```
output[0] = 0.0f;  
for(i = 1; i < N; ++i) {  
    output[i] = output[i-1] + input[i-1];  
}
```

Exclusive Scan

- In general:

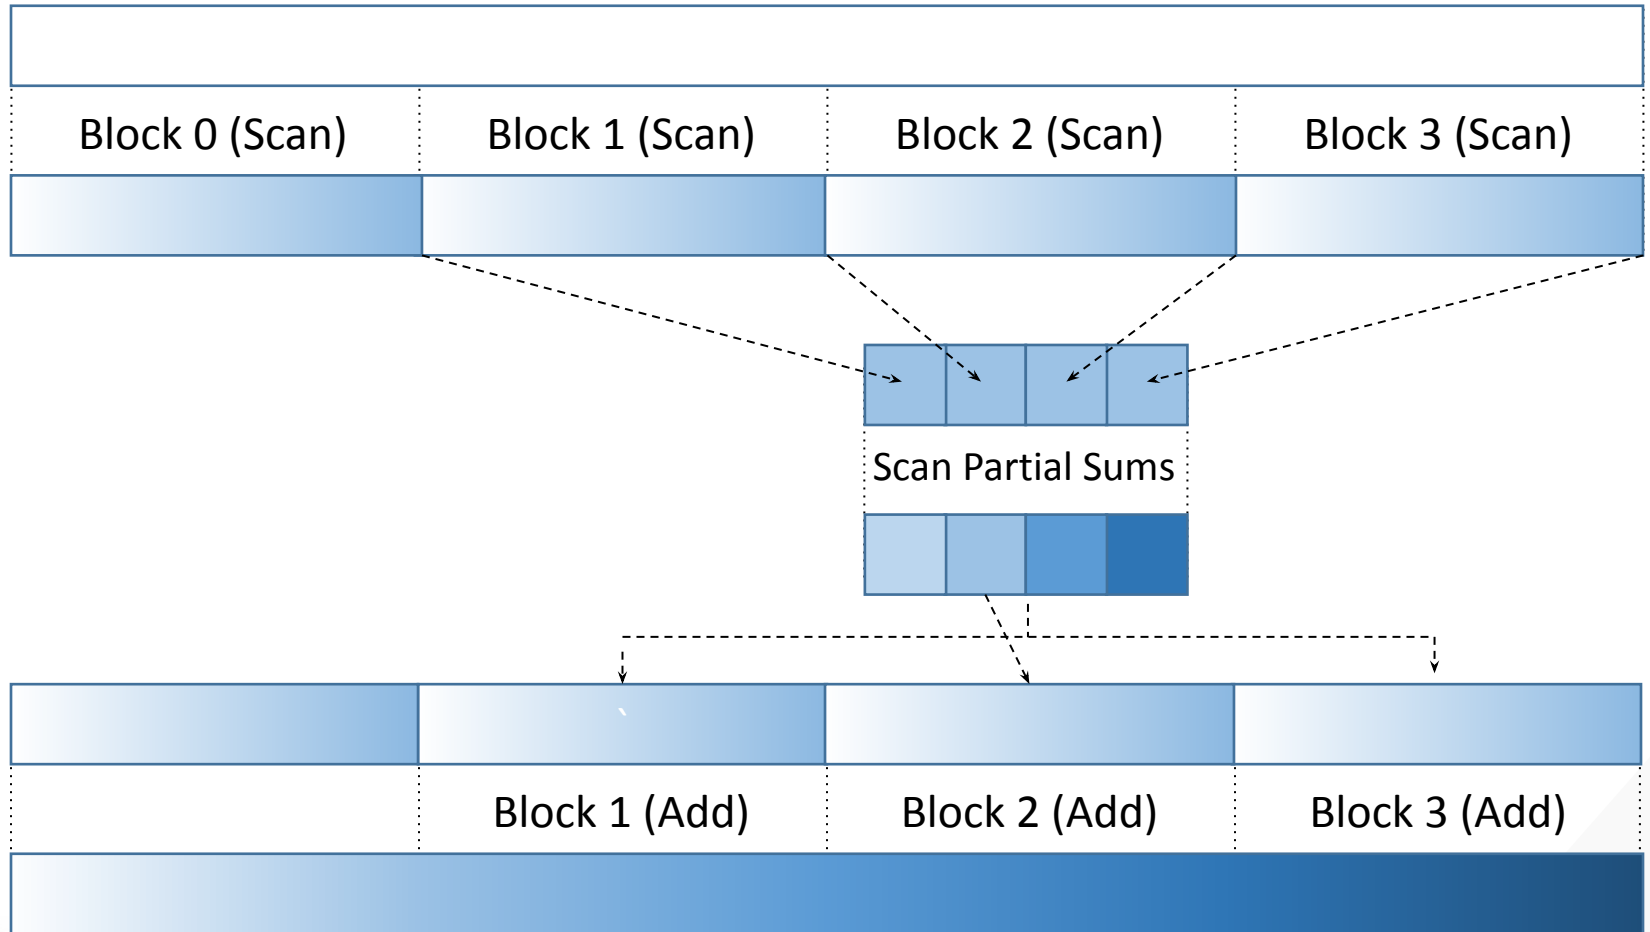
```
output[0] = input[0];  
for(i = 1; i < N; ++i) {  
    output[i] = f(output[i-1], input[i]);  
}
```

Inclusive Scan


```
output[0] = IDENTITY;  
for(i = 1; i < N; ++i) {  
    output[i] = f(output[i-1], input[i-1]);  
}
```

Exclusive Scan


- Parallel scan requires synchronization across parallel workers
- Approach: **segmented scan**
 - Every thread block scans a segment
 - Scan the segments' partial sums
 - Add each segment's scanned partial sum to the next segment




For now, we will focus on implementing a parallel scan in each block




A parallel reduction tree for the last element gives some others as a byproduct




Another reduction tree gives us more elements




Keep doing reduction trees until we get all answers




Keep doing reduction trees until we get all answers




Overlap the trees and do them simultaneously




One thread for each element





One thread for each element

Optimization: load once to a **shared memory** buffer and perform successive reads and writes to the same array can be done in shared memory



```
unsigned int i = blockIdx.x*blockDim.x + threadIdx.x;
```

```
__shared__ float buffer_s[BLOCK_DIM];  
buffer_s[threadIdx.x] = input[i];  
__syncthreads();
```

```
for(unsigned int stride = 1; stride <= BLOCK_DIM/2; stride *= 2) {  
    if(threadIdx.x >= stride) {  
        buffer_s[threadIdx.x] += buffer_s[threadIdx.x - stride];  
    }  
    __syncthreads();  
}
```

Incorrect!

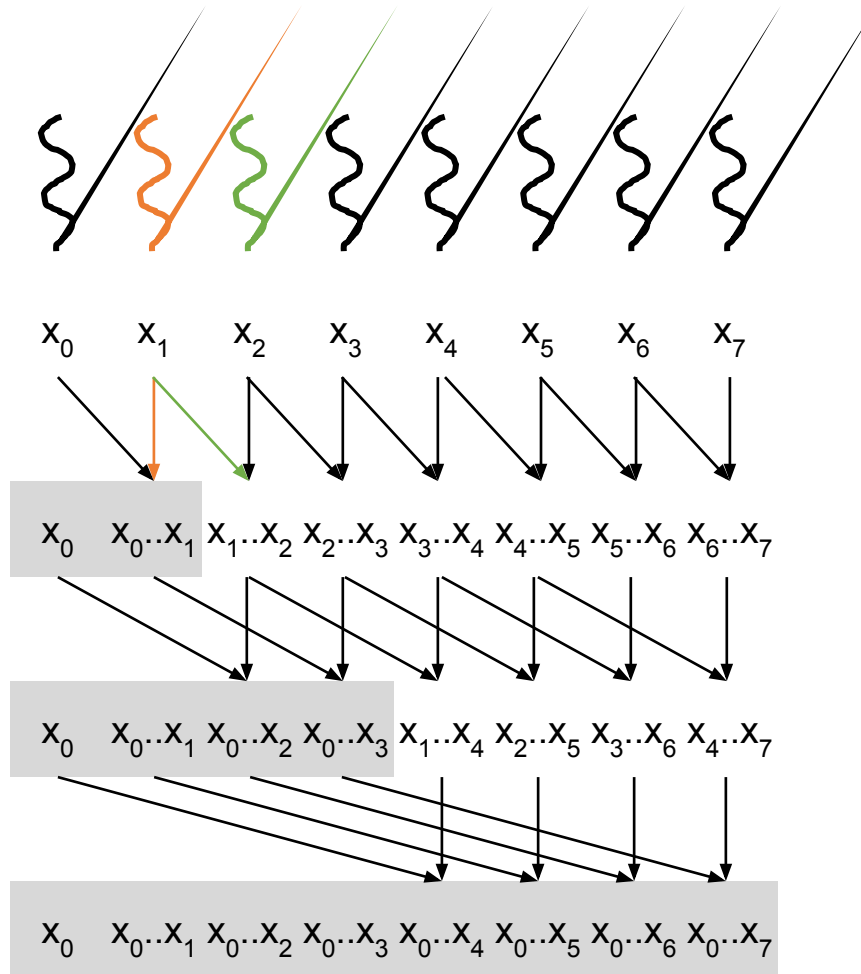
Different threads are reading and writing the same data location without synchronizing

```
if(threadIdx.x == BLOCK_DIM - 1) {  
    partialSums[blockIdx.x] = buffer_s[threadIdx.x];  
}
```

```
output[i] = buffer_s[threadIdx.x];
```

Thread 1 may update value at index 1 before thread 2 reads it

Solution: wait for everyone to read before updating



```
unsigned int i = blockIdx.x*blockDim.x + threadIdx.x;

__shared__ float buffer_s[BLOCK_DIM];
buffer_s[threadIdx.x] = input[i];
__syncthreads();

for(unsigned int stride = 1; stride <= BLOCK_DIM/2; stride *= 2) {
    float v;
    if(threadIdx.x >= stride) {
        v = buffer_s[threadIdx.x - stride];
    }
    __syncthreads(); _____ Wait for everyone to read
                                before writing
    if(threadIdx.x >= stride) {
        buffer_s[threadIdx.x] += v;
    }
    __syncthreads();
}

if(threadIdx.x == BLOCK_DIM - 1) {
    partialSums[blockIdx.x] = buffer_s[threadIdx.x];
}

output[i] = buffer_s[threadIdx.x];
```



```
unsigned int i = blockIdx.x*blockDim.x + threadIdx.x;
```

```
__shared__ float buffer_s[BLOCK_DIM];
buffer_s[threadIdx.x] = input[i];
__syncthreads();
```

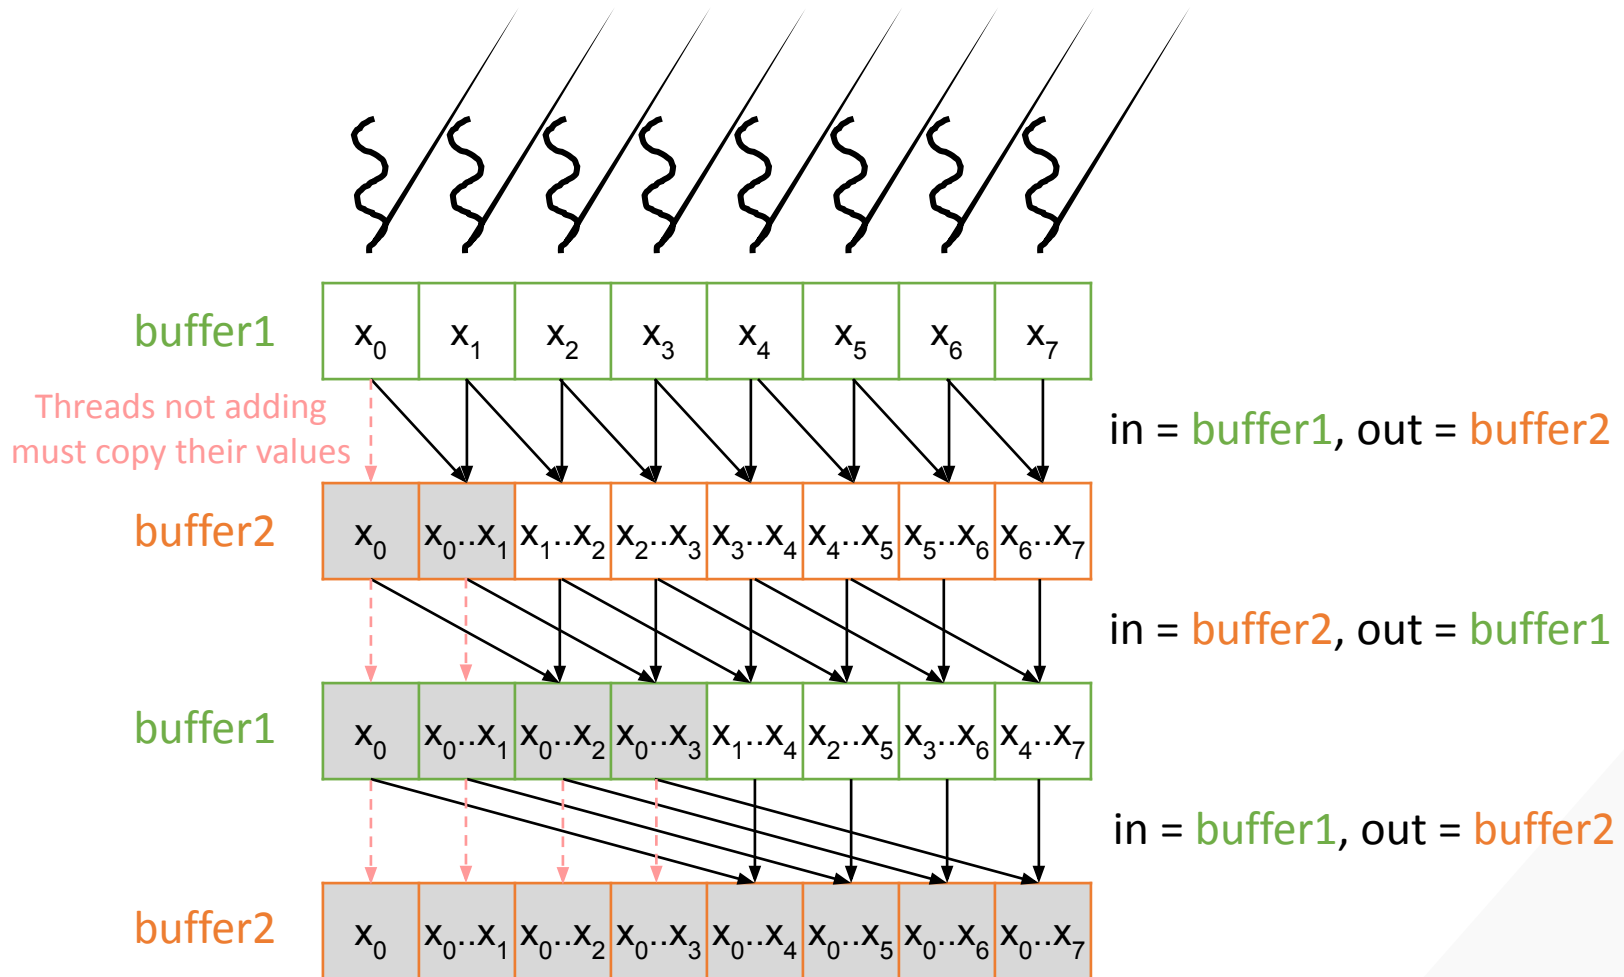
```
for(unsigned int stride = 1; stride <= BLOCK_DIM/2; stride *= 2) {
    float v;
    if(threadIdx.x >= stride) {
        v = buffer_s[threadIdx.x - stride];
    }
    __syncthreads();
    if(threadIdx.x >= stride) {
        buffer_s[threadIdx.x] += v;
    }
    __syncthreads();
}
```

This synchronization enforces a **false dependence**
(we only need to finish reading before others write
because we are using the same buffer)

This synchronization enforces a **true dependence**
(we must finish writing before others can read)

```
if(threadIdx.x == BLOCK_DIM - 1) {
    partialSums[blockIdx.x] = buffer_s[threadIdx.x];
}
```

```
output[i] = buffer_s[threadIdx.x];
```



Optimization: eliminate the synchronization that enforces a false dependence by using separate buffers for reading and writing, and alternate the buffers each iteration (called **double buffering**)

```
unsigned int i = blockIdx.x*blockDim.x + threadIdx.x;

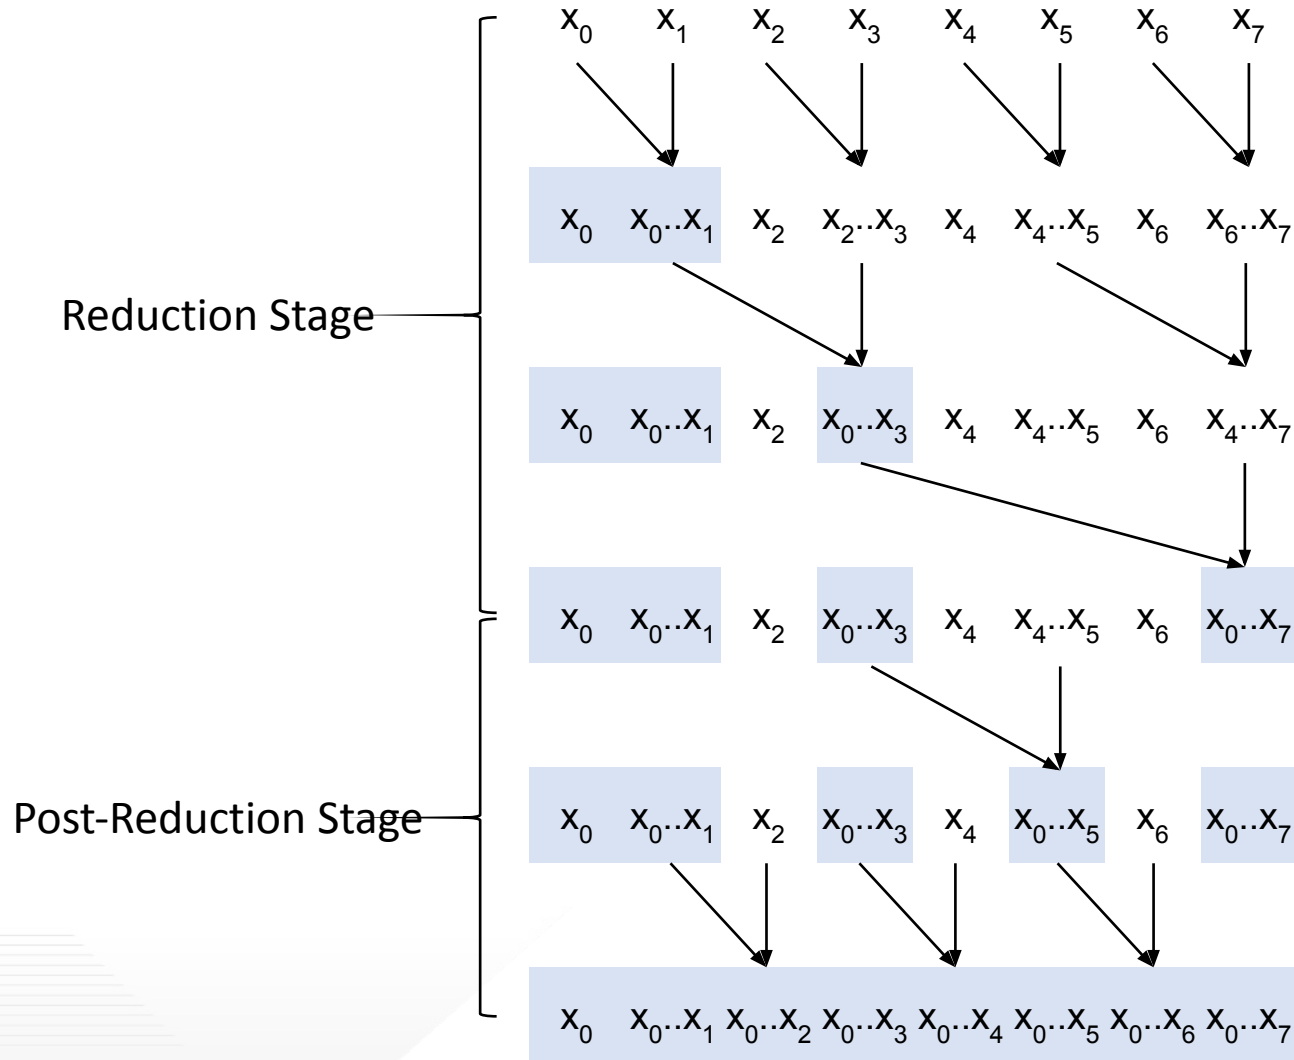
__shared__ float buffer1_s[BLOCK_DIM];
__shared__ float buffer2_s[BLOCK_DIM];
float* inBuffer_s = buffer1_s;
float* outBuffer_s = buffer2_s;
inBuffer_s[threadIdx.x] = input[i];
__syncthreads();

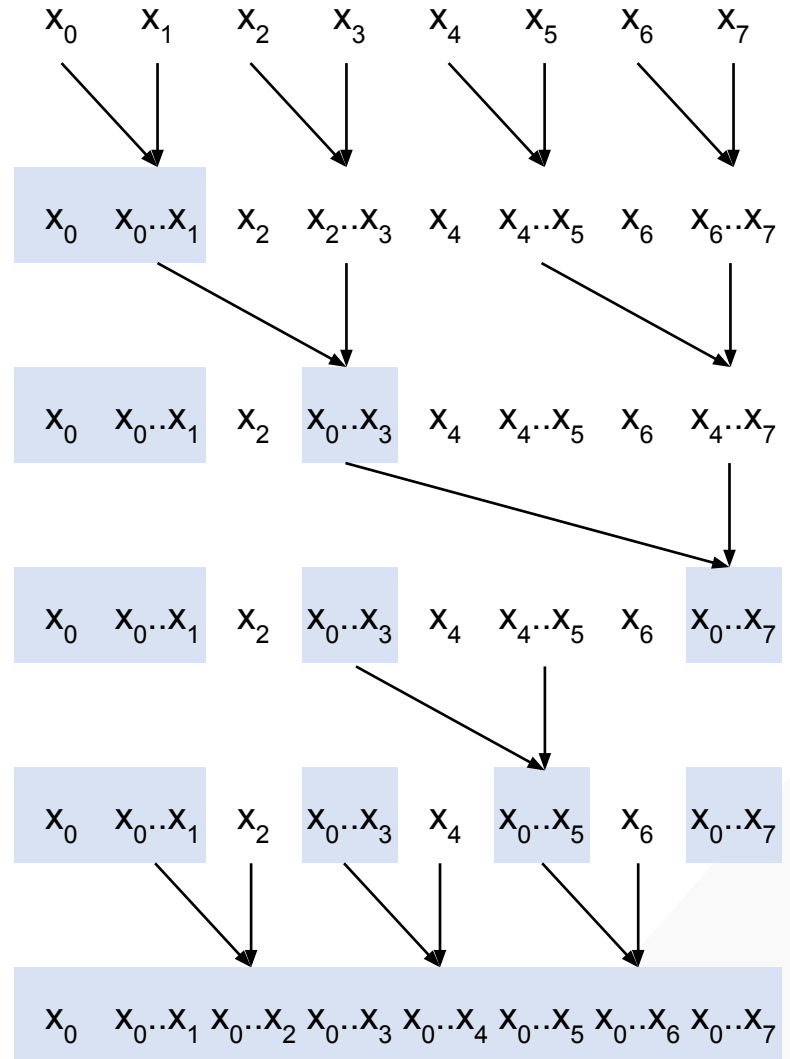

for(unsigned int stride = 1; stride <= BLOCK_DIM/2; stride *= 2) {
    if(threadIdx.x >= stride) {
        outBuffer_s[threadIdx.x] =
            inBuffer_s[threadIdx.x] + inBuffer_s[threadIdx.x - stride];
    } else {
        outBuffer_s[threadIdx.x] = inBuffer_s[threadIdx.x];
    }
    __syncthreads();
    float* tmp = inBuffer_s;
    inBuffer_s = outBuffer_s;
    outBuffer_s = tmp;
}

if(threadIdx.x == BLOCK_DIM - 1) {
    partialSums[blockIdx.x] = inBuffer_s[threadIdx.x];
}

output[i] = inBuffer_s[threadIdx.x];
```

- A parallel algorithm is **work-efficient** if it performs the same amount of work as the corresponding sequential algorithm
- Scan work efficiency
 - Sequential scan performs N additions
 - Kogge-Stone parallel scan performs:
 - $\log(N)$ steps, $N - 2^{\text{step}}$ operations per step
 - Total: $(N-1) + (N-2) + (N-4) + \dots + (N-N/2)$
 $= N * \log(N) - (N-1) = O(N * \log(N))$ operations
 - Algorithm is not work efficient
- If resources are limited, parallel algorithm will be slow because of low work efficiency





- Recall: Kogge-Stone
 - **$\log(N)$ steps**
 - **$O(N \cdot \log(N))$ operations**
- Brent-Kung
 - Reduction stage:
 - $\log(N)$ steps
 - $N/2 + N/4 + \dots + 4 + 2 + 1 = N-1$ operations
 - Post-Reduction stage:
 - $\log(N)-1$ steps
 - $(2-1) + (4-1) + \dots + (N/2-1) = (N-2) - (\log(N)-1)$
 - Total:
 - **$2 \cdot \log(N) - 1$ steps**
 - $(N-1) + (N-2) - (\log(N)-1) = 2 \cdot N - \log(N) - 2 = O(N)$ operations
- Brent-Kung takes **more steps** but is **more work-efficient**
- So which one is faster?

- Wen-mei W. Hwu, David B. Kirk, and Izzat El Hajj. *Programming Massively Parallel Processors: A Hands-on Approach*. Morgan Kaufmann, 2022.