CS 380 - GPU and GPGPU Programming
Lecture 23: GPU Parallel Prefix Sum / Scan

Markus Hadwiger, KAUST

Reading Assignment #14 (until Dec 6)

Read (required):

- Warp Shuffle Functions
  - CUDA Programming Guide 11.5, Appendix B.22
- CUDA Cooperative Groups
  - CUDA Programming Guide 11.5, Appendix C
  - `https://developer.nvidia.com/blog/cooperative-groups/`
- Programming Tensor Cores
  - CUDA Programming Guide 11.5, Appendix B.24 (Warp matrix functions)
  - `https://developer.nvidia.com/blog/programming-tensor-cores-cuda-9/`

Read (optional):

- CUDA Warp-Level Primitives
  - `https://developer.nvidia.com/blog/using-cuda-warp-level-primitives/`
- Warp-aggregated atomics
  - `https://developer.nvidia.com/blog/cuda-pro-tip-optimized-filtering-warp-aggregated-atomics/`

GPU Parallel Prefix Sum

Parallel Prefix Sum (Scan)

- **Definition:**

  The all-prefix-sums operation takes a binary associative operator $\oplus$ with identity I, and an array of n elements

  $$[a_0, a_1, \ldots, a_{n-1}]$$

  and returns the ordered set

  $$[I, a_0, (a_0 \oplus a_1), \ldots, (a_0 \oplus a_1 \oplus \ldots \oplus a_{n-2})].$$

Exclusive scan: last input element is not included in the result

- **Example:**
  if $\oplus$ is addition, then scan on the set

  [3 1 7 0 4 1 6 3]

  returns the set

  [0 3 4 11 11 15 16 22]

(From Blelloch, 1990, "Prefix Sums and Their Applications)

Applications of Scan

- **Scan is a simple and useful parallel building block**
  - Convert recurrences from sequential :

```
for(j=1;j<n;j++)
   out[j] = out[j-1] + f(j);
```

  - into parallel:

```
forall(j) { temp[j] = f(j) };
scan(out, temp);
```

- **Useful for many parallel algorithms:**
  - radix sort
  - quicksort
  - String comparison
  - Lexical analysis
  - Stream compaction
  - Polynomial evaluation
  - Solving recurrences
  - Tree operations
  - Range Histograms
  - Etc.

Scan on the CPU

```
void scan( float* scanned, float* input, int length)
{
  scanned[0] = 0;
  for(int i = 1; i < length; ++i)
  {
    scanned[i] = input[i-1] + scanned[i-1];
  }
}
```

- **Just add each element to the sum of the elements before it**
- **Trivial, but sequential**
- **Exactly *n* adds: optimal in terms of work efficiency**

Prefix Sum Application
- Compaction -

Parallel Data Compaction

- **Given an array of marked values**

3	1	7	4	2	1	5	6	3	1
1	0	1	0	0	0	0	1	0	0

- **Output the compacted list of marked values**

3	7	6

Using Prefix Sum

- **Calculate prefix sum on index array**

3	1	7	4	2	1	5	6	3	1
1	0	1	0	0	0	0	1	0	0
0	1	1	2	2	2	2	2	3	3

- **For each marked value lookup the destination index in the prefix sum**

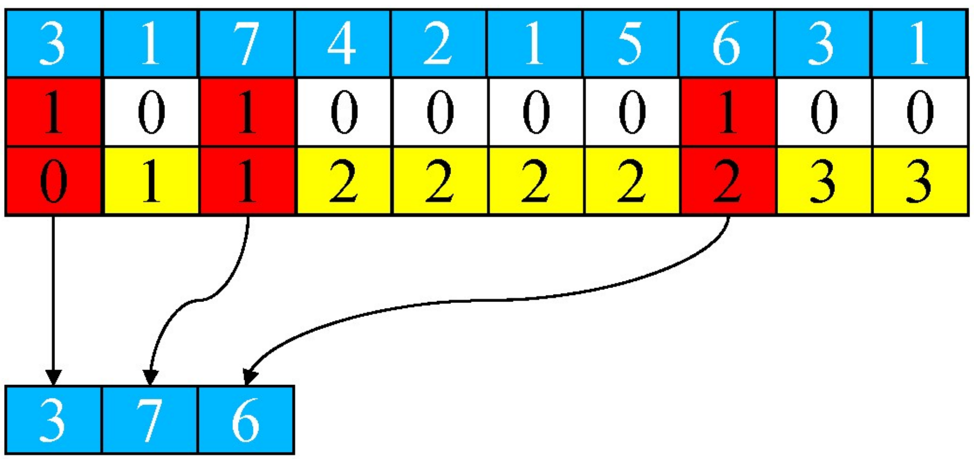

- **Parallel write to separate destination elements**

Prefix Sum Application
- Range Histogram -

Range Histogram

- **A histogram calculate the occurance of each value in an array.**

  $$h[i] = |J| \quad J=\{j \mid v[j] = i\}$$

- **Range query: number over elements in interval [a,b].**

- **Slow answer:**

```
hrange = 0;
for (i = a; i<=b; ++i)
    hrange += h[i];
```

Fast Range Histogram

- **Compute prefix sum of histogram**
- **Fast answer:**

```
hrange = pref[B] - pref[A];
```

$$= \sum_{0}^{B} h[i] - \sum_{0}^{A} h[i] = \sum_{A}^{B} h[i]$$

Prefix Sum Application
- Summed Area Tables -

Summed Area Tables

- Per texel, store sum from (0, 0) to (u, v)

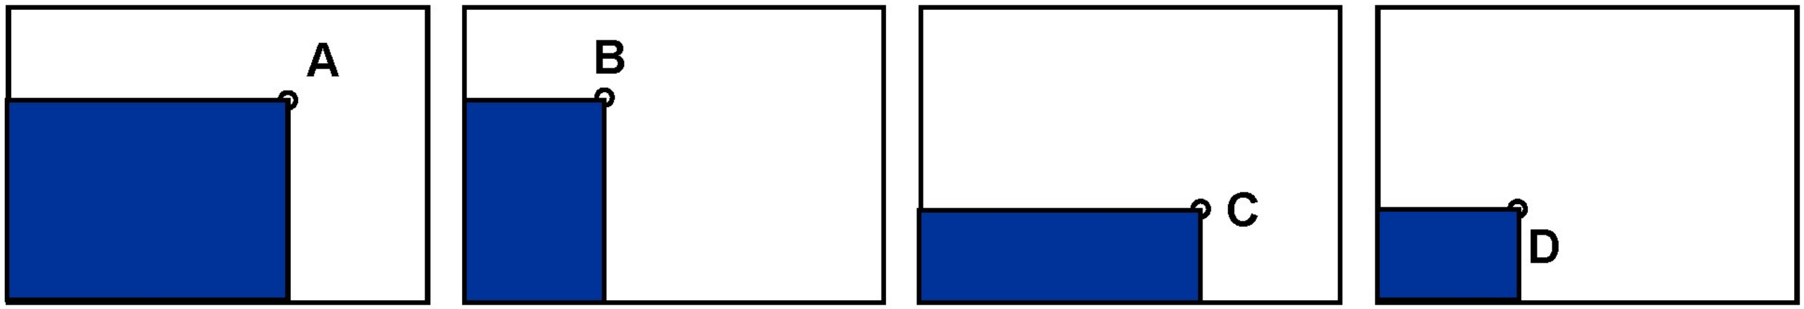

- Many bits per texel (sum !)
- Evaluation of 2D integrals in constant time!

$$\int_{Bx}^{Ax}\int_{Cy}^{Ay} I(x,y)\,dx\,dy = A - B - C + D$$

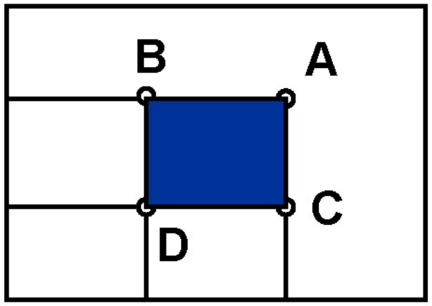

Summed Area Table with Prefix Sums

- One possible way:
- Compute prefix sum horizontally

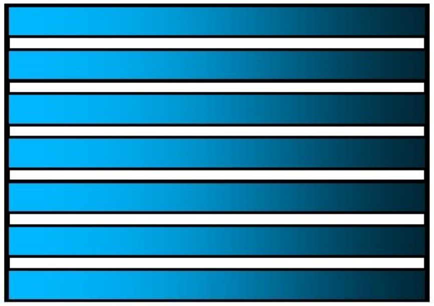

- … then vertically on the result

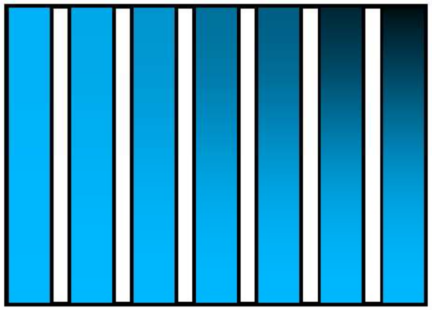

Work Efficiency

Guy E. Blelloch and Bruce M. Maggs:
Parallel Algorithms, 2004 (https://www.cs.cmu.edu/~guyb/papers/BM04.pdf)

In designing a parallel algorithm, it is more important to make it efficient than to make it asymptotically fast. The efficiency of an algorithm is determined by the total number of operations, or work that it performs. On a sequential machine, an algorithm's work is the same as its time. On a parallel machine, the work is simply the processor-time product. Hence, an algorithm that takes time t on a P-processor machine performs work W = Pt. In either case, the work roughly captures the actual cost to perform the computation, assuming that the cost of a parallel machine is proportional to the number of processors in the machine.

We call an algorithm work-efficient (or just efficient) if it performs the same amount of work, to within a constant factor, as the fastest known sequential algorithm.

For example, a parallel algorithm that sorts n keys in O(sqrt(n) log(n)) time using sqrt(n) processors is efficient since the work, O(n log(n)), is as good as any (comparison-based) sequential algorithm.

However, a sorting algorithm that runs in O(log(n)) time using n^2 processors is not efficient.

The first algorithm is better than the second - even though it is slower - because its work, or cost, is smaller. Of course, given two parallel algorithms that perform the same amount of work, the faster one is generally better.

Vector Reduction

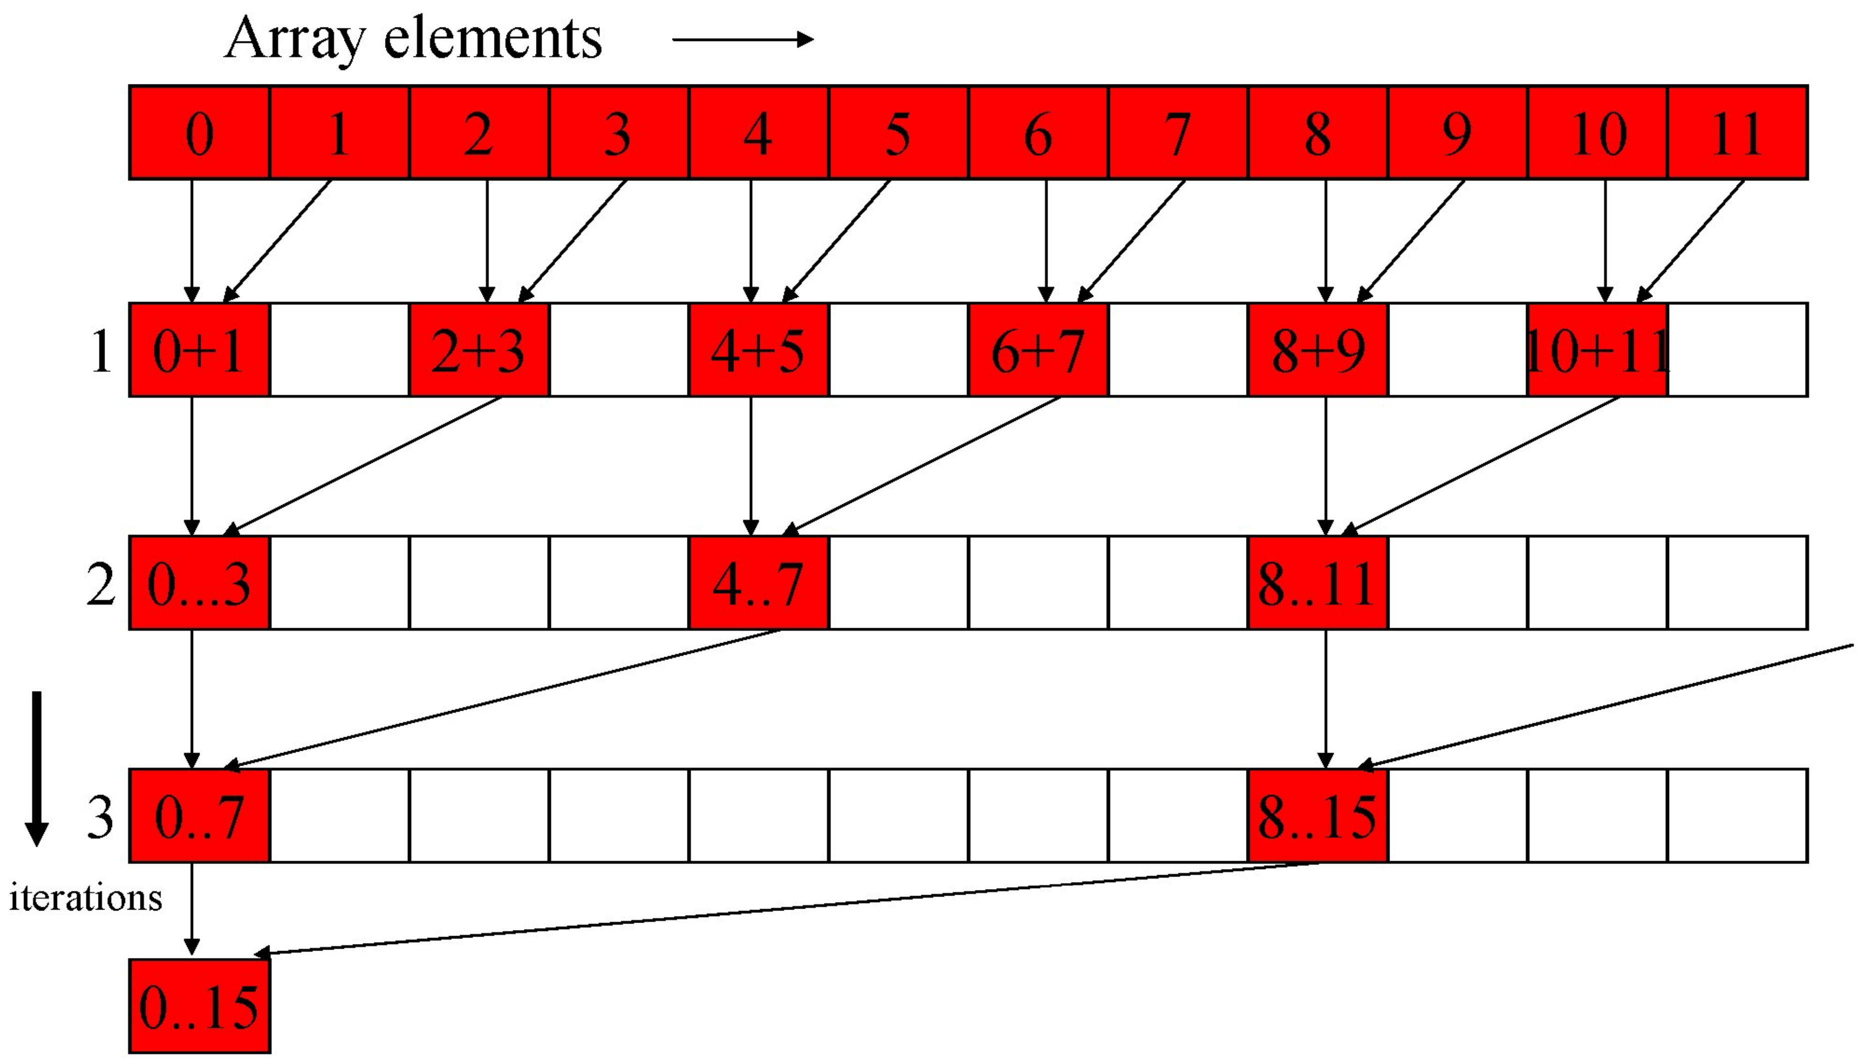

Typical Parallel Programming Pattern

- **log(n) steps**

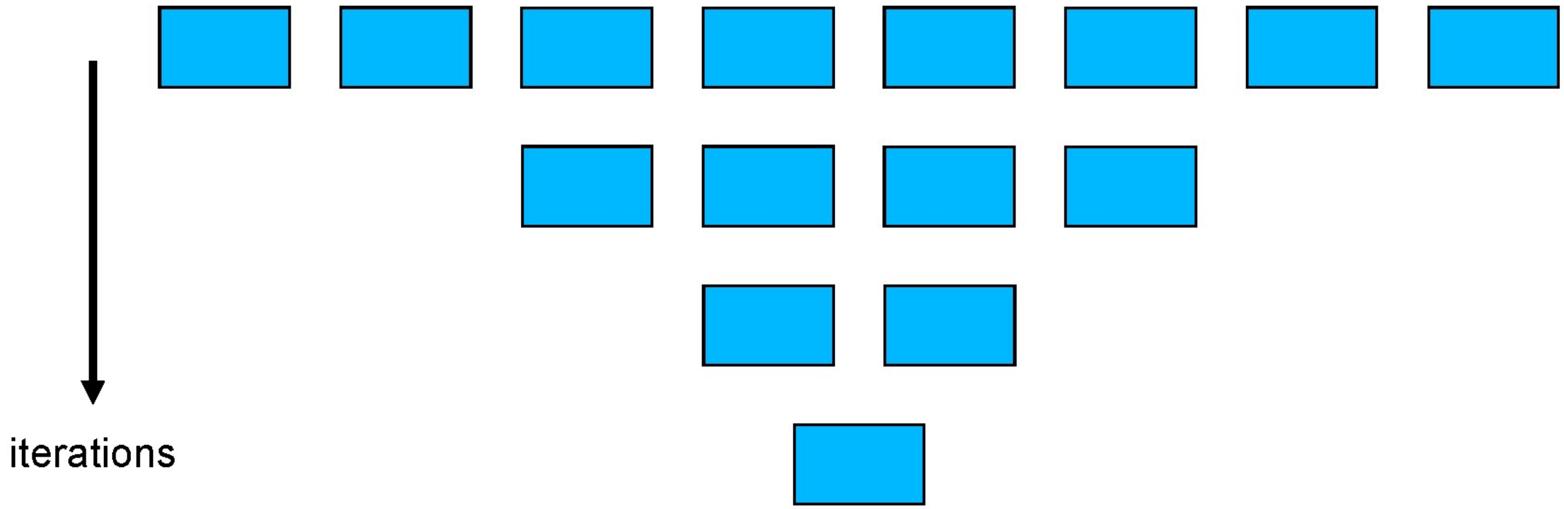

Helpful fact for counting nodes of full binary trees:
If there are N leaf nodes, there will be N-1 non-leaf nodes

Courtesy John Owens

Kogge-Stone Scan

Circuit family

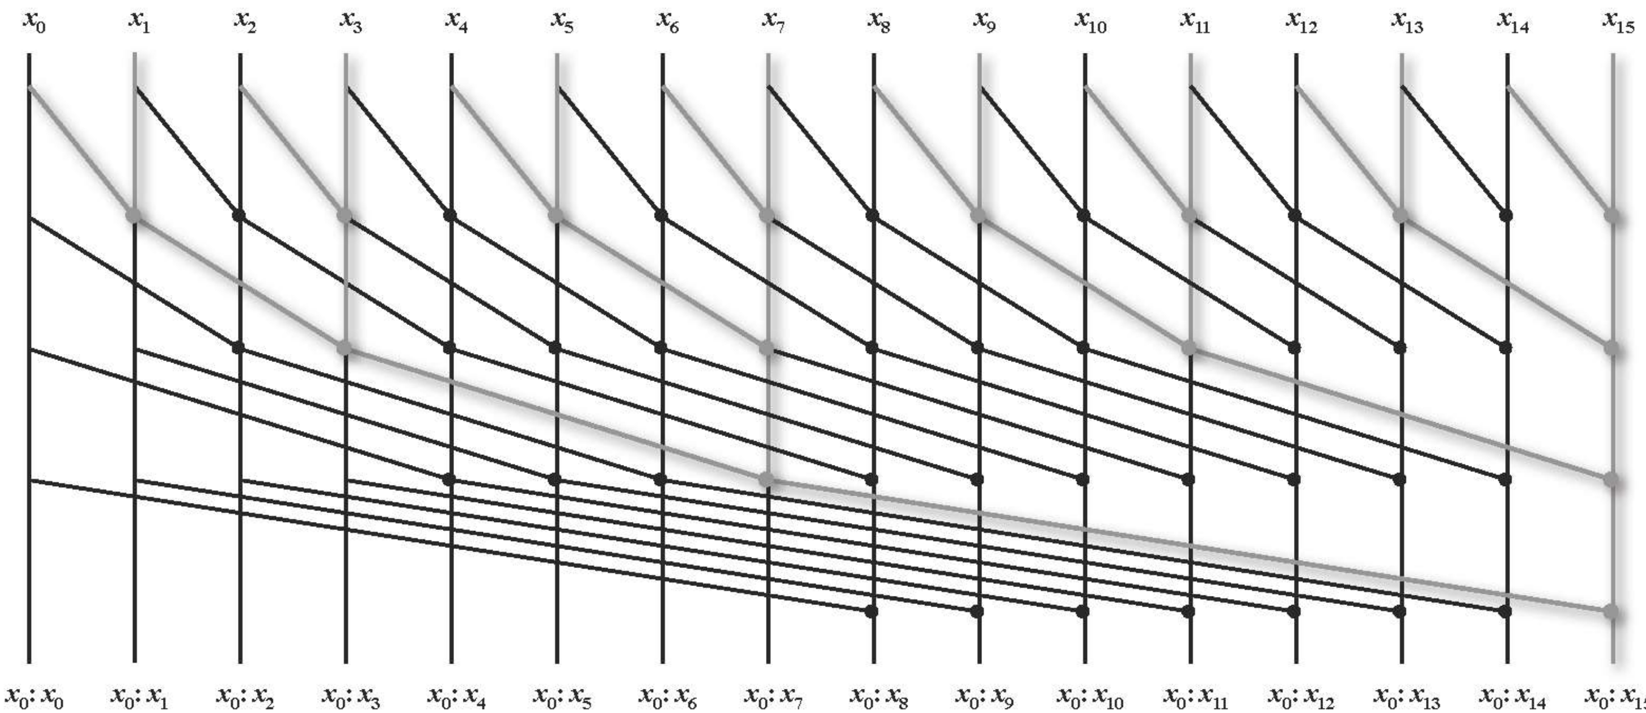

A Parallel Algorithm for the Efficient Solution of a General Class of Recurrence Equations, Kogge and Stone, 1973

See “carry lookahead” adders vs. “ripple carry” adders

$O(n \log n)$ Scan

Courtesy John Owens

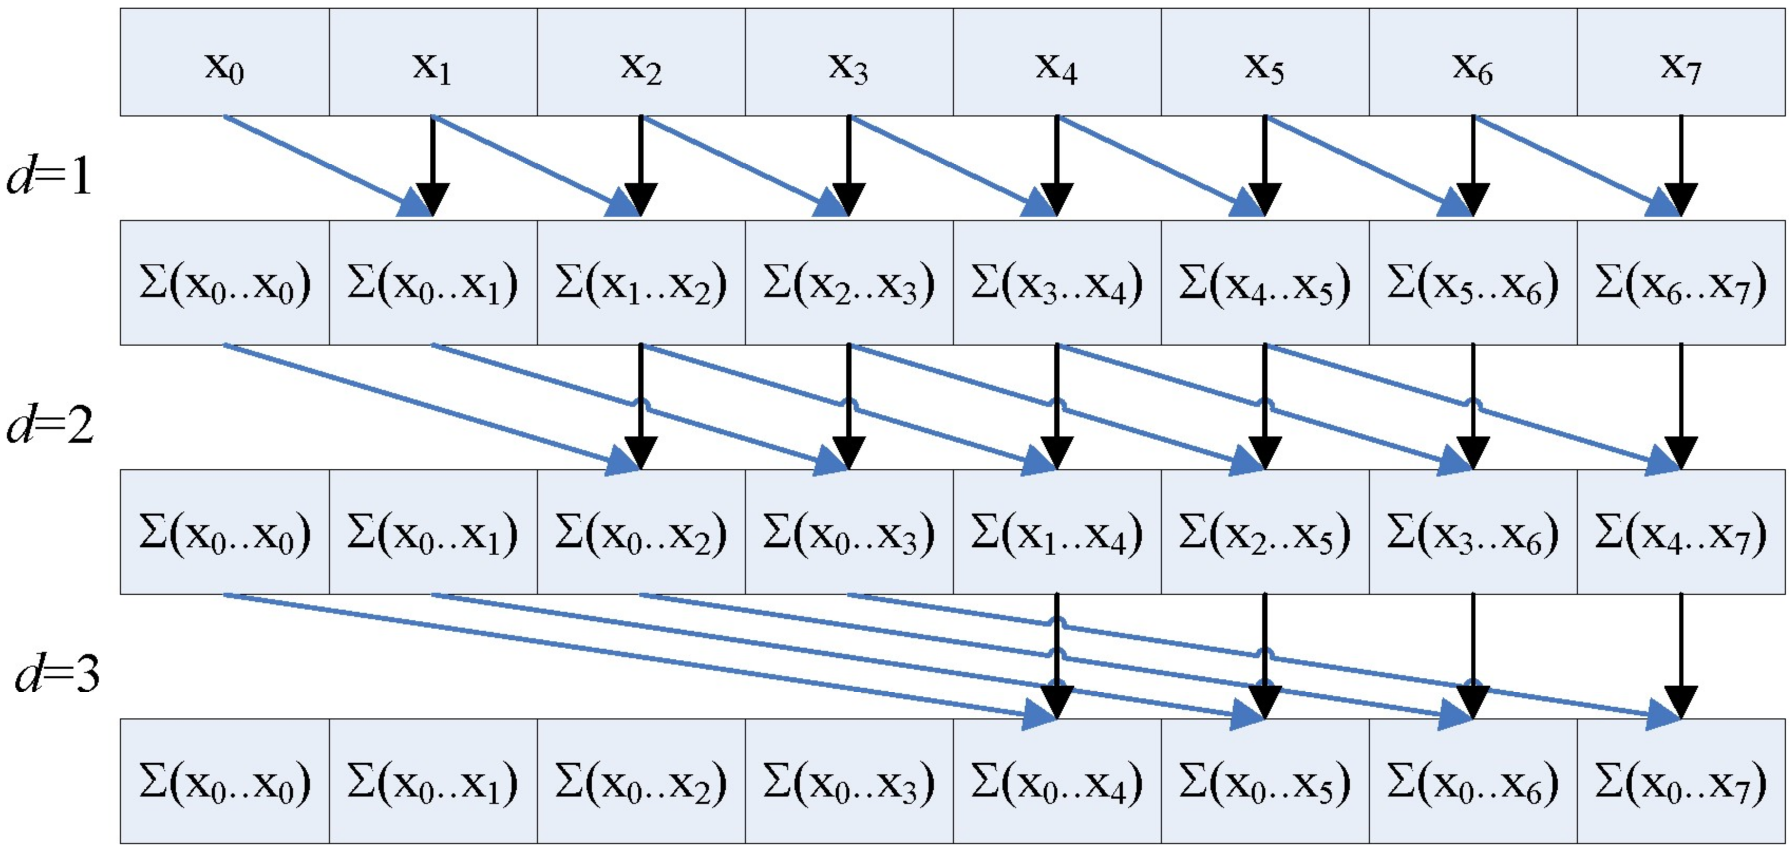

- Step efficient (log *n* steps)
- Not work efficient (*n* log *n* work)
- Requires barriers at each step (WAR dependencies)

Courtesy John Owens

Hillis-Steele Scan Implementation

No WAR conflicts, $O(2N)$ storage

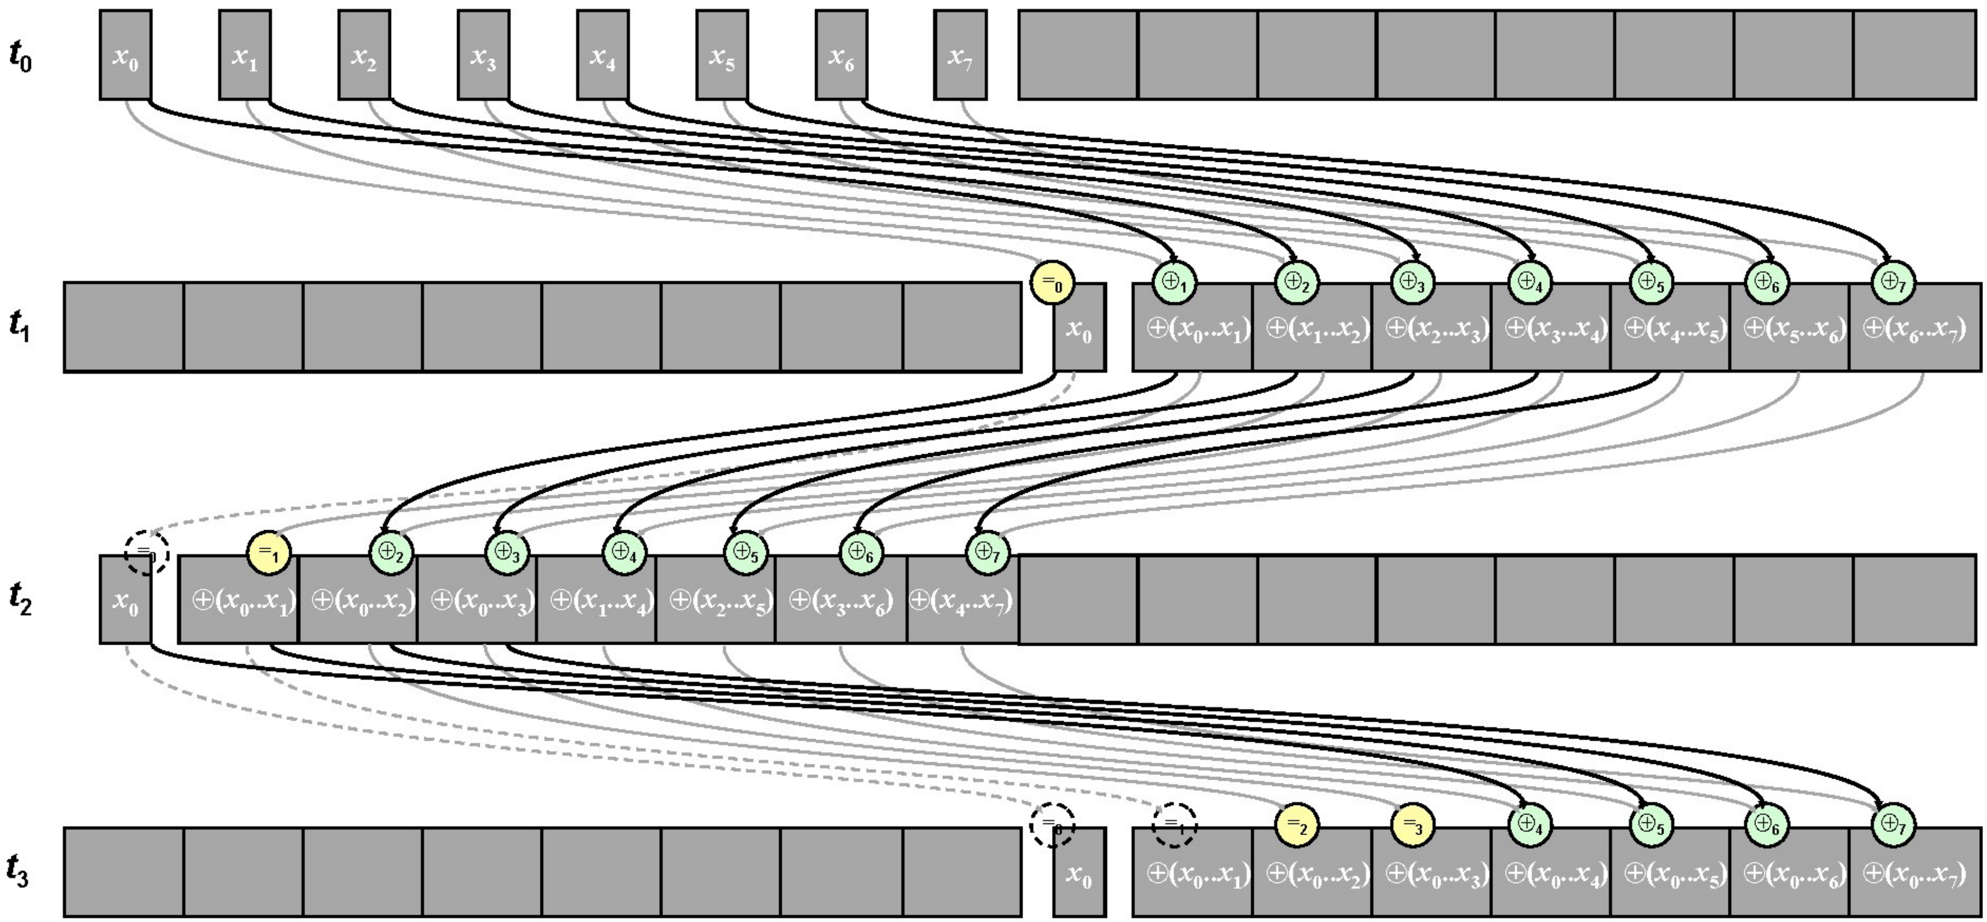

A First-Attempt Parallel Scan Algorithm

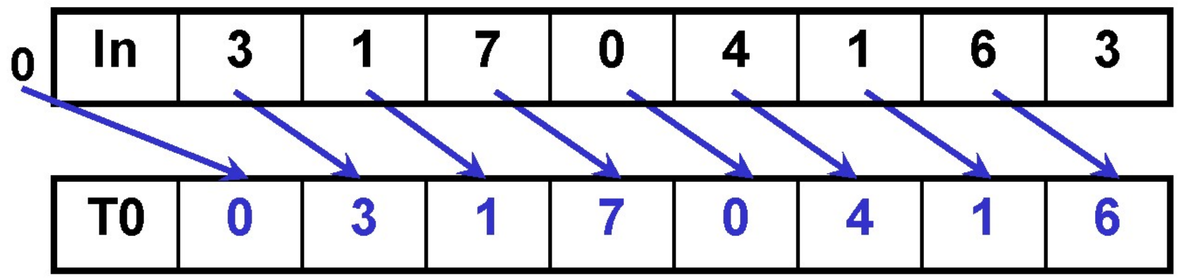

Each thread reads one value from the input array in device memory into shared memory array T0. Thread 0 writes 0 into shared memory array.

1. Read input from device memory to shared memory. Set first element to zero and shift others right by one.

A First-Attempt Parallel Scan Algorithm

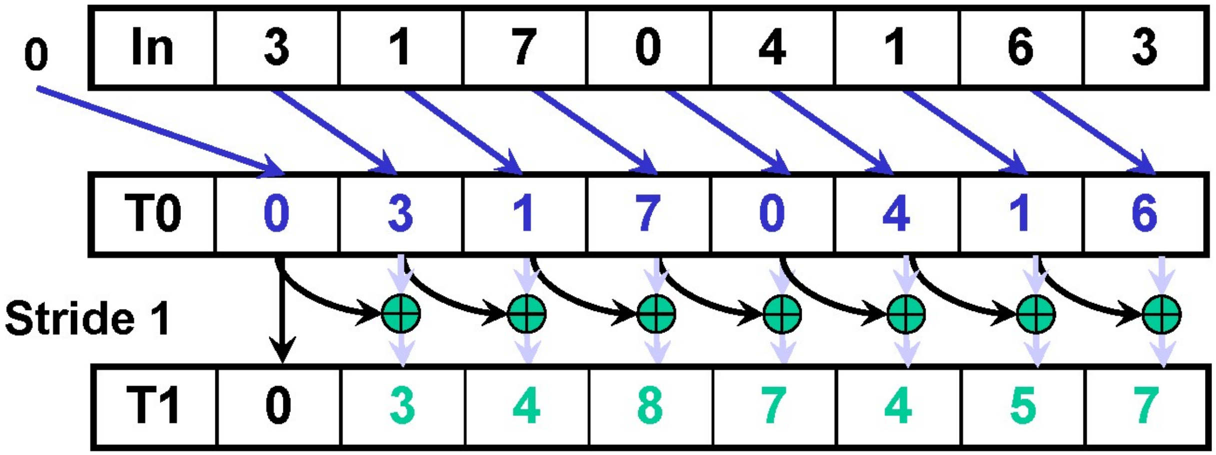

1. (previous slide)
2. Iterate log(n) times: Threads *stride* to *n:* Add pairs of elements *stride* elements apart. Double *stride* at each iteration. (note must double buffer shared mem arrays)

Iteration #1
Stride = 1

• Active threads: *stride* to *n*-1 (*n-stride* threads)
• Thread *j* adds elements *j* and *j-stride* from T0 and writes result into shared memory buffer T1 (ping-pong)

A First-Attempt Parallel Scan Algorithm

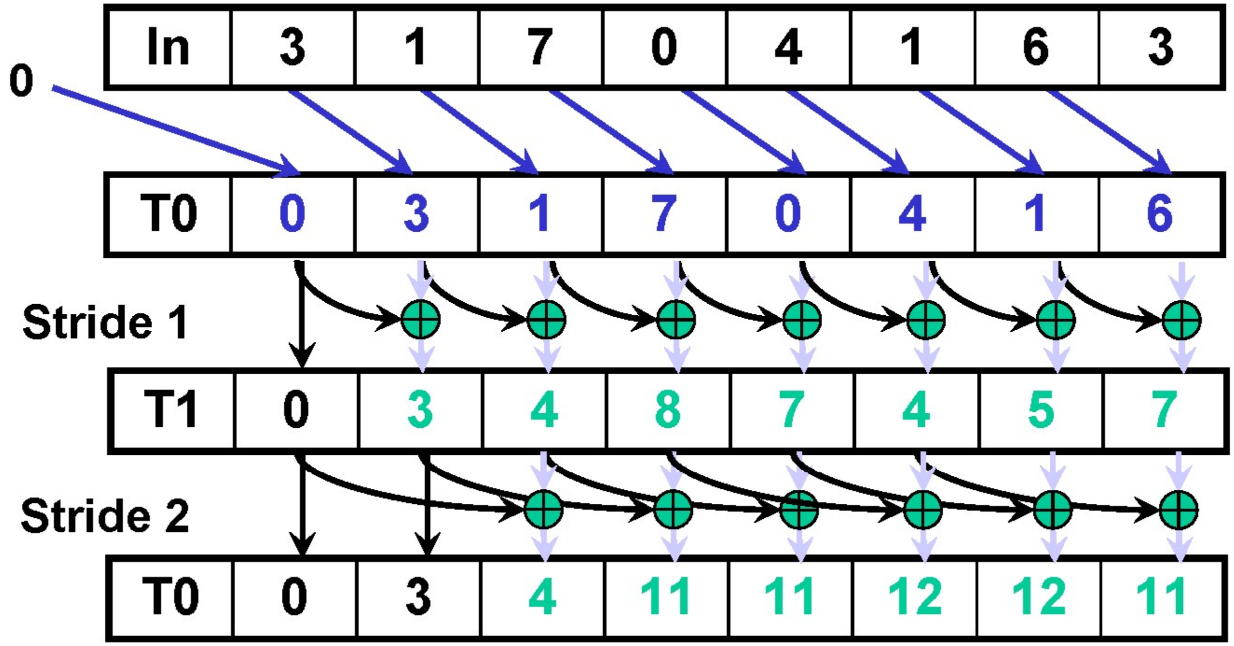

Iteration #2
Stride = 2

1. Read input from device memory to shared memory. Set first element to zero and shift others right by one.

2. Iterate log(n) times: Threads *stride* to *n:* Add pairs of elements *stride* elements apart. Double *stride* at each iteration. (note must double buffer shared mem arrays)

A First-Attempt Parallel Scan Algorithm

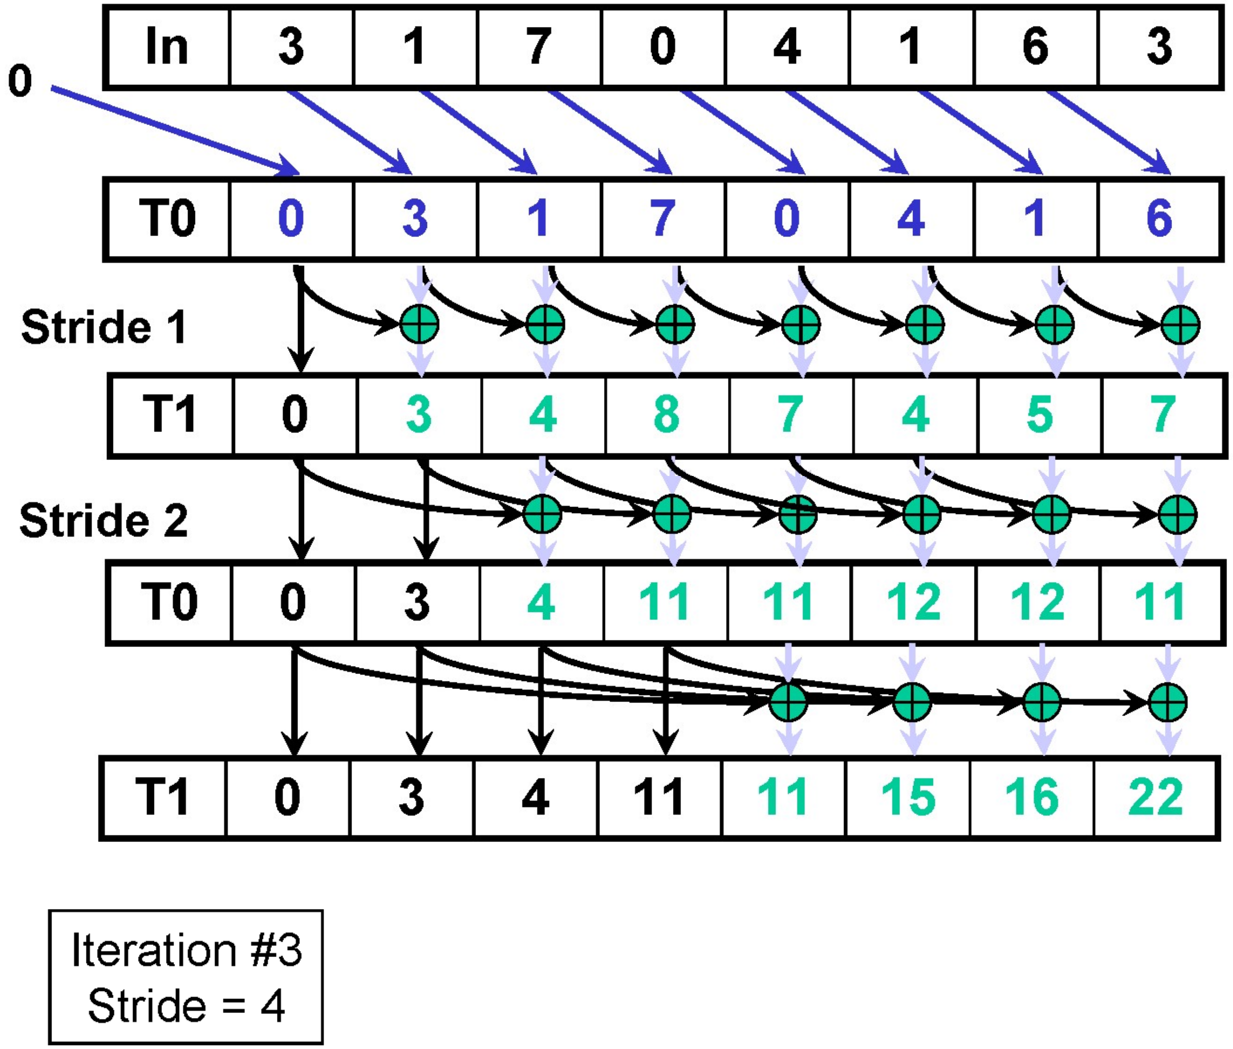

1. Read input from device memory to shared memory. Set first element to zero and shift others right by one.

2. Iterate log(n) times: Threads *stride* to *n:* Add pairs of elements *stride* elements apart. Double *stride* at each iteration. (note must double buffer shared mem arrays)

A First-Attempt Parallel Scan Algorithm

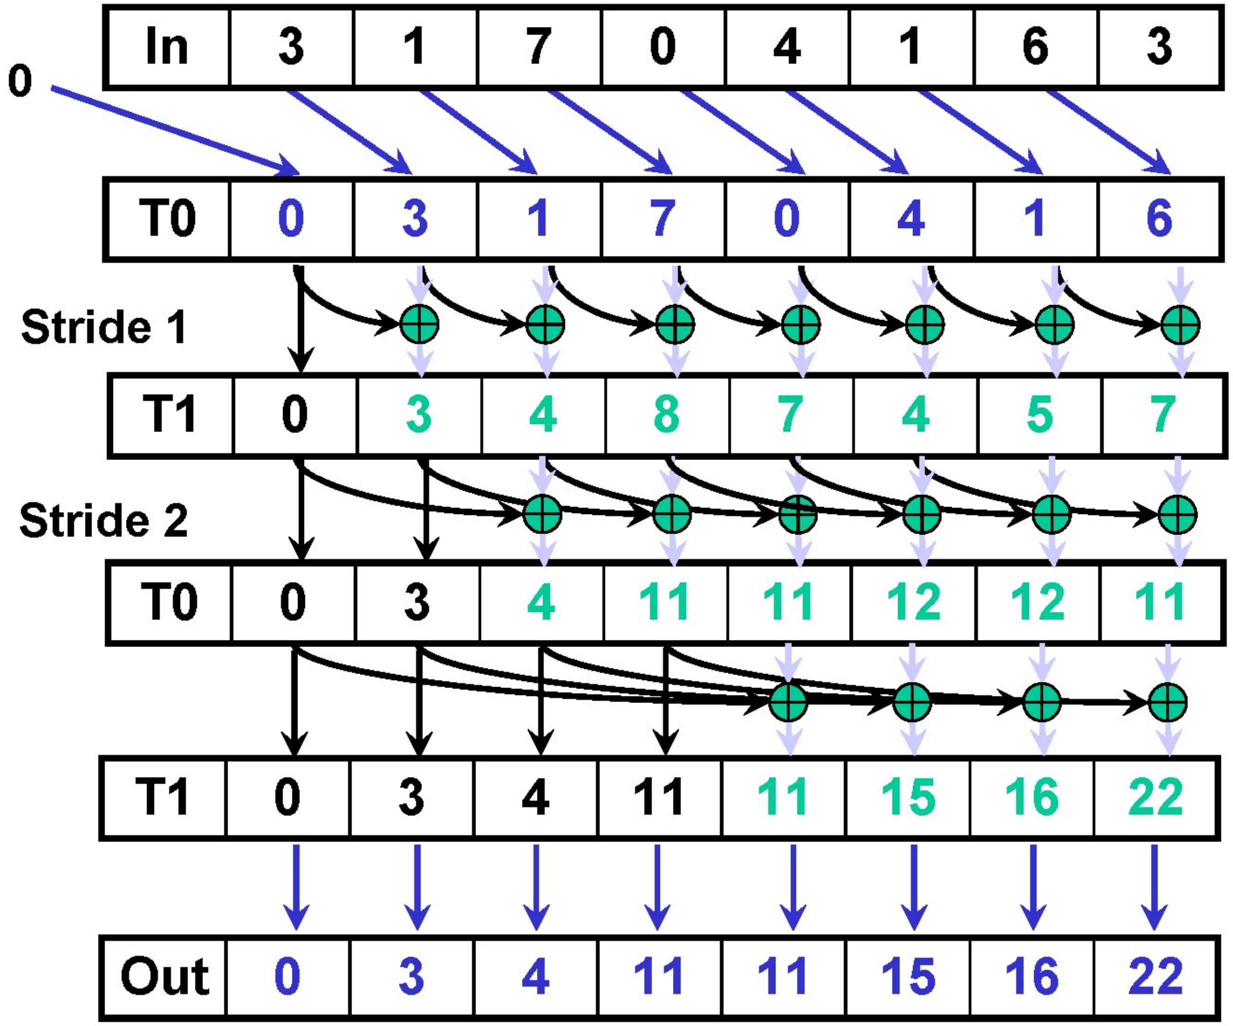

1. Read input from device memory to shared memory. Set first element to zero and shift others right by one.

2. Iterate log(n) times: Threads *stride* to *n:* Add pairs of elements *stride* elements apart. Double *stride* at each iteration. (note must double buffer shared mem arrays)

3. Write output to device memory.

Work Efficiency Considerations

- **The first-attempt Scan executes log(n) parallel iterations**
  - Total adds: n * (log(n) – 1) + 1 → O(n*log(n)) work
- **This scan algorithm is not very work efficient**
  - Sequential scan algorithm does *n* adds
  - A factor of log(n) hurts: 20x for 10^6 elements!
- **A parallel algorithm can be slow when execution resources are saturated due to low work efficiency**

Balanced Trees

- **For improving efficiency**
- **A common parallel algorithm pattern:**
  - Build a balanced binary tree on the input data and sweep it to and from the root
  - Tree is not an actual data structure, but a concept to determine what each thread does at each step
- **For scan:**
  - Traverse down from leaves to root building partial sums at internal nodes in the tree
    - Root holds sum of all leaves
  - Traverse back up the tree building the scan from the partial sums

Typical Parallel Programming Pattern

- **2 log(n) steps**

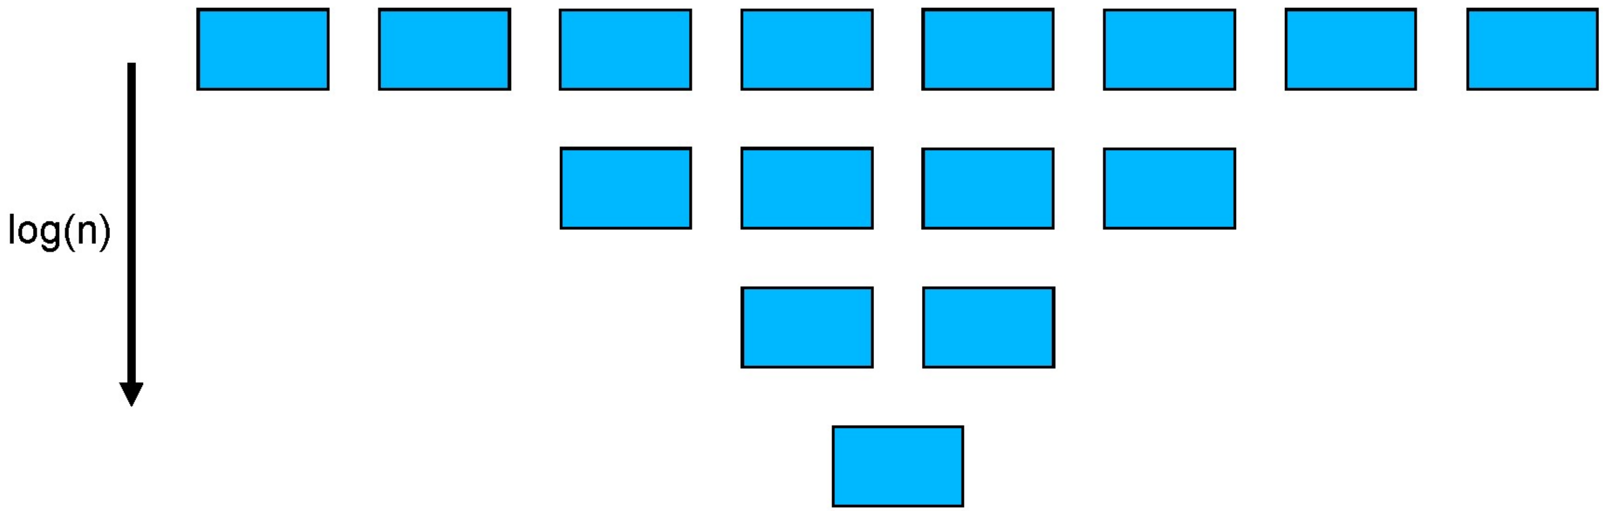

Typical Parallel Programming Pattern

- **2 log(n) steps**

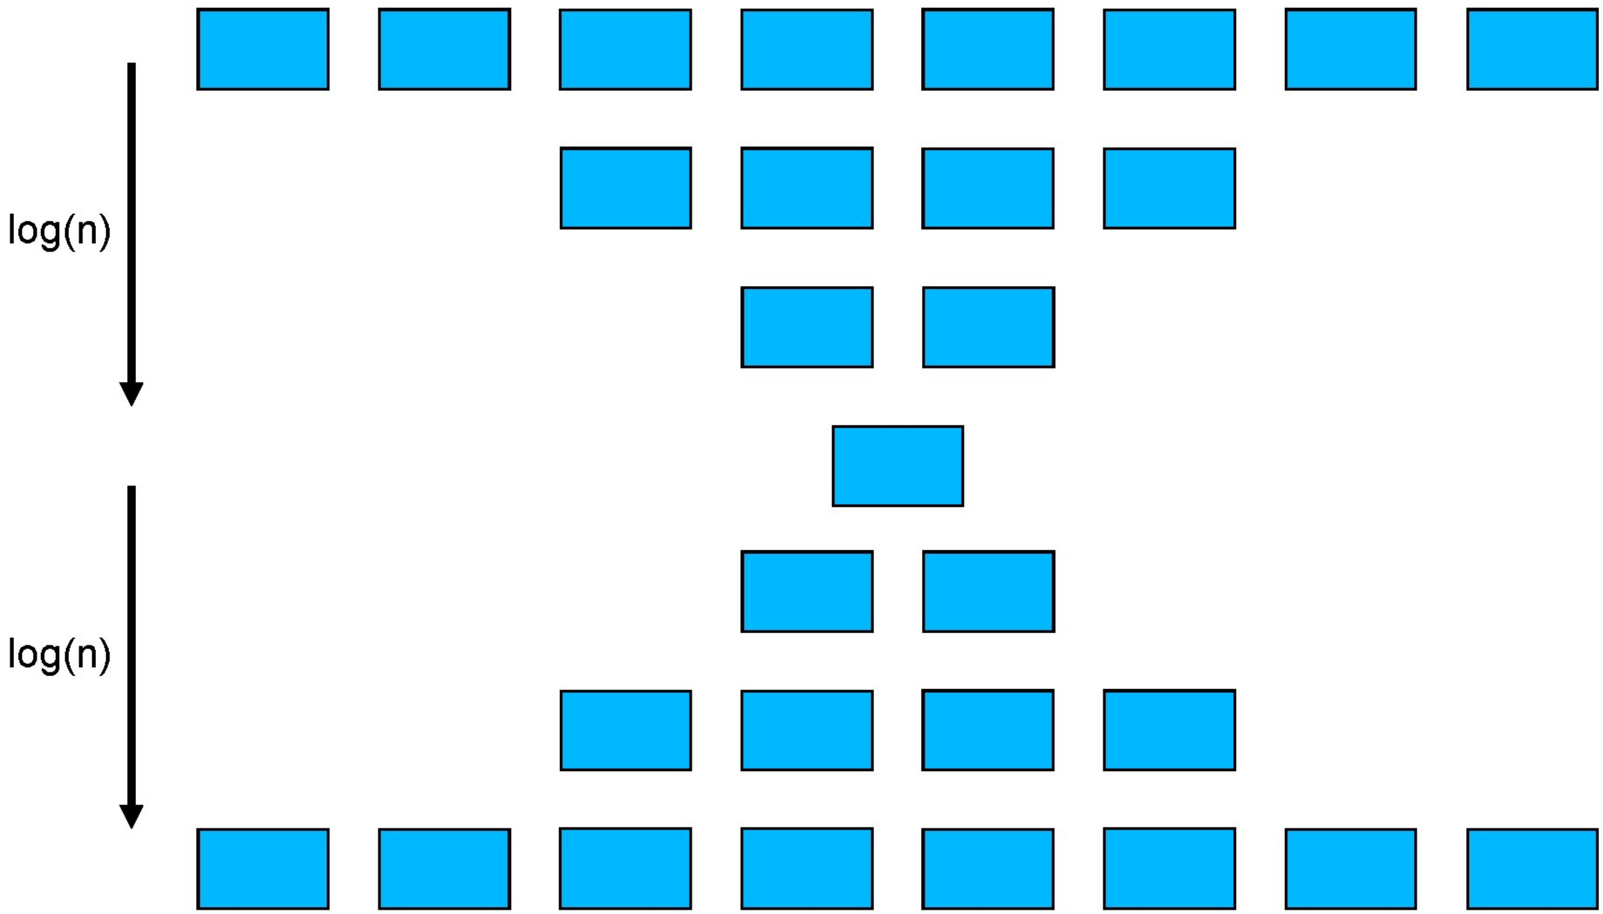

Courtesy John Owens

Brent Kung Scan

Circuit family

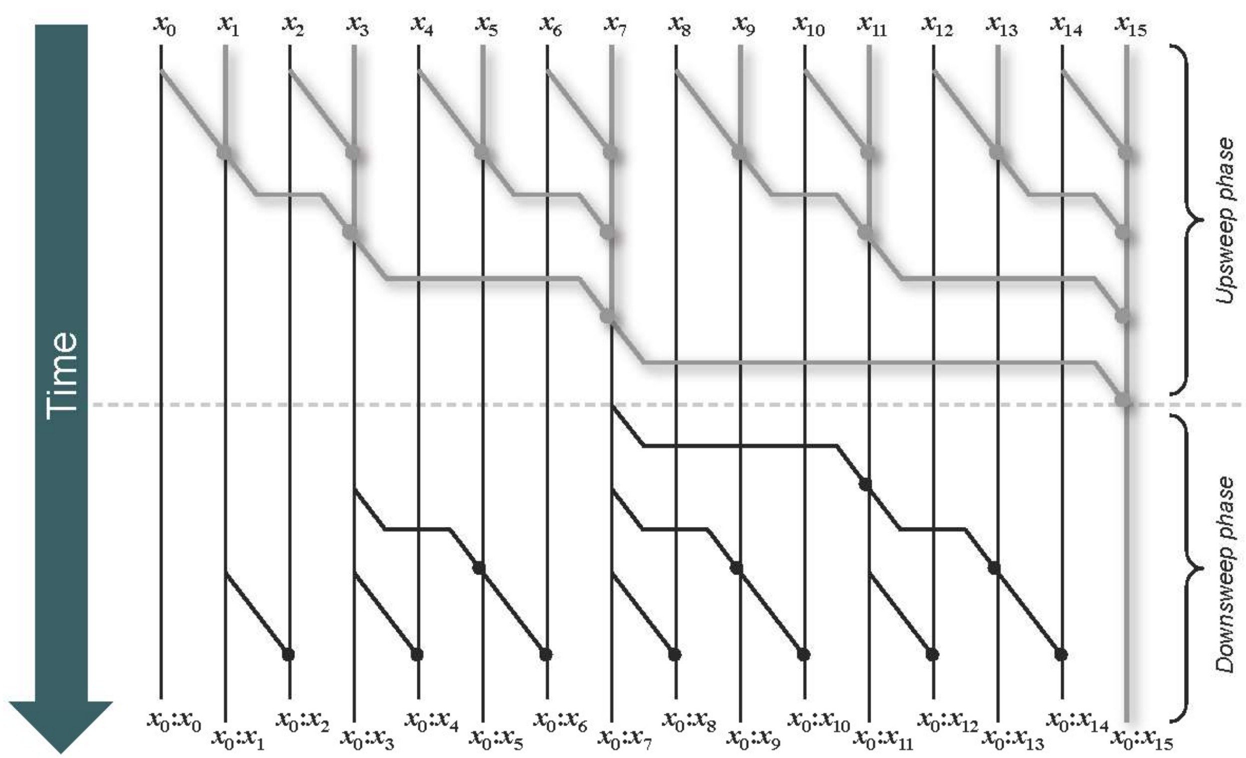

A Regular Layout for Parallel Adders, Brent and Kung, 1982

$O(n)$ Scan [Blelloch]

Courtesy John Owens

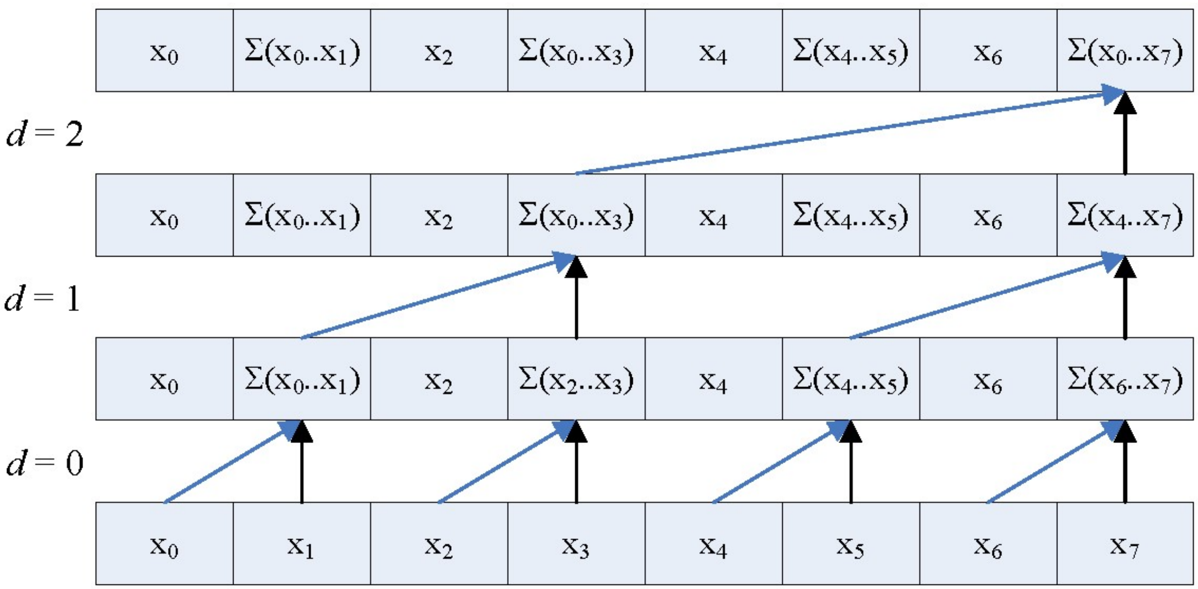

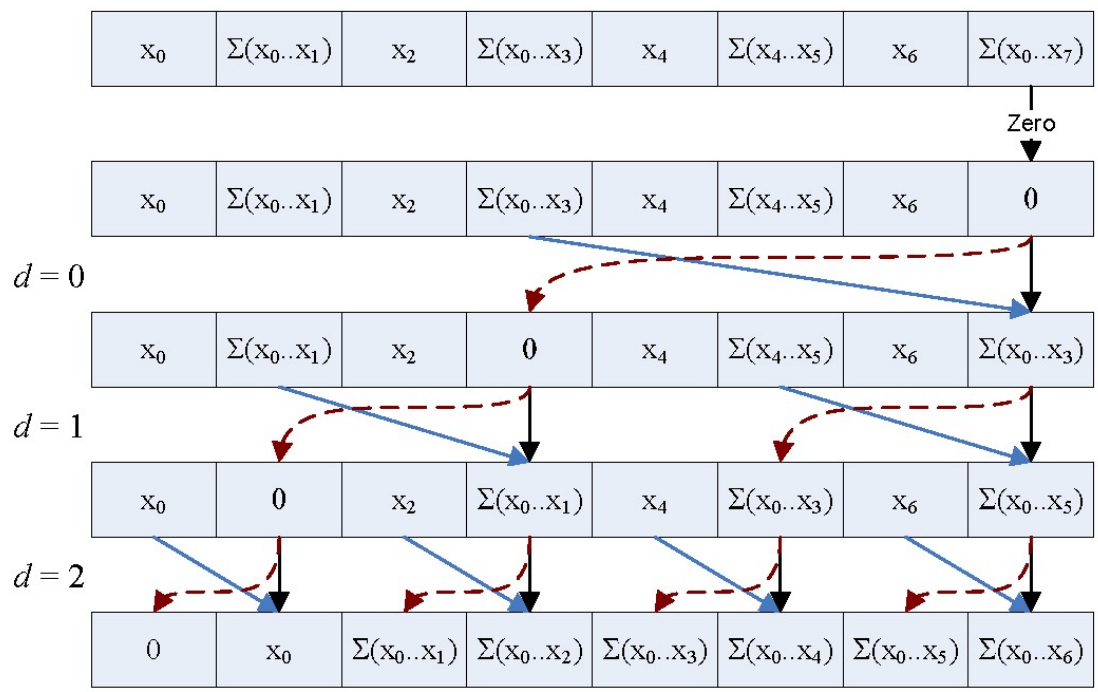

- Work efficient ($O(n)$ work)
- Bank conflicts, and lots of 'em

Build the Sum Tree

T	3	1	7	0	4	1	6	3

Assume array is already in shared memory

Build the Sum Tree

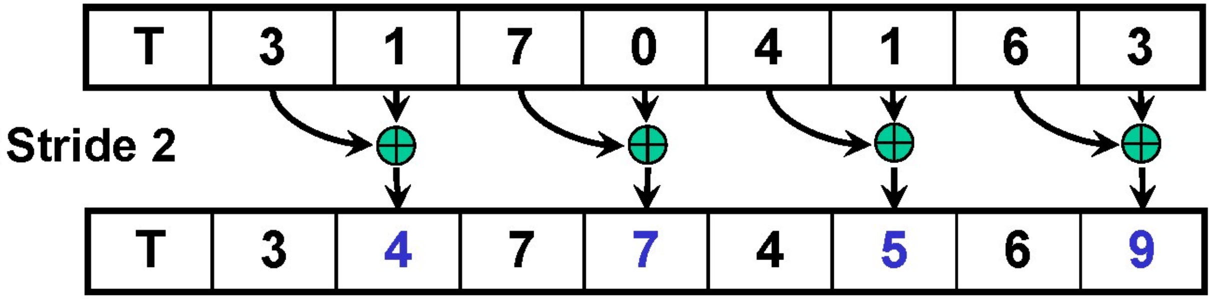

Iteration 1, *n*/2 threads

Each ⊕ corresponds to a single thread.

Iterate log(n) times. Each thread adds value *stride* / 2 elements away to its own value.

Build the Sum Tree

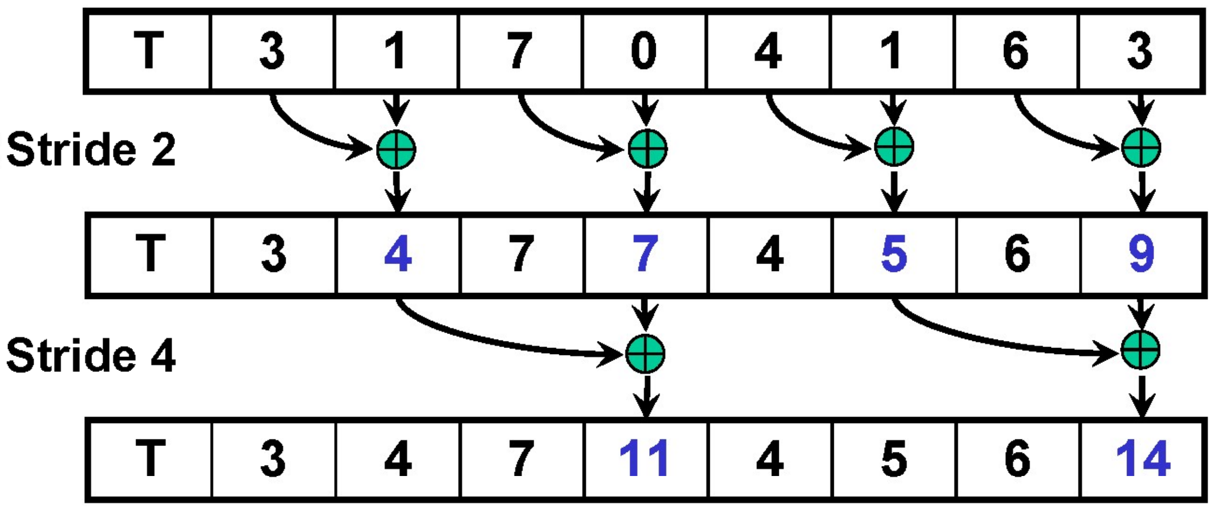

Iteration 2, *n*/4 threads

Each ⊕ corresponds to a single thread.

Iterate log(n) times. Each thread adds value *stride* / 2 elements away to its own value.

Build the Sum Tree

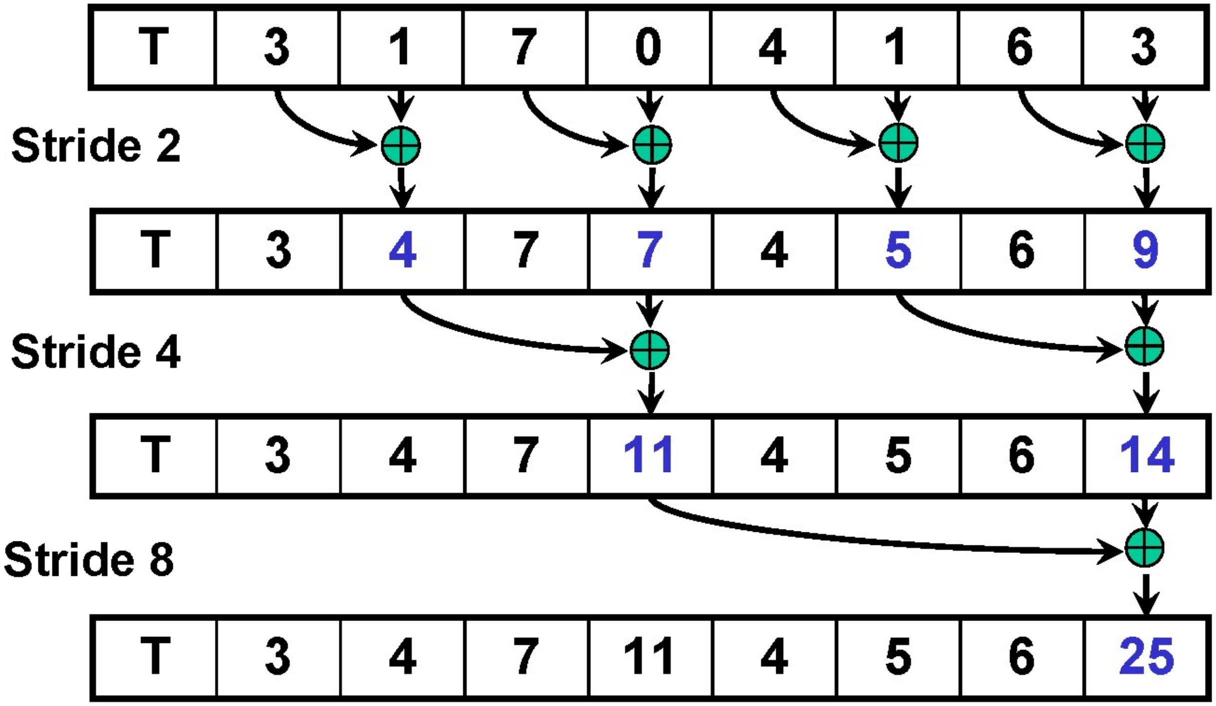

Iteration log(*n*), 1 thread

Each ⊕ corresponds to a single thread.

Iterate log(n) times. Each thread adds value *stride* / 2 elements away to its own value.

Note that this algorithm operates in-place: no need for double buffering

Down-Sweep Variant 1: Exclusive Scan

T	3	4	7	11	4	5	6	0

We now have an array of partial sums. Since this is an exclusive scan, set the last element to zero. It will propagate back to the first element.

Build Scan From Partial Sums

T	3	4	7	11	4	5	6	0

Build Scan From Partial Sums

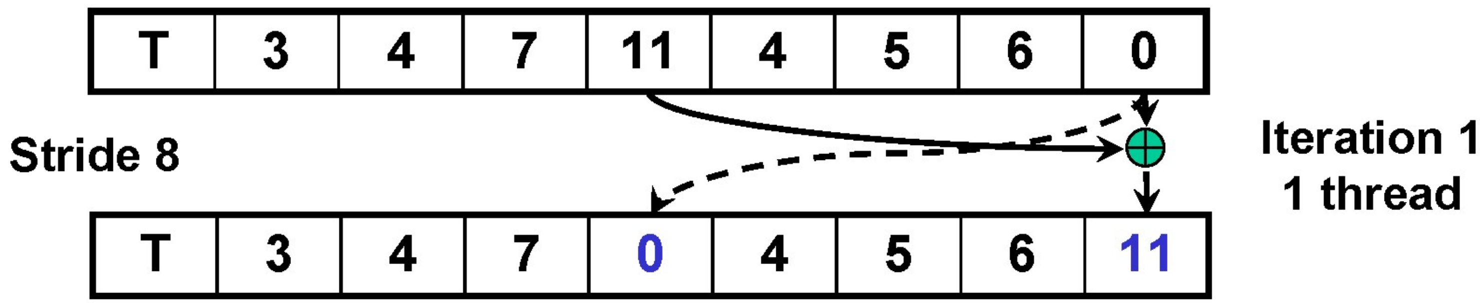

Each ⊕ corresponds to a single thread.

Iterate log(n) times. Each thread adds value *stride* / 2 elements away to its own value. and sets the value *stride* elements away to its own *previous* value.

Build Scan From Partial Sums

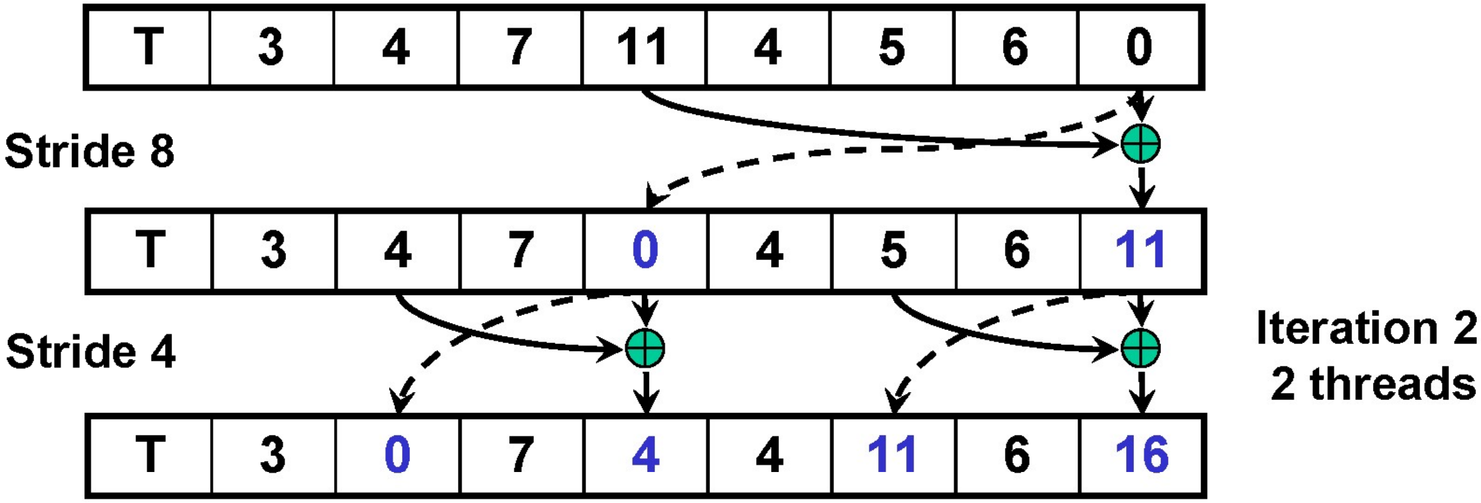

Each ⊕ corresponds to a single thread.

Iterate log(n) times. Each thread adds value *stride* / 2 elements away to its own value. and sets the value *stride* / 2 elements away to its own *previous* value.

Build Scan From Partial Sums

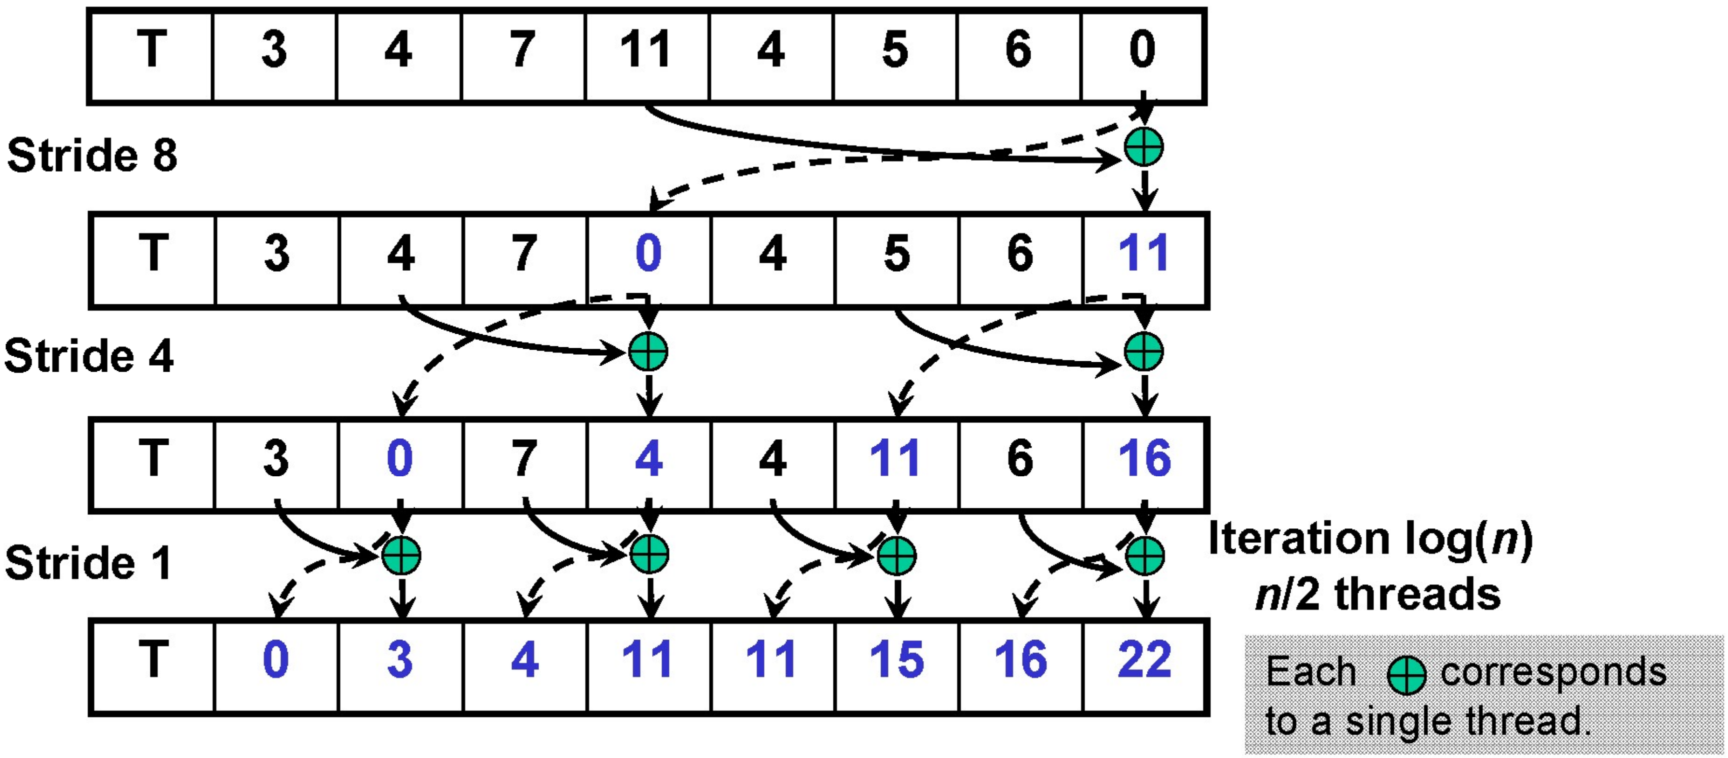

Done! We now have a completed scan that we can write out to device memory.

Total steps: 2 * log(n).
Total work: 2 * (n-1) adds = $O(n)$ **Work Efficient!**

Down-Sweep Variant 2: Inlusive Scan

T	3	4	7	11	4	5	6	25

We now have an array of partial sums. Let’s propagate the sums back.

Build Scan From Partial Sums

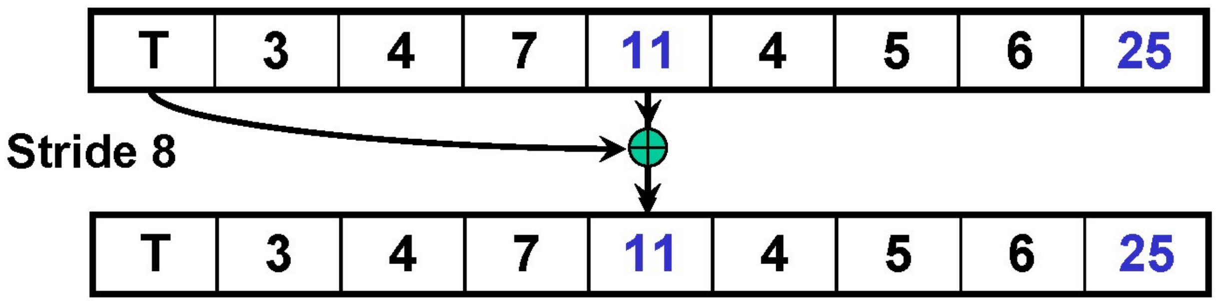

no operation

Each $\oplus$ corresponds to a single thread.

Iterate log(n) times. Each thread adds value *stride* / 2 elements away to its own value. First element adds zero.

Build Scan From Partial Sums

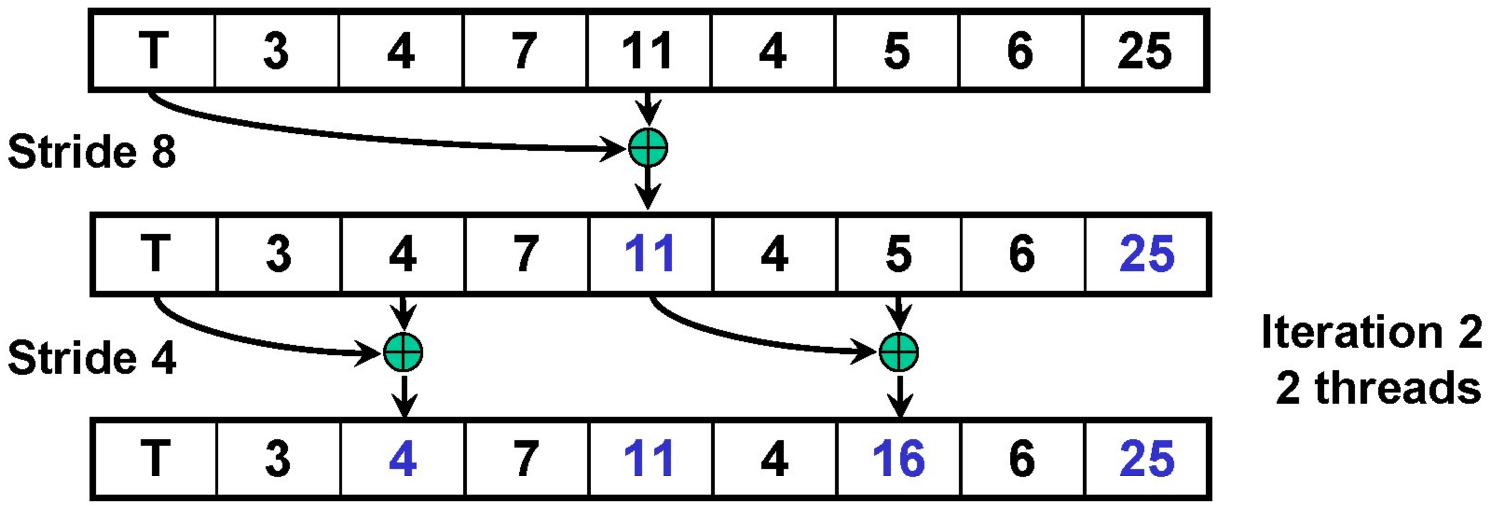

Each ⊕ corresponds to a single thread.

Iterate log(n) times. Each thread adds value *stride* / 2 elements away to its own value. First element adds zero.

Build Scan From Partial Sums

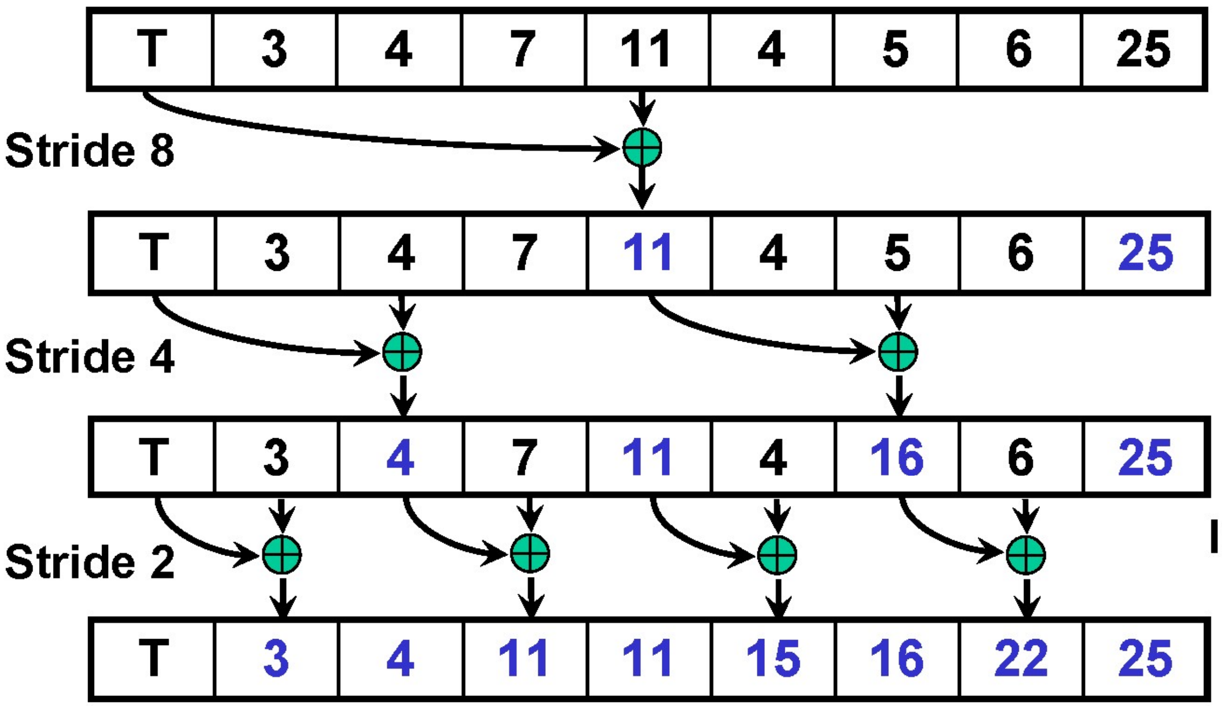

Iteration log(*n*)
***n*/2 threads**

Each ⊕ corresponds to a single thread.

Done! We now have a completed scan that we can write out to device memory.

Total steps: 2 * log(*n*).
Total work: < 2 * (*n*-1) adds = *O*(*n*) **Work Efficient!**

Application to Large Arrays

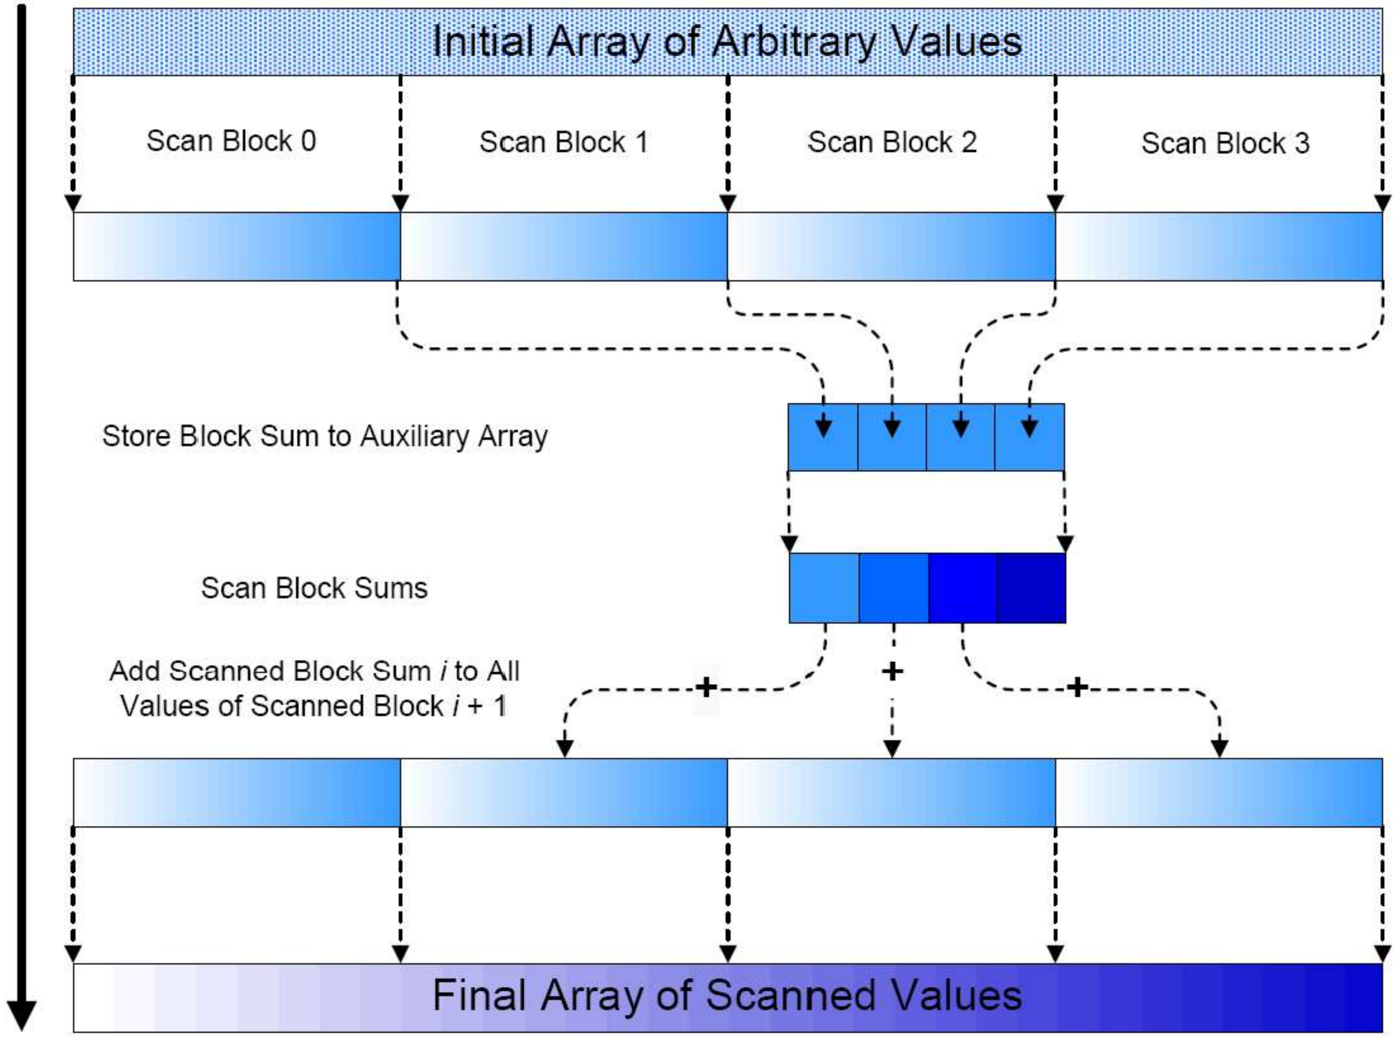

[Mark Harris]

Scan papers

Courtesy John Owens

- Daniel Horn, Stream Reduction Operations for GPGPU Applications, GPU Gems 2, Chapter 36, pp. 573–589, March 2005.
- Shubhabrata Sengupta, Aaron E. Lefohn, and John D. Owens. A Work-Efficient Step-Efficient Prefix Sum Algorithm. In Proceedings of the 2006 Workshop on Edge Computing Using New Commodity Architectures, pages D–26–27, May 2006
- Mark Harris, Shubhabrata Sengupta, and John D. Owens.Parallel Prefix Sum (Scan) with CUDA. In Hubert Nguyen, editor, GPU Gems 3, chapter 39, pages 851–876. Addison Wesley, August 2007.
- Shubhabrata Sengupta, Mark Harris, Yao Zhang, and John D. Owens. Scan Primitives for GPU Computing. In Graphics Hardware 2007, pages 97–106, August 2007.
- Y. Dotsenko, N. K. Govindaraju, P. Sloan, C. Boyd, and J. Manferdelli, "Fast scan algorithms on graphics processors," in ICS '08: Proceedings of the 22nd Annual International Conference on Supercomputing, 2008, pp. 205–213.
- Shubhabrata Sengupta, Mark Harris, Michael Garland, and John D. Owens. Efficient Parallel Scan Algorithms for many-core GPUs. In Jakub Kurzak, David A. Bader, and Jack Dongarra, editors, Scientific Computing with Multicore and Accelerators, Chapman & Hall/CRC Computational Science, chapter 19, pages 413–442. Taylor & Francis, January 2011.
- D. Merrill and A. Grimshaw, Parallel Scan for Stream Architectures. Technical Report CS2009-14, Department of Computer Science, University of Virginia, 2009, 54pp.
- Shengen Yan, Guoping Long, and Yunquan Zhang. 2013. StreamScan: fast scan algorithms for GPUs without global barrier synchronization. In Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '13). ACM, New York, NY, USA, 229-238.

Thank you.

- Hendrik Lensch, Robert Strzodka
- John Owens
- NVIDIA