

departamento de informática
FACULDADE DE CIÊNCIAS E TECNOLOGIA
UNIVERSIDADE NOVA DE LISBOA

Map and Reduce Patterns

Concurrency and Parallelism — 2019-20

Master in Computer Science

(Mestrado Integrado em Eng. Informática)

Joao Lourenço <joao.lourenco@fct.unl.pt>

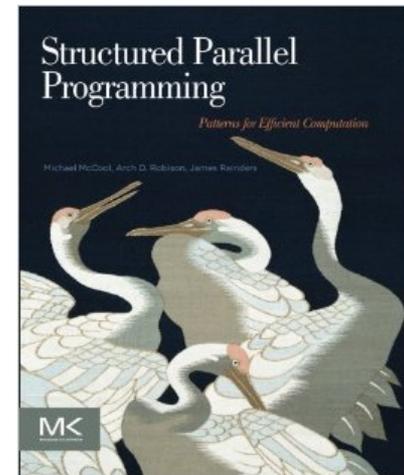
Source: Parallel Computing, CIS 410/510, Department of Computer and Information Science

Outline

- Map pattern
 - Optimizations
 - sequences of Maps
 - code Fusion
 - cache Fusion
 - Related Patterns
 - Example: Scaled Vector Addition (SAXPY)
- Reduce
 - Example: Dot Product

– Bibliography:

- **Chapters 4 and 5** of book McCool M., Arch M., Reinders J.; Structured Parallel Programming: Patterns for Efficient Computation; Morgan Kaufmann (2012); ISBN: 978-0-12-415993-8



Mapping

- “Do the same thing many times”

```
foreach i in foo:
```

```
    do_something (i)
```

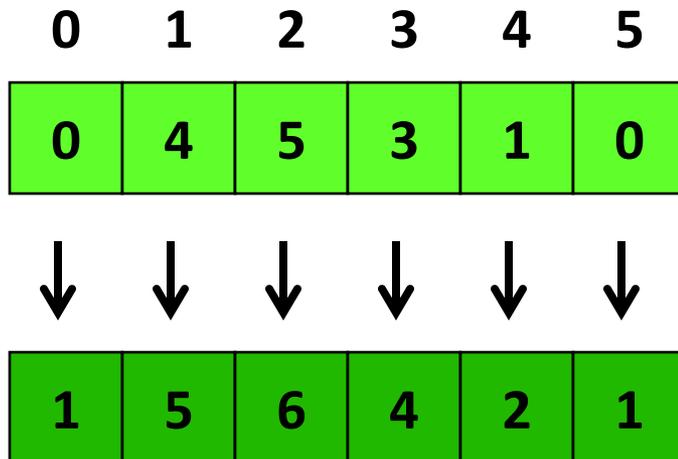
- Well-known higher order function in languages like ML, Haskell, Scala

$$\text{map} : \forall ab.(a \rightarrow b) \text{List}\langle a \rangle \rightarrow \text{List}\langle b \rangle$$

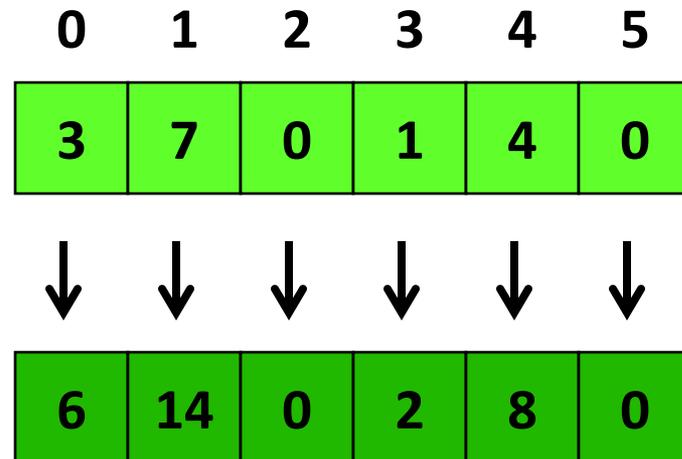
applies a function to each element in a list and returns a list of results

Example Maps

Add 1 to every item in an array



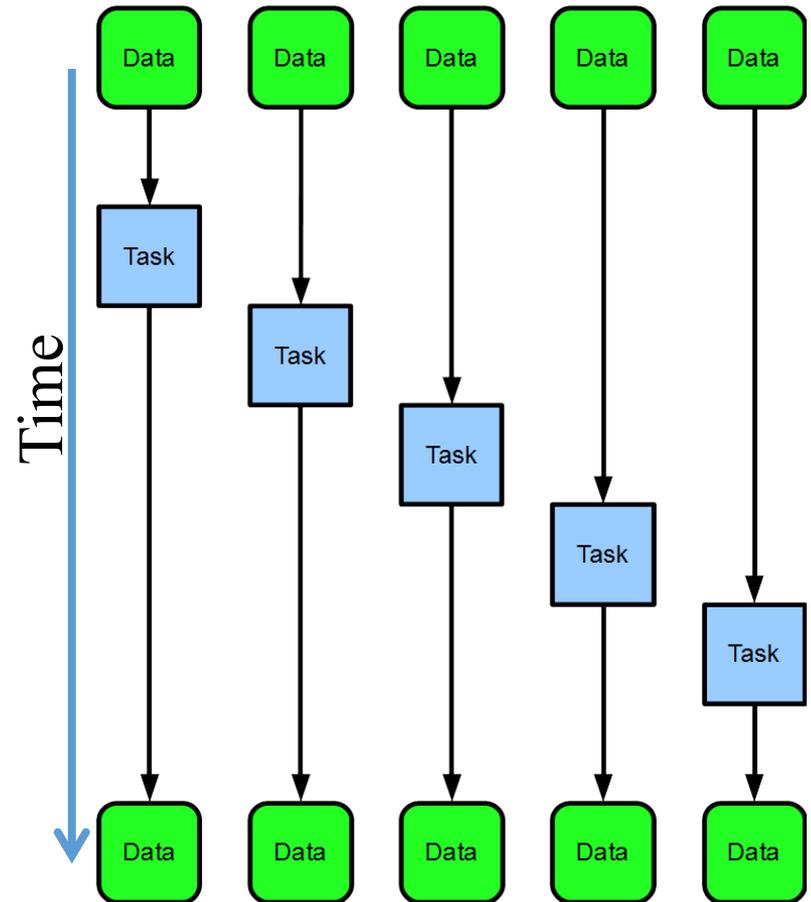
Double every item in an array



Key Point: An operation is a map if it can be applied to each element without knowledge of its neighbors.

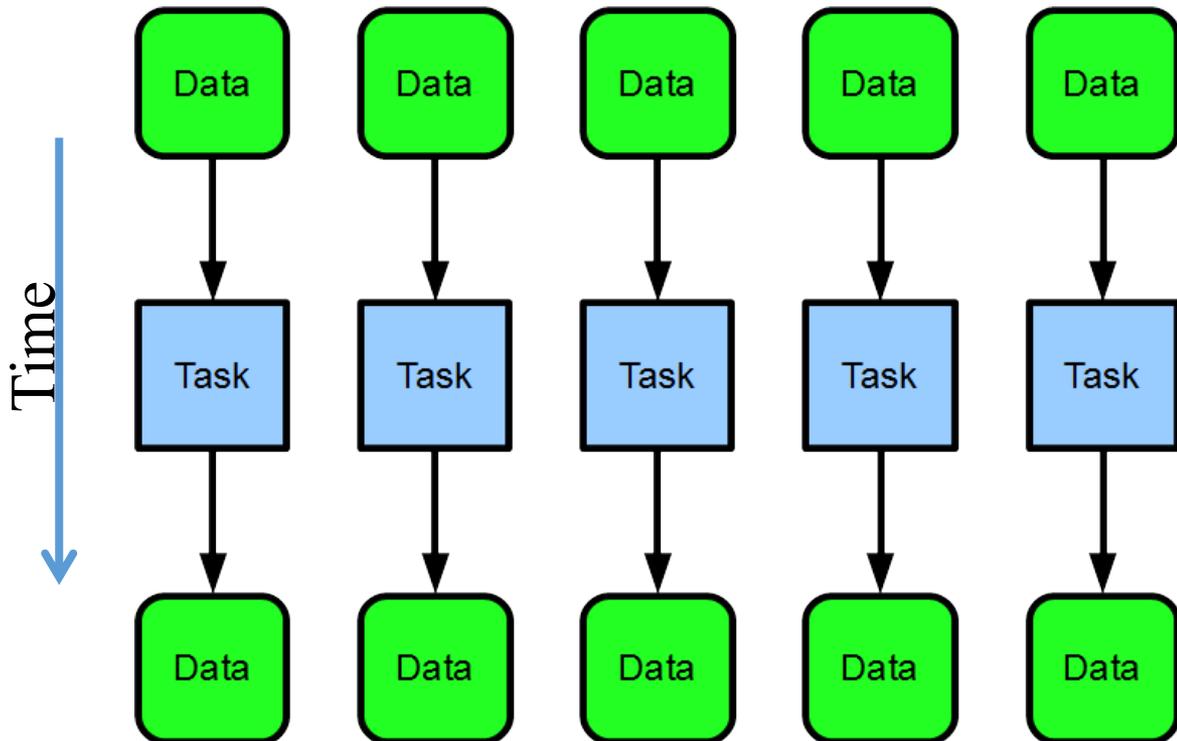
Sequential Map

```
for(int n=0; n < array.length; ++n) {  
    process(array[n]);  
}
```



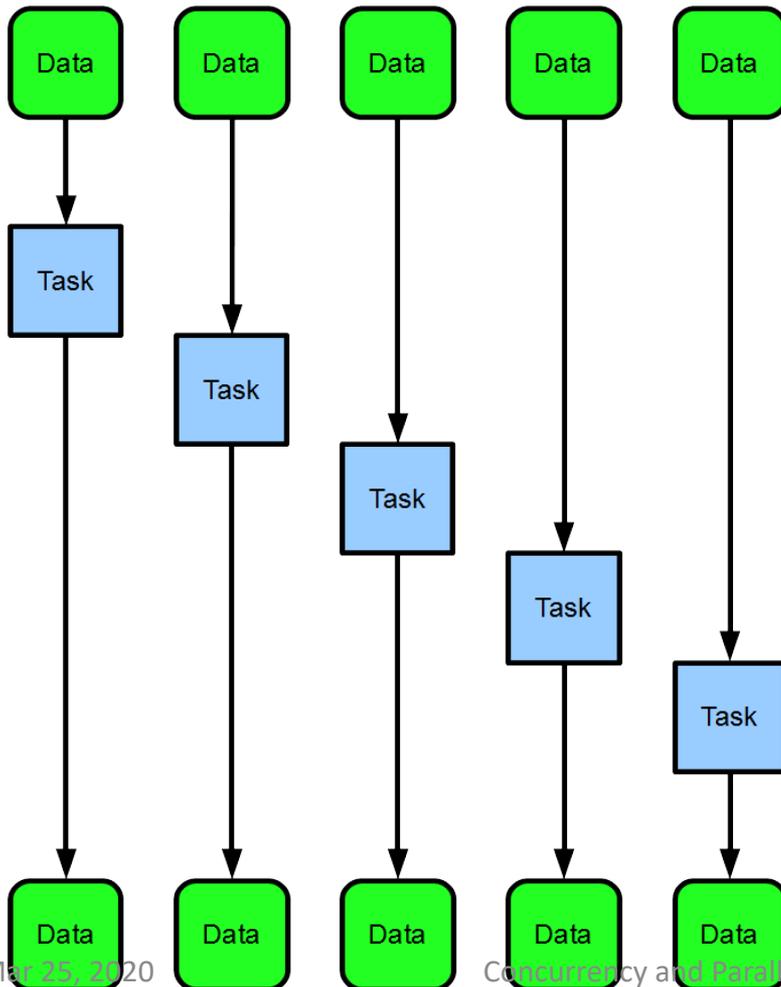
Parallel Map

```
parallel_for_each(x in array) {  
    process(x);  
}
```

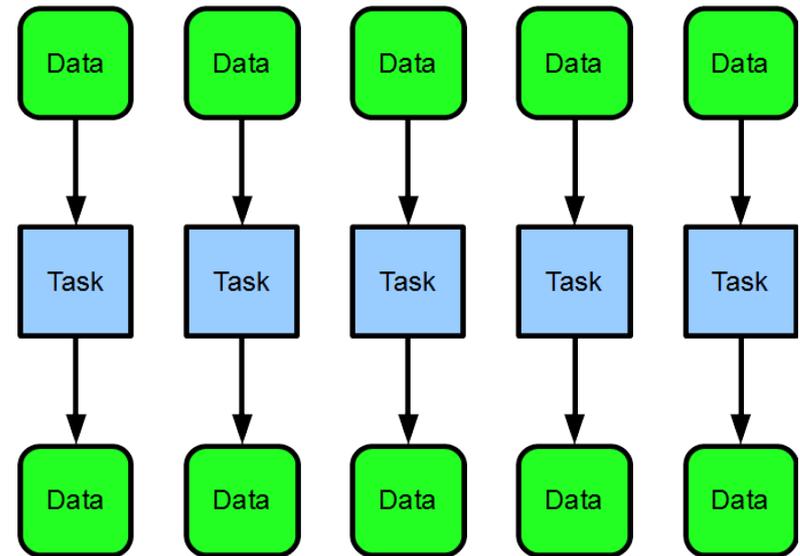


Comparing Maps

Serial Map



Parallel Map



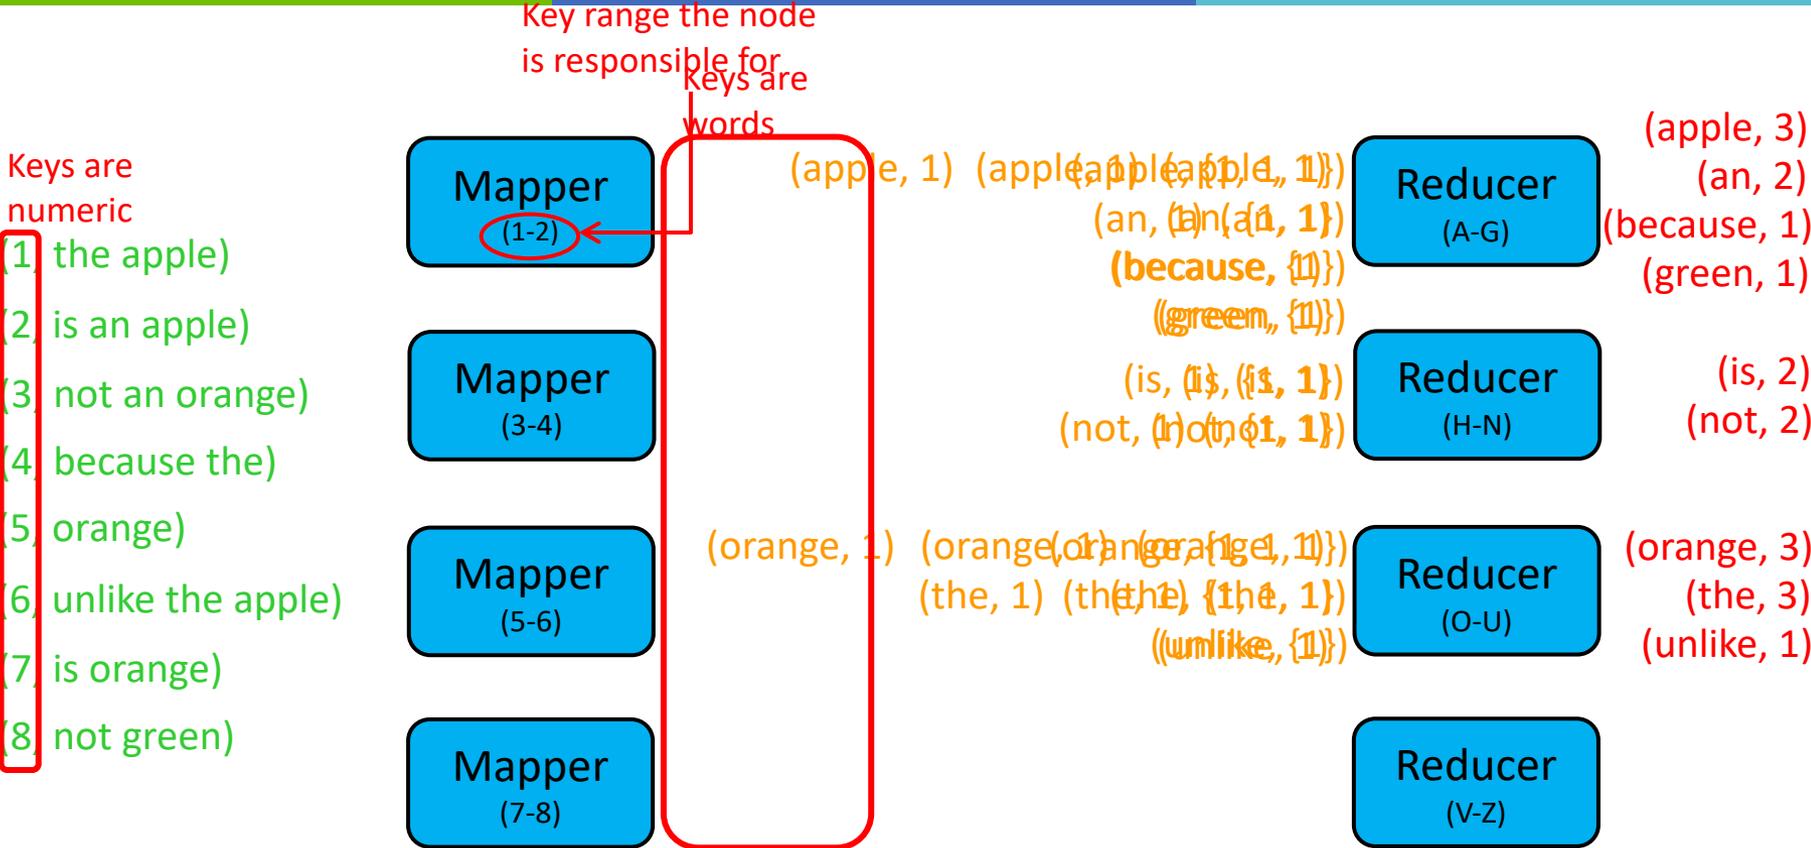
Time



Speedup

The space here is speedup. With the parallel map, the program finished execution early, while the serial map is still running.

Simple example: Word count



- 1 Each mapper receives some of the KV-pairs as input
- 2 The mappers process the KV-pairs one by one
- 3 Each KV-pair output by the mapper is sent to the reducer that is responsible for it
- 4 The infrastructure sort their input by key and group it
- 5 The reducers process their input one group at a time

Independence

- The key to (embarrassing) parallelism is independence

Warning: No shared state!

Map function should be “pure” (or “pure-ish”) and should not modify shared states

- Modifying shared state breaks perfect independence
- Results of accidentally violating independence:
 - non-determinism
 - data-races
 - undefined behavior
 - segfaults

Implementation and API

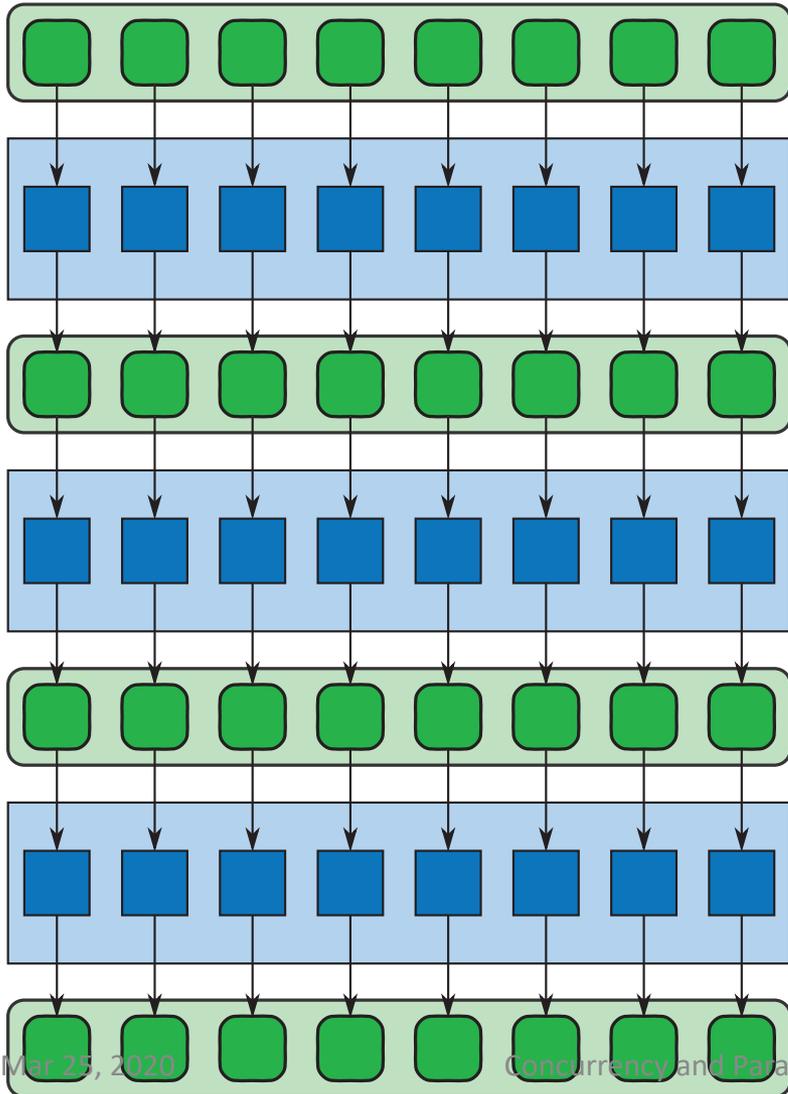
- OpenMP and CilkPlus contain a parallel **for** language construct
- Map is a mode of use of parallel **for**
- Some languages (CilkPlus, Matlab, Fortran) provide **array notation** which makes some maps more concise

Array Notation

```
A[:] = A[:] * 5;
```

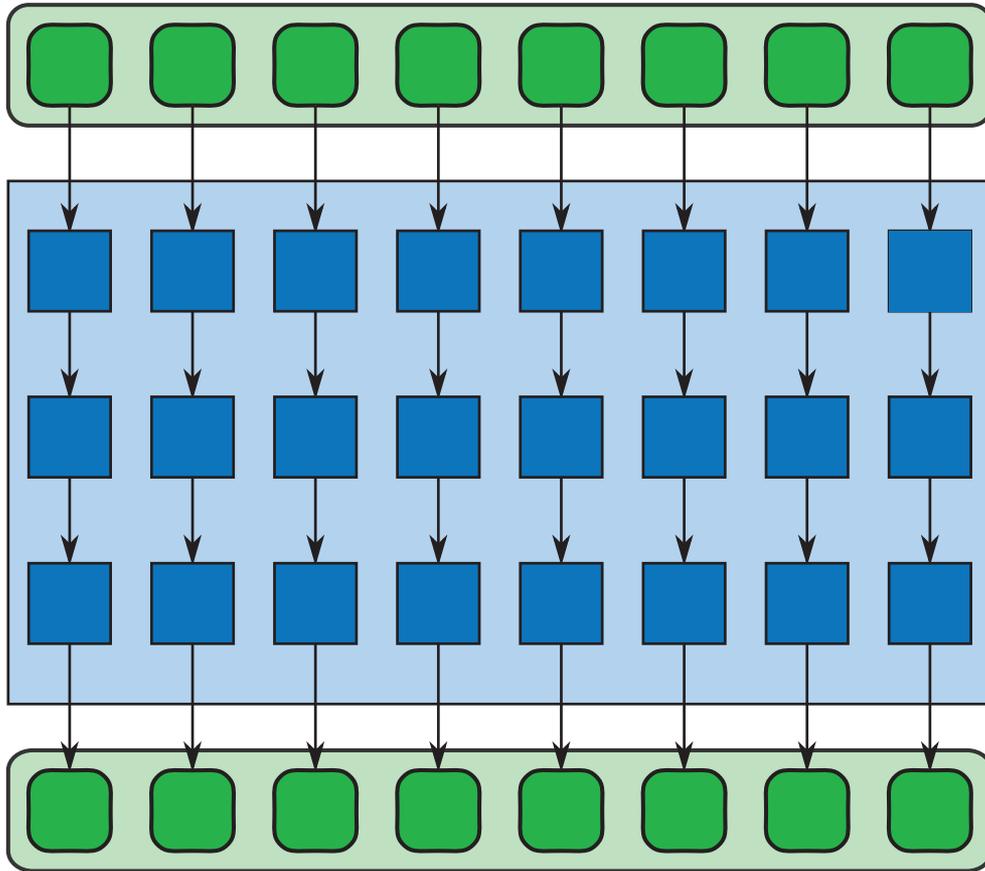
is CilkPlus array notation for “multiply every element in A by 5”

Optimization – Sequences of Maps



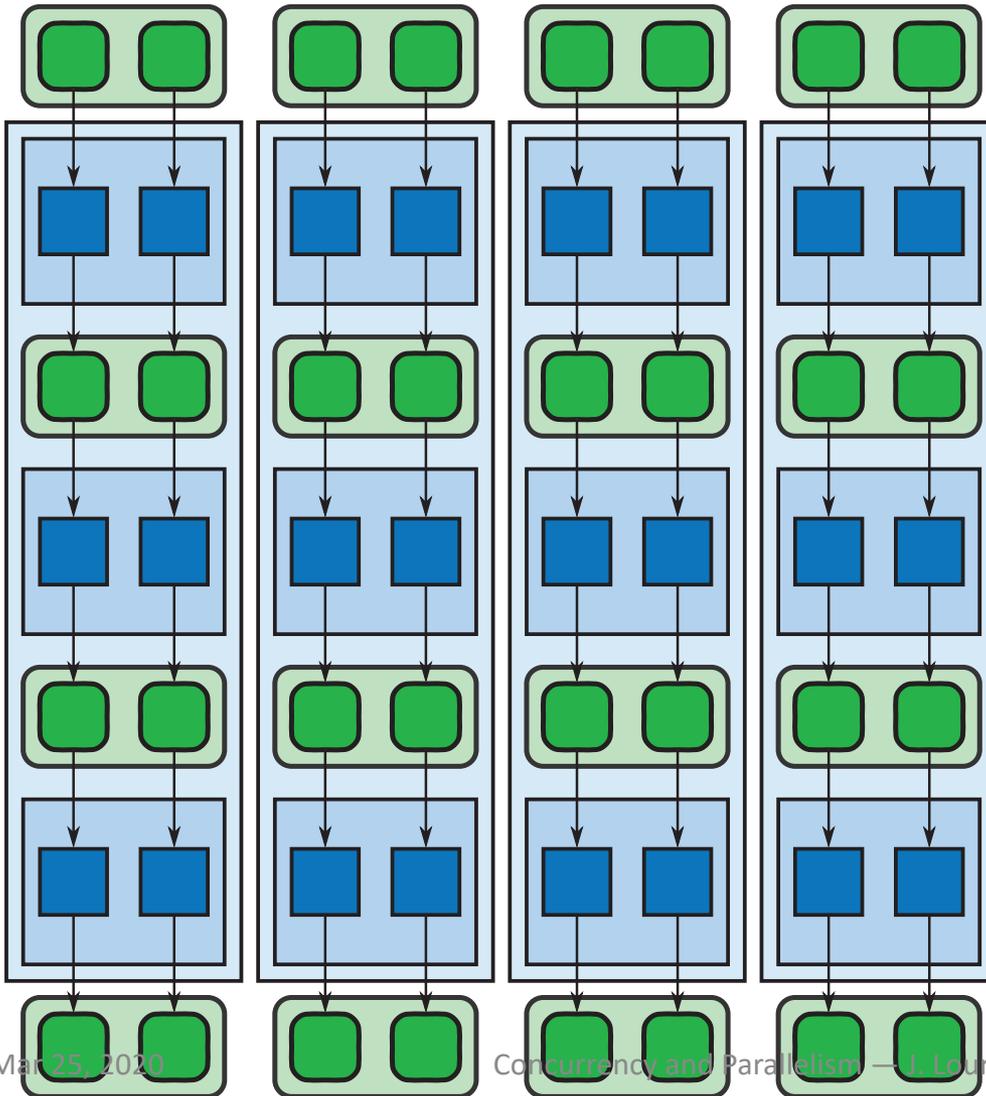
- Often several map operations occur in sequence
 - Vector math consists of many small operations such as additions and multiplications applied as maps
- A naïve implementation may write each intermediate result to memory, wasting memory BW and likely overwhelming the cache

Optimization – Code Fusion



- Can sometimes “fuse” together the operations to perform them at once
- Adds arithmetic intensity, reduces memory/cache usage
- Ideally, operations can be performed using registers alone

Optimization – Cache Fusion



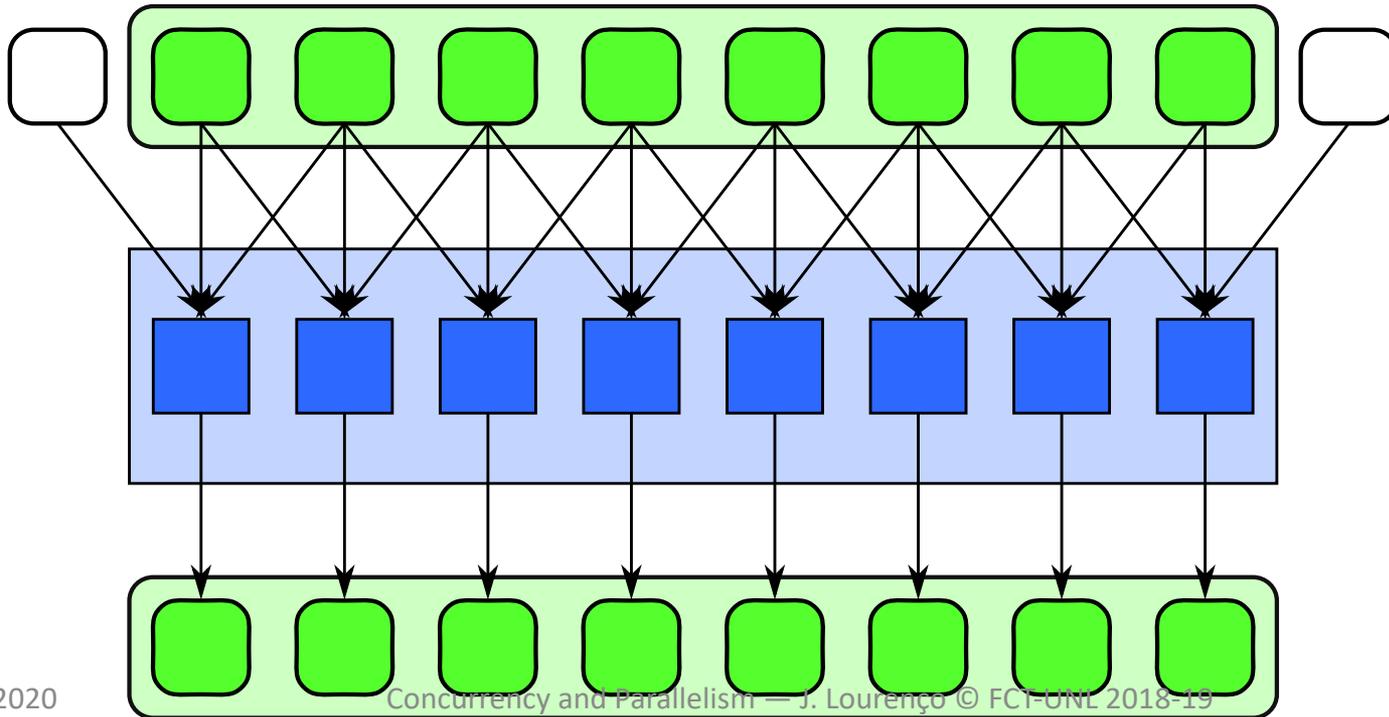
- Sometimes impractical to fuse together the map operations
- Can instead break the work into blocks, giving each CPU one block at a time
- Hopefully, operations use cache alone

Related Patterns

- Three patterns related to map are now discussed here:
 - Stencil
 - Workpile
 - Divide-and-Conquer

Stencil

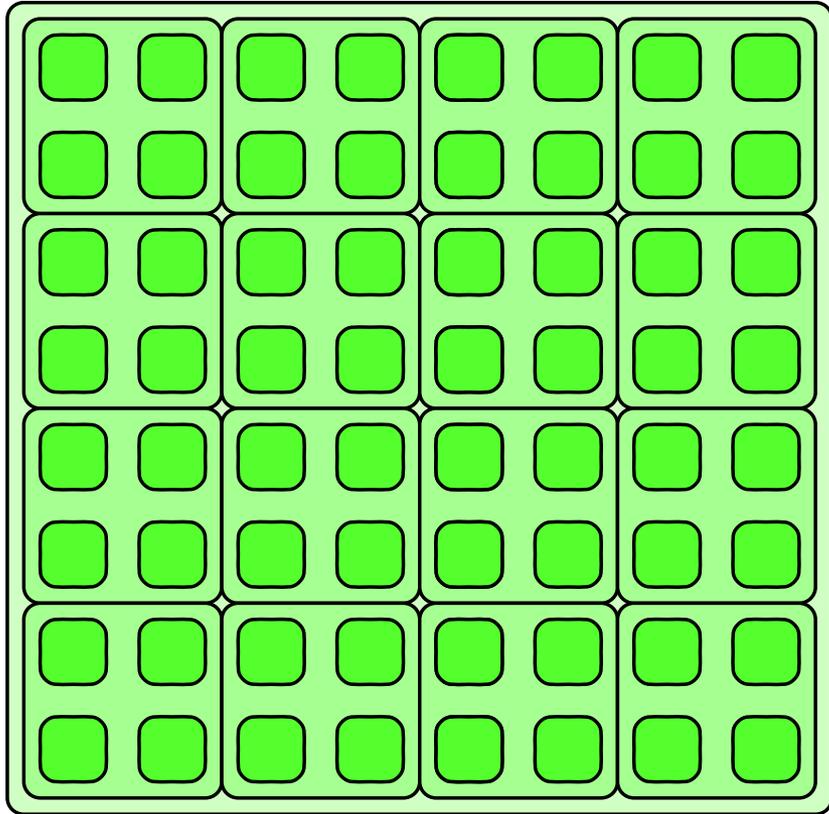
- Each instance of the map function accesses neighbors of its input, offset from its usual input
- Common in imaging and PDE solvers



Workpile

- Work items can be added to the map while it is in progress, from inside map function instances
- Work grows and is consumed by the map
- Workpile pattern terminates when no more work is available

Divide-and-Conquer



- Applies if a problem can be divided into smaller sub-problems recursively until a base case is reached that can be solved serially

Example: Scaled Vector Addition (SAXPY)

- $y \leftarrow ax + y$
 - Scales vector x by a and adds it to vector y
 - Result is stored in input vector y
- Comes from the BLAS (Basic Linear Algebra Subprograms) library
- **Every element in vector x and vector y are independent**

What does $y \leftarrow ax + y$ look like?

	0	1	2	3	4	5	6	7	8	9	10	11
a	4	4	4	4	4	4	4	4	4	4	4	4
*												
x	2	4	2	1	8	3	9	5	5	1	2	1
+												
y	3	7	0	1	4	0	0	4	5	3	1	0
	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
y	11	23	8	5	36	12	36	49	50	7	9	4

Visual: $y \leftarrow ax + y$

	0	1	2	3	4	5	6	7	8	9	10	11
a	4	4	4	4	4	4	4	4	4	4	4	4
*	2	4	2	1	8	3	9	5	5	1	2	1
+												
y	3	7	0	1	4	0	0	4	5	3	1	0
	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
y	11	23	8	5	36	12	36	49	50	7	9	4

Twelve processors used \rightarrow one for each element in the vector

Visual: $y \leftarrow ax + y$

	0	1	2	3	4	5	6	7	8	9	10	11
a	4	4	4	4	4	4	4	4	4	4	4	4
*	2	4	2	1	8	3	9	5	5	1	2	1
+												
y	3	7	0	1	4	0	0	4	5	3	1	0
	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
y	11	23	8	5	36	12	36	49	50	7	9	4

Six processors used → one for every two elements in the vector

Visual: $y \leftarrow ax + y$

	0	1	2	3	4	5	6	7	8	9	10	11
a	4	4	4	4	4	4	4	4	4	4	4	4
*	2	4	2	1	8	3	9	5	5	1	2	1
+												
y	3	7	0	1	4	0	0	4	5	3	1	0
	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
y	11	23	8	5	36	12	36	49	50	7	9	4

Two processors used → one for every six elements in the vector

Serial SAXPY Implementation

```
1 void saxpy_serial(  
2     size_t n,           // the number of elements in the vectors  
3     float a,           // scale factor  
4     const float x[],   // the first input vector  
5     float y[]          // the output vector and second input vector  
6 ) {  
7     for (size_t i = 0; i < n; ++i)  
8         y[i] = a * x[i] + y[i];  
9 }
```

Cilk Plus SAXPY Implementation

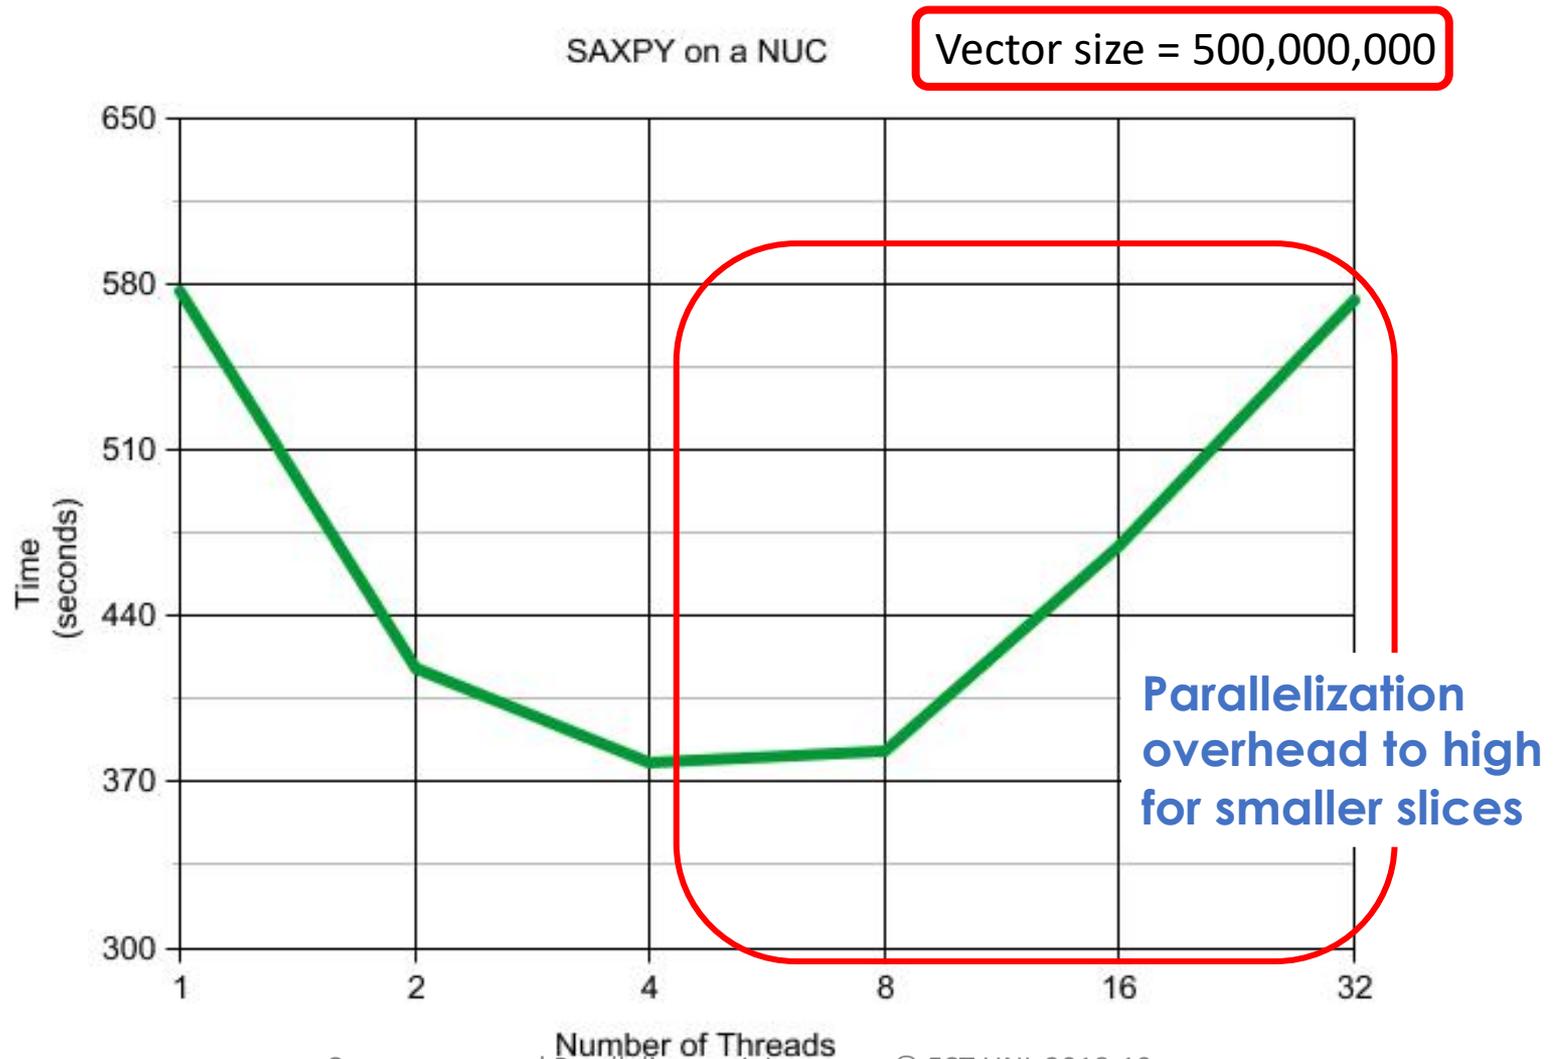
```
1 void saxpy_cilk(  
2     int n,           // the number of elements in the vectors  
3     float a,        // scale factor  
4     float x[],      // the first input vector  
5     float y[]       // the output vector and second input vector  
6 ) {  
7     cilk_for (int i = 0; i < n; ++i)  
8         y[i] = a * x[i] + y[i];  
9 }
```

$y[0:n] = a * x[0:n] + y[0:n]$

OpenMP SAXPY Implementation

```
1 void saxpy_openmp(  
2     int n,          // the number of elements in the vectors  
3     float a,       // scale factor  
4     float x[],     // the first input vector  
5     float y[]      // the output vector and second input vector  
6 ) {  
7     #pragma omp parallel for  
8         for (int i = 0; i < n; ++i)  
9             y[i] = a * x[i] + y[i];  
10 }
```

OpenMP SAXPY Performance

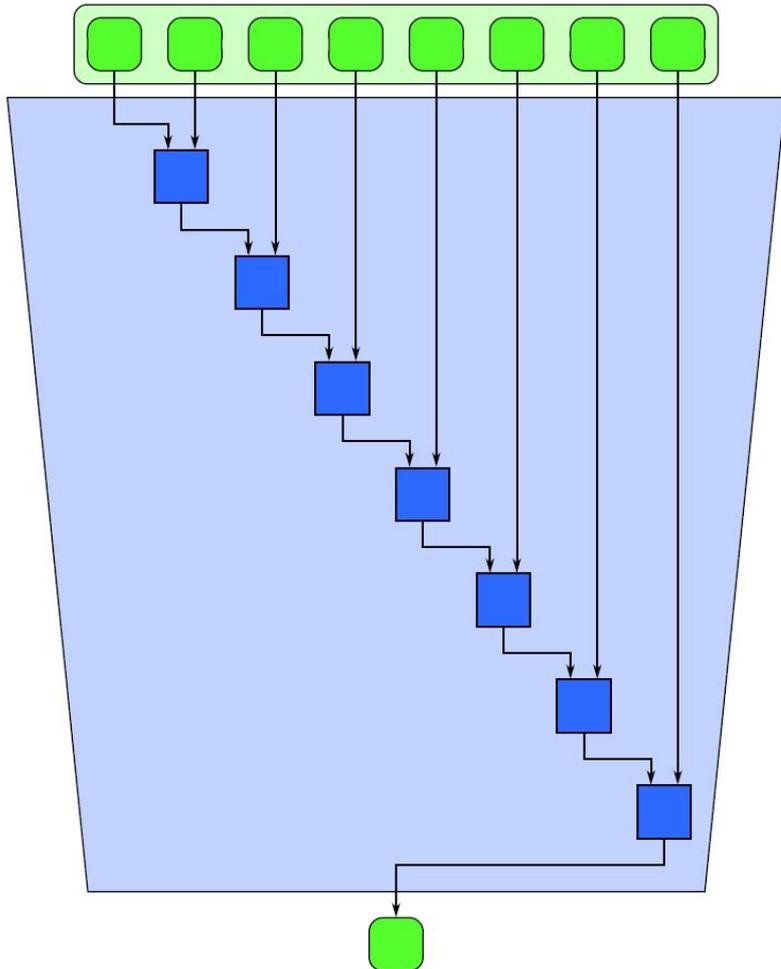


Reduce

- **Reduce** is used to combine a collection of elements into one summary value
- A combiner function combines elements pairwise
- A combiner function only needs to be *associative* to be parallelizable
- Example combiner functions:
 - Addition
 - Multiplication
 - Maximum / Minimum

Reduce

Serial Reduction

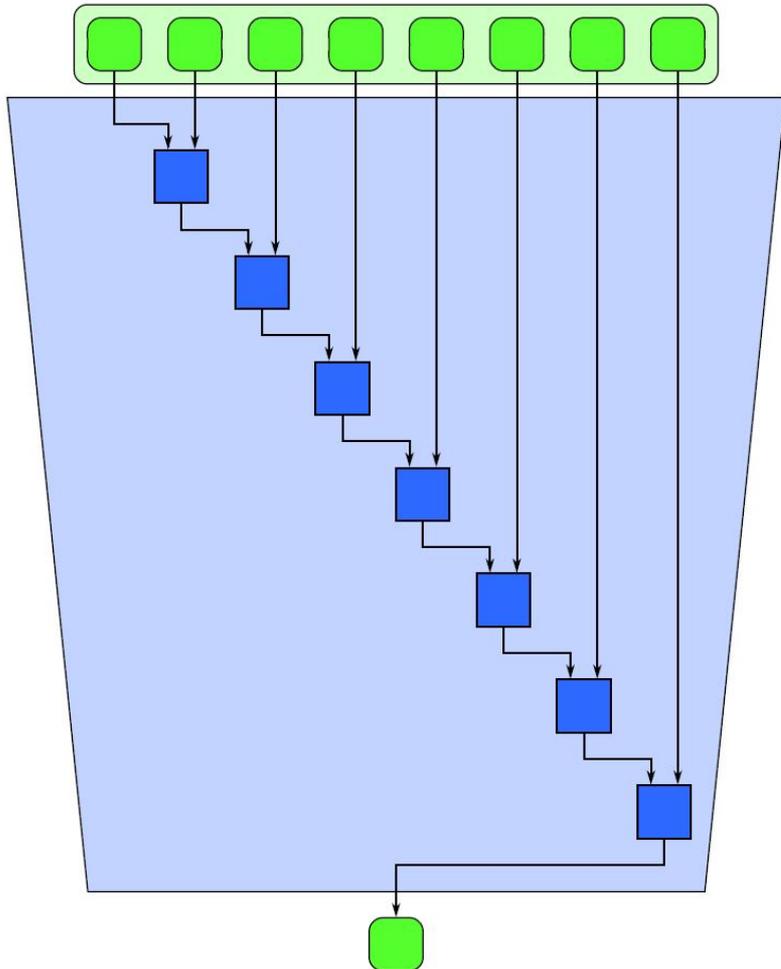


```
1  template<typename T>
2  T reduce(
3      T (*f)(T,T), //combiner function
4      size_t n, // number of elements in input array
5      T a[] // input array
6  ) {
7      assert(n > 0);
8      T accum = a[0];
9      for (size_t i = 1; i < n; i++) {
10         accum = f(accum, a[i]);
11     }
12     return accum;
13 }
```

The input array cannot be empty!

Reduce

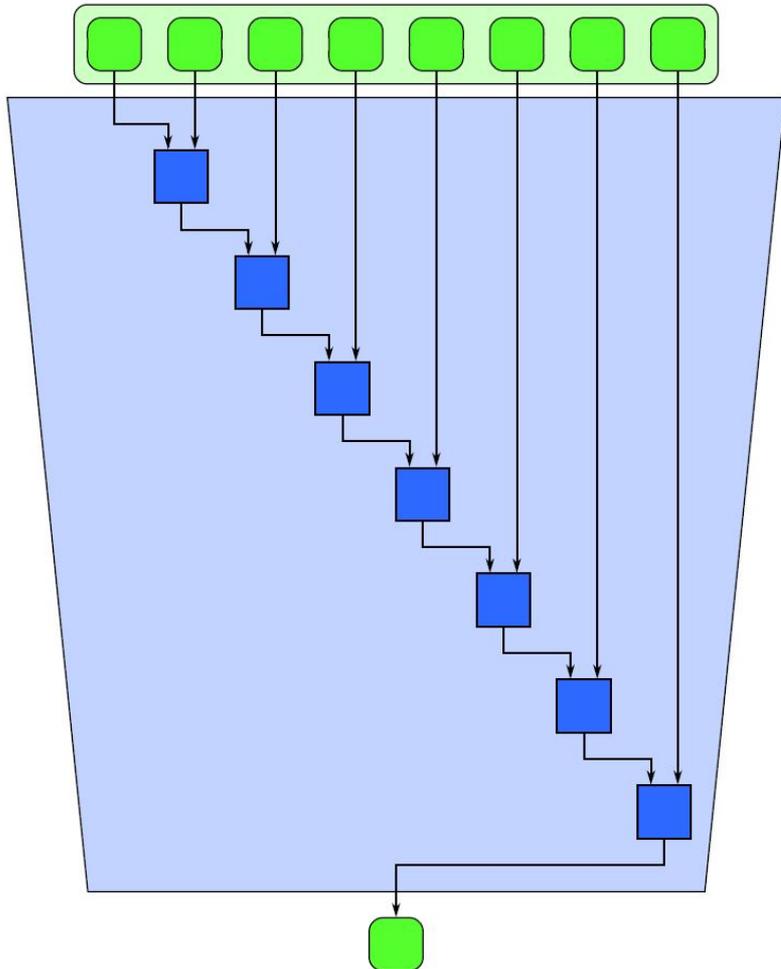
Serial Reduction



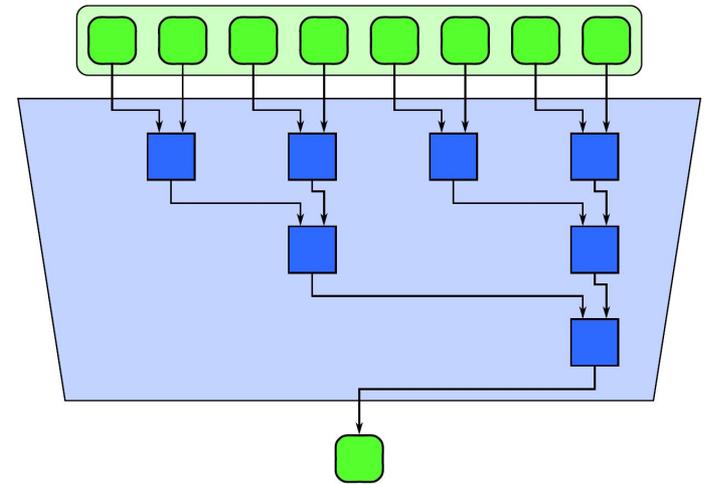
```
1  template<typename T>
2  T reduce(
3      T (*f)(T,T), // combiner function
4      size_t n, // number of elements in input array
5      T a[], // input array
6      T identity // identity of combiner function
7  ) {
8      T accum = identity;
9      for (size_t i = 0; i < n; ++i) {
10         accum = f(accum, a[i]);
11     }
12     return accum;
13 }
```

Reduce

Serial Reduction



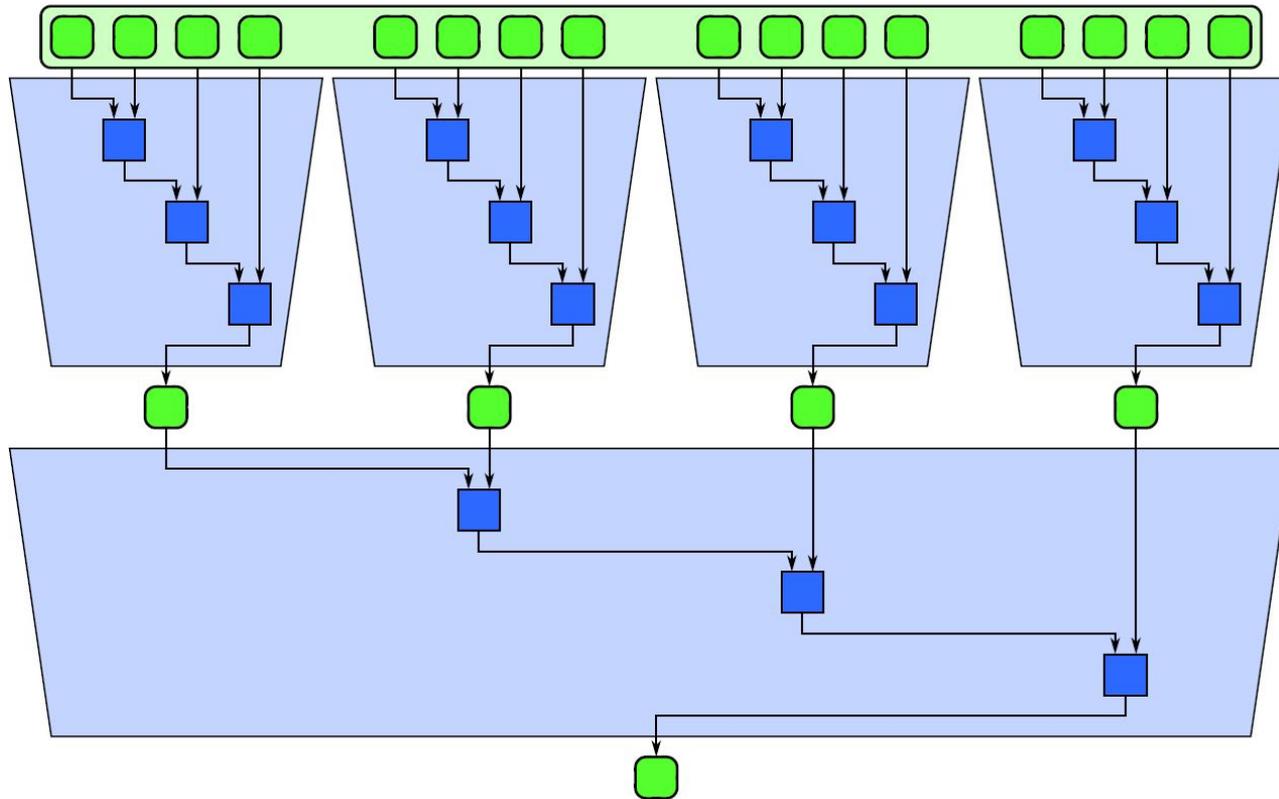
Parallel Reduction



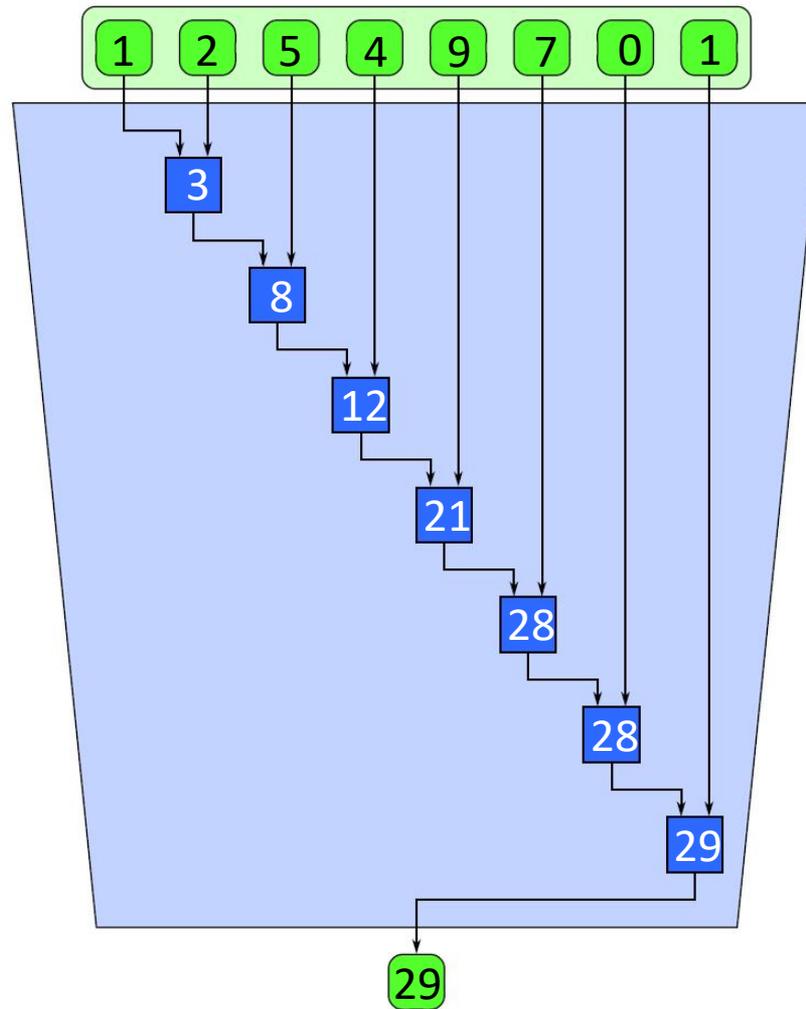
Implementation later...

Reduce

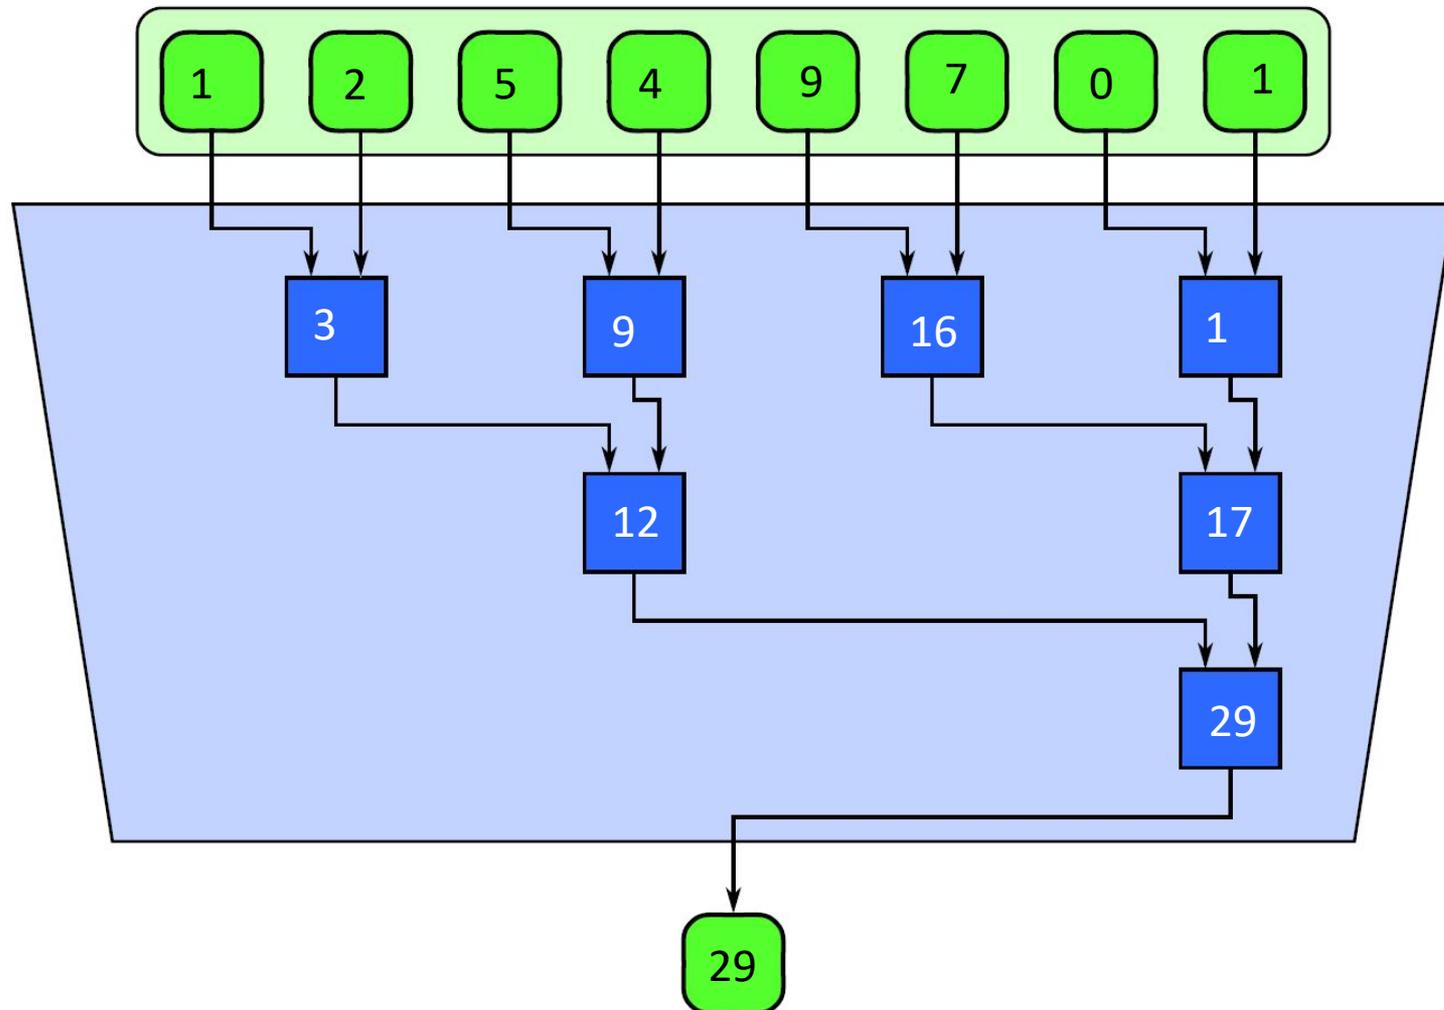
- **Tiling** is used to break chunks of work up for workers to reduce serially



Reduce – Add Example

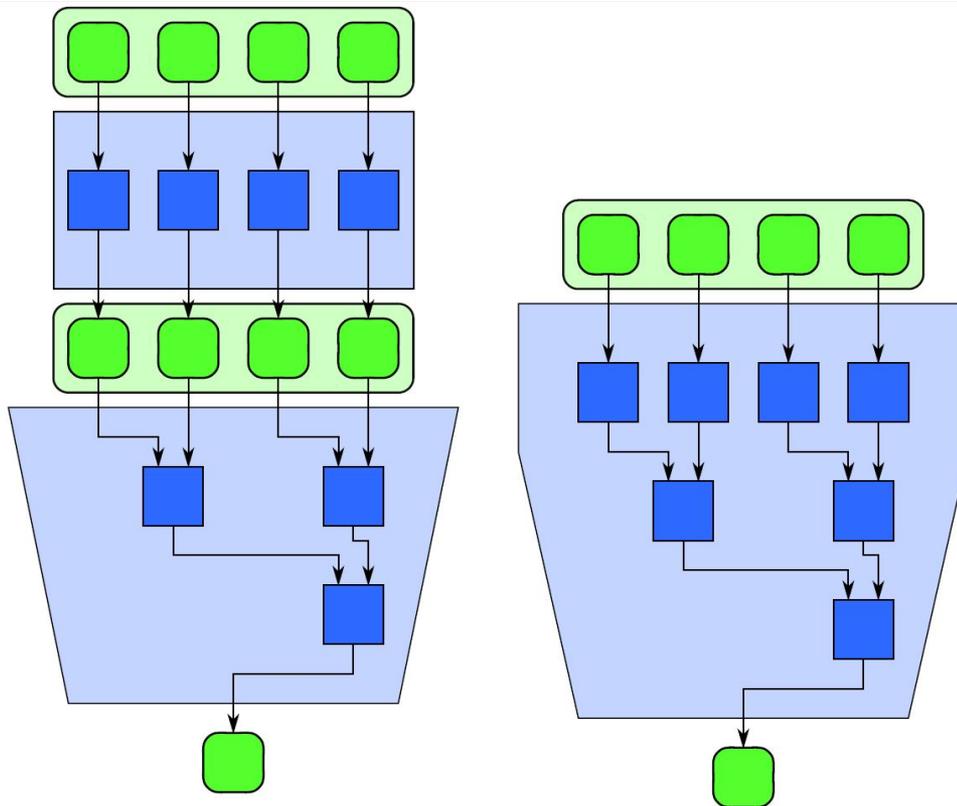


Reduce – Add Example



Reduce

- We can “fuse” the map and reduce patterns



Reduce

- Precision can become a problem with reductions on floating point data
- Different orderings of floating-point data can change the reduction value

Reduce Example: Dot Product

- 2 vectors of same length
- Map (x) to multiply the components
- Then reduce with (+) to get the final answer

$$a \cdot b = \sum_{i=0}^{n-1} a_i b_i$$

Also: $\vec{a} \cdot \vec{b} = |\vec{a}| \cos(\theta) |\vec{b}|$

Dot Product – Example Uses

- Essential operation in physics, graphics, video games,...
- Gaming analogy: in Mario Kart, there are “boost pads” on the ground that increase your speed
 - red vector is your speed (x and y direction)
 - blue vector is the orientation of the boost pad (x and y direction). Larger numbers are more power.



$$Total = speed_x \cdot boost_x + speed_y \cdot boost_y$$

Ref: <http://betterexplained.com/articles/vector-calculus-understanding-the-dot-product/>

Dot Product – Example Uses

- How much boost will you get? For the analogy, imagine the pad multiplies your speed:
- If you come in going 0, you'll get nothing
- If you cross the pad perpendicularly, you'll get 0 [just like the banana obliteration, it will give you 0x boost in the perpendicular direction]



$$Total = speed_x \cdot boost_x + speed_y \cdot boost_y$$

Ref: <http://betterexplained.com/articles/vector-calculus-understanding-the-dot-product/>

Dot Product – Serial implem.

```
1 float sprod(  
2     size_t n,  
3     const float a[],  
4     const float b[]  
5 ) {  
6     float res = 0.0f;  
7     for (size_t i = 0; i < n; i++) {  
8         res += a[i] * b[i];  
9     }  
10    return res;  
11 }
```

$$a \cdot b = \sum_{i=0}^{n-1} a_i b_i$$

Dot Product – Cilk+ with Array Notation

$$a \cdot b = \sum_{i=0}^{n-1} a_i b_i$$

```
1 float cilkplus_sprod(  
2     size_t n,  
3     const float a[],  
4     const float b[]  
5 ) {  
6     return __sec_reduce_add(a[0:n] * b[0:n]);  
7 }
```

Dot Product – OpenMP

```
1 float openmp_sprod(  
2     size_t n,  
3     const float *a,  
4     const float *b  
5 ) {  
6     float res = 0.0f;  
7     #pragma omp parallel for reduction(+:res)  
8     for (size_t i = 0; i < n; i++) {  
9         res += a[i] * b[i];  
10    }  
11    return res;  
12 }
```

$$a \cdot b = \sum_{i=0}^{n-1} a_i b_i$$

The END
