

Parallel Prefix

CSE 332 – Section 8

Slides by James Richie Sulaeman



Parallel Prefix

Parallel Prefix

Parallel prefix is a type of programming problem where:

- A given array of elements needs to be **processed in parallel**
- Each element is **combined with its predecessors** through some operation
 - i.e. in parallel prefix sum, each element is summed with its predecessors

Allows for efficient computation of certain operations on large datasets since it enables multiple threads to work simultaneously on different portions of the data.

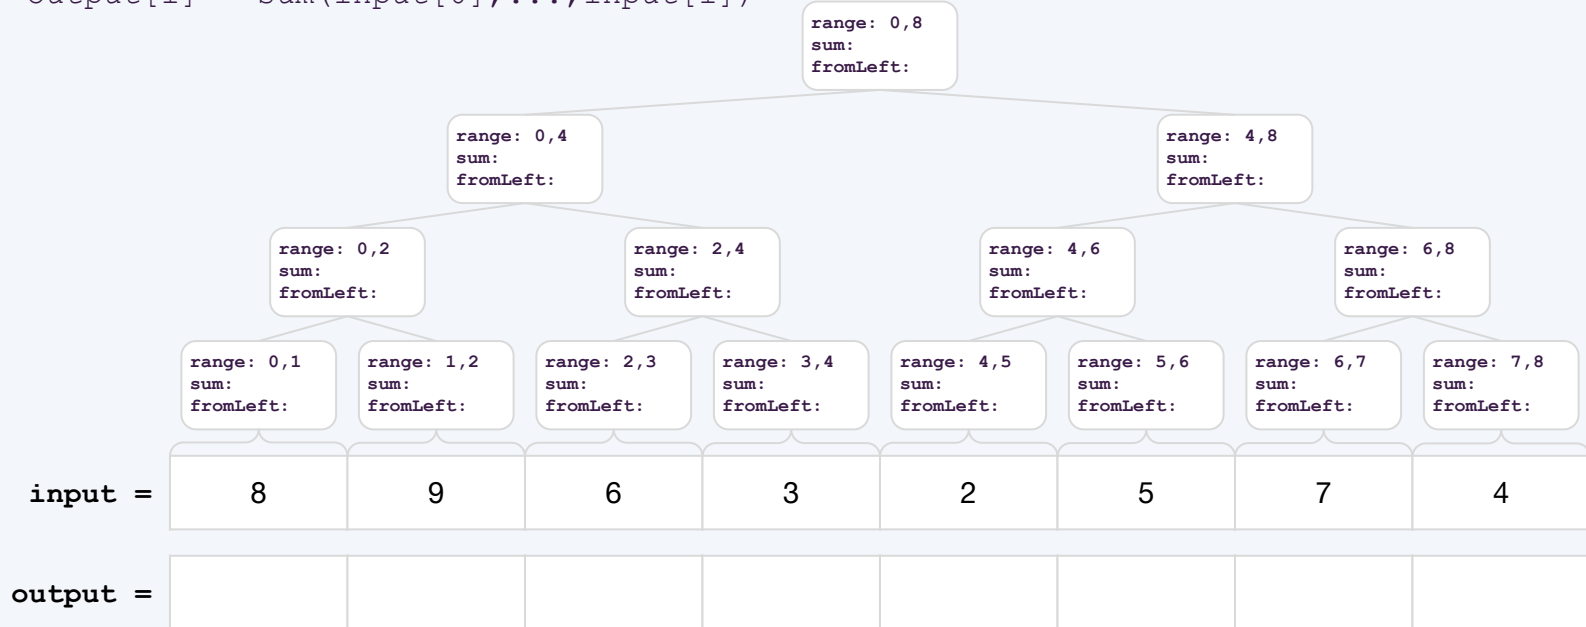
Problem 0

Parallel Prefix Sum

Problem 0

1. Divide problem into parallel pieces
cutoff = 1

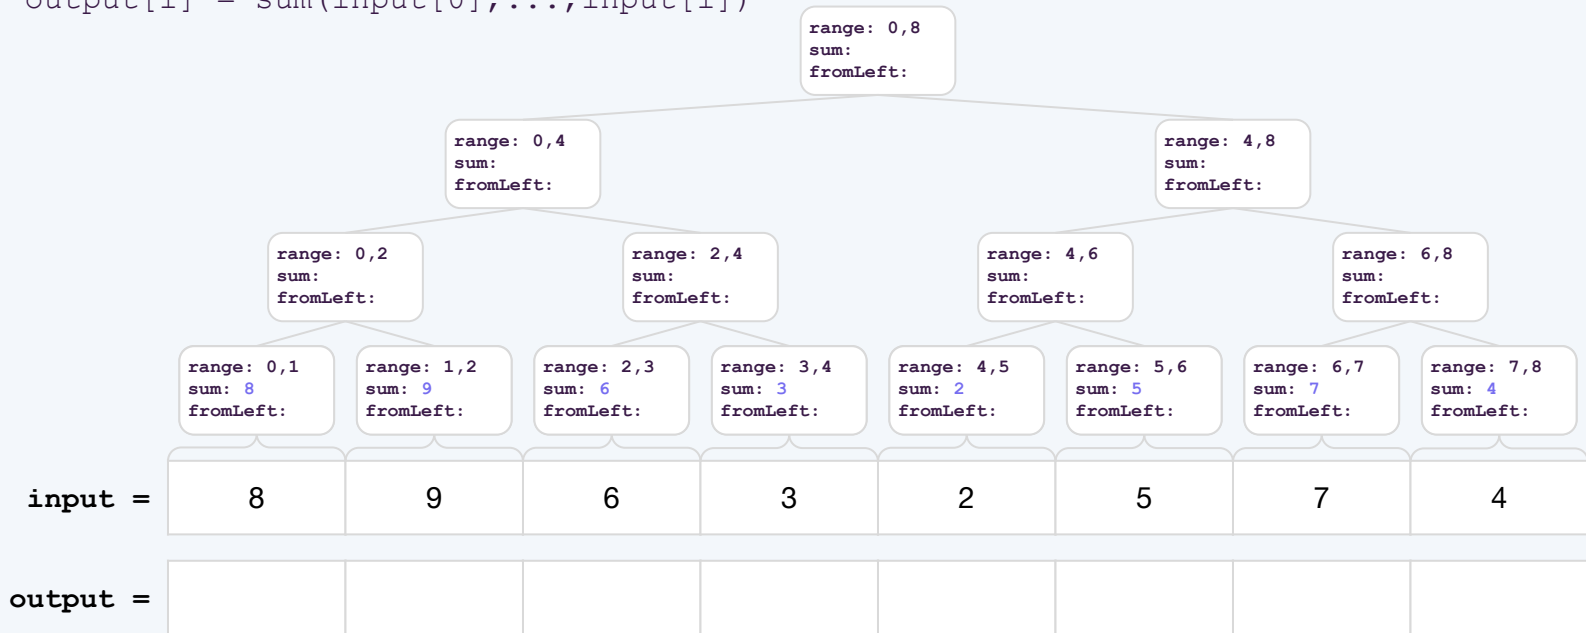
Given an array `input`, output an array such that
`output[i] = sum(input[0], ..., input[i])`



Problem 0

1. Divide problem into parallel pieces
cutoff = 1
2. Find sum at cutoff
leaves[i].sum = input[i]

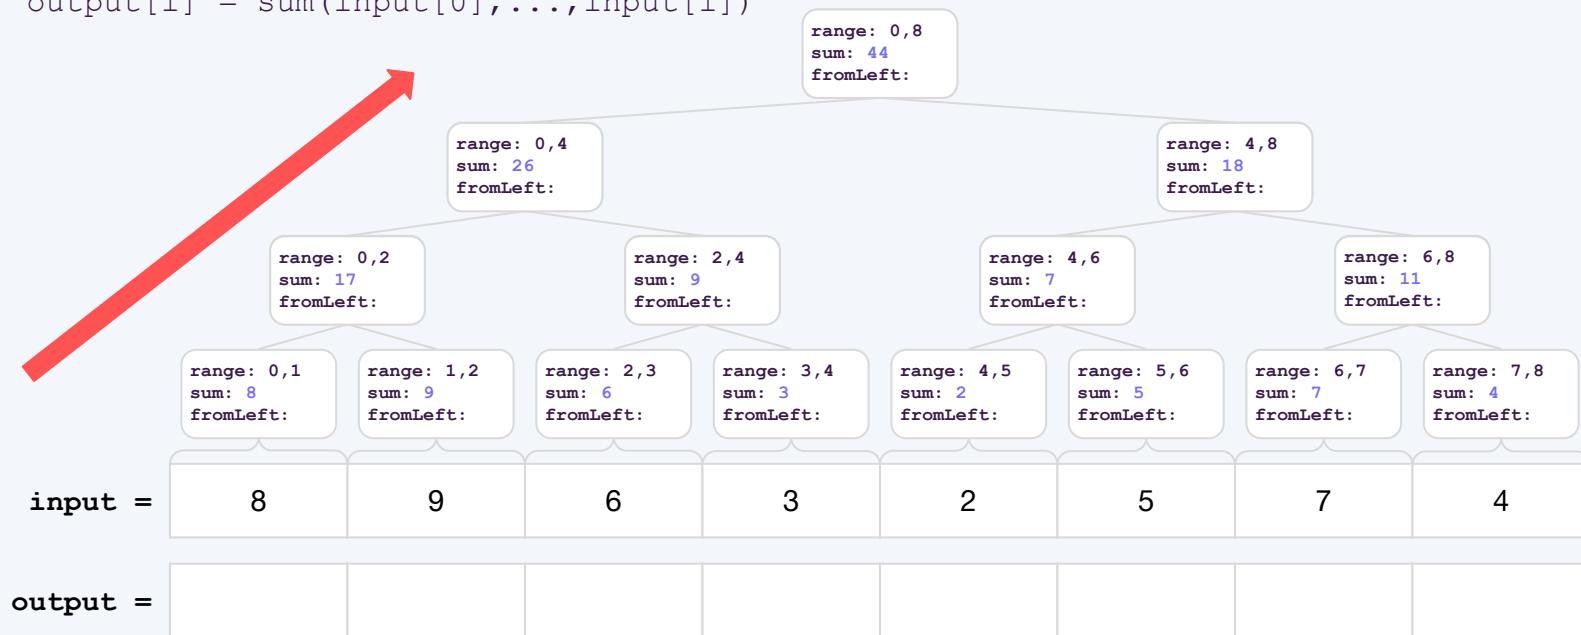
Given an array `input`, output an array such that
`output[i] = sum(input[0], ..., input[i])`



Problem 0

1. Divide problem into parallel pieces
cutoff = 1
2. Find sum at cutoff
leaves[i].sum = input[i]
3. Propagate sum up
parent.sum = left.sum + right.sum

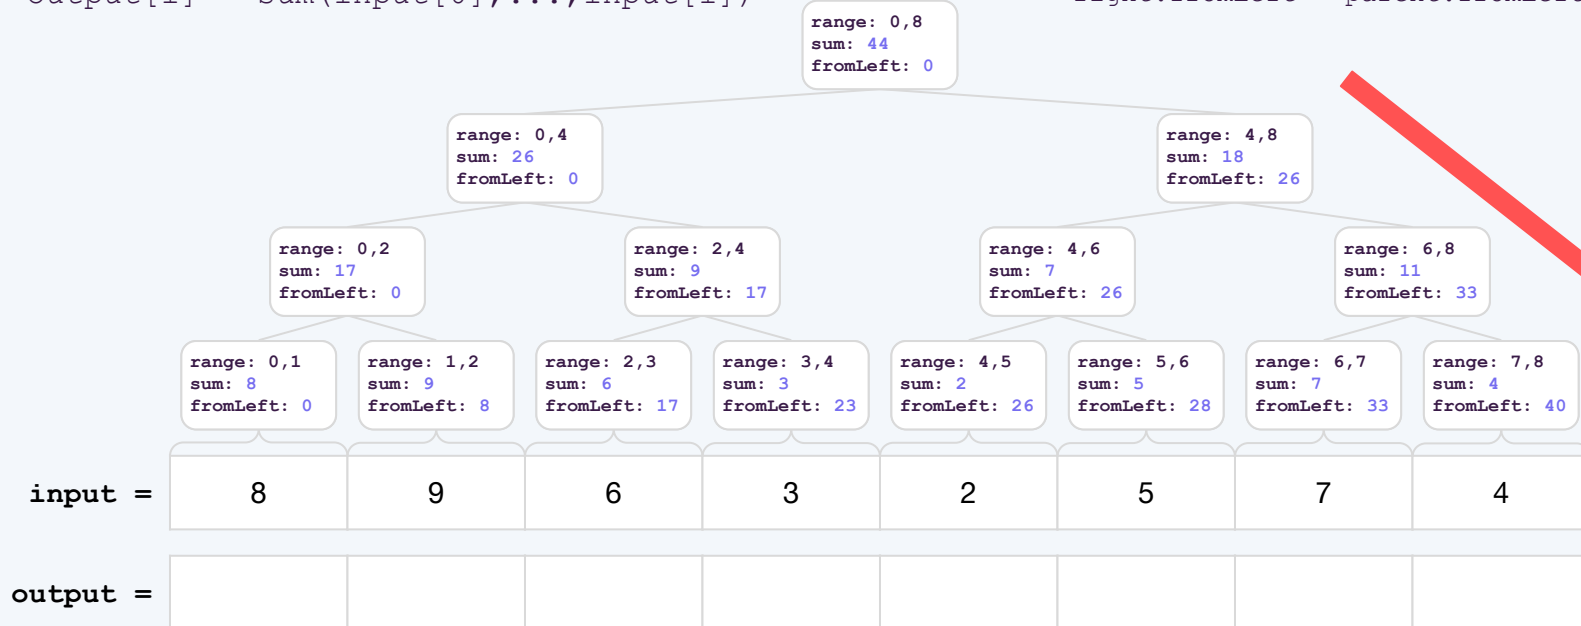
Given an array `input`, output an array such that
`output[i] = sum(input[0], ..., input[i])`



Problem 0

Given an array `input`, output an array such that
`output[i] = sum(input[0], ..., input[i])`

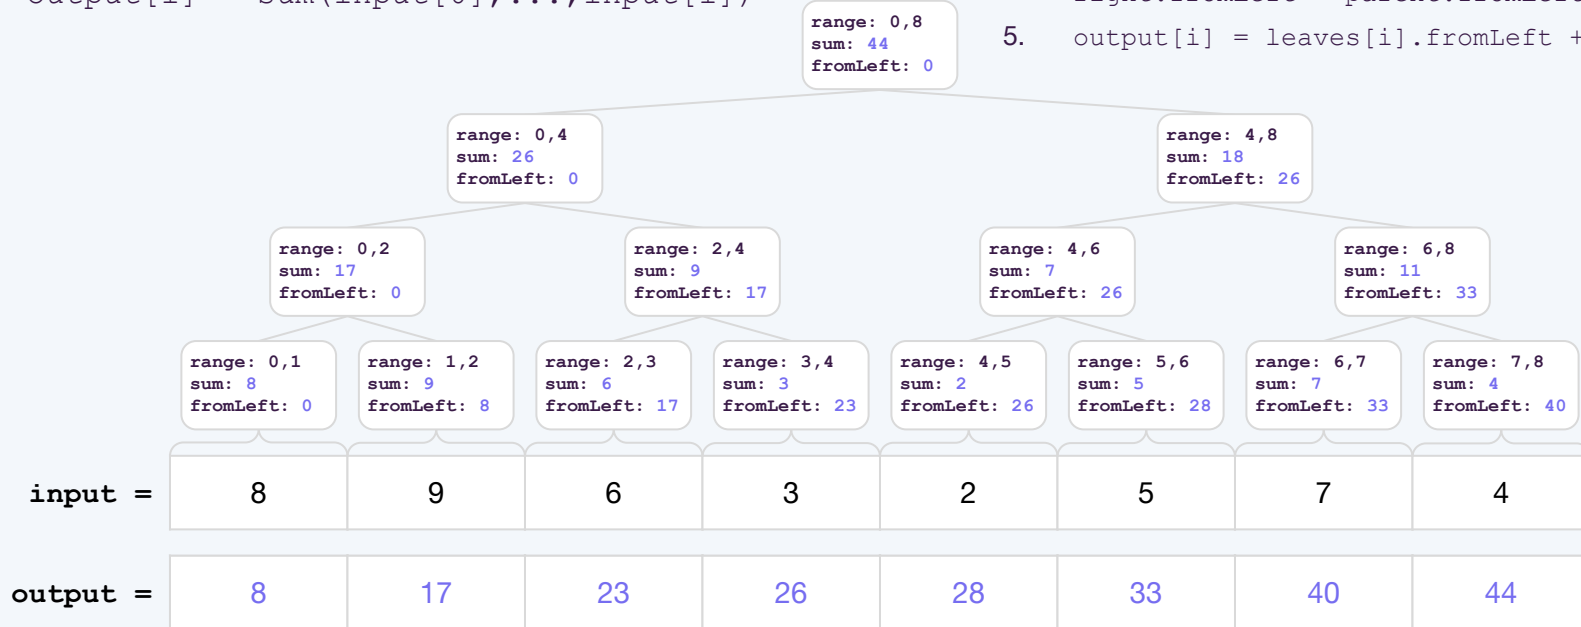
1. Divide problem into parallel pieces
`cutoff = 1`
2. Find sum at cutoff
`leaves[i].sum = input[i]`
3. Propagate sum up
`parent.sum = left.sum + right.sum`
4. Propagate `fromLeft` down
`left.fromLeft = parent.fromLeft`
`right.fromLeft = parent.fromLeft + left.sum`



Problem 0

Given an array `input`, output an array such that
`output[i] = sum(input[0], ..., input[i])`

1. Divide problem into parallel pieces
`cutoff = 1`
2. Find sum at cutoff
`leaves[i].sum = input[i]`
3. Propagate sum up
`parent.sum = left.sum + right.sum`
4. Propagate `fromLeft` down
`left.fromLeft = parent.fromLeft`
`right.fromLeft = parent.fromLeft + left.sum`
5. `output[i] = leaves[i].fromLeft + input[i]`



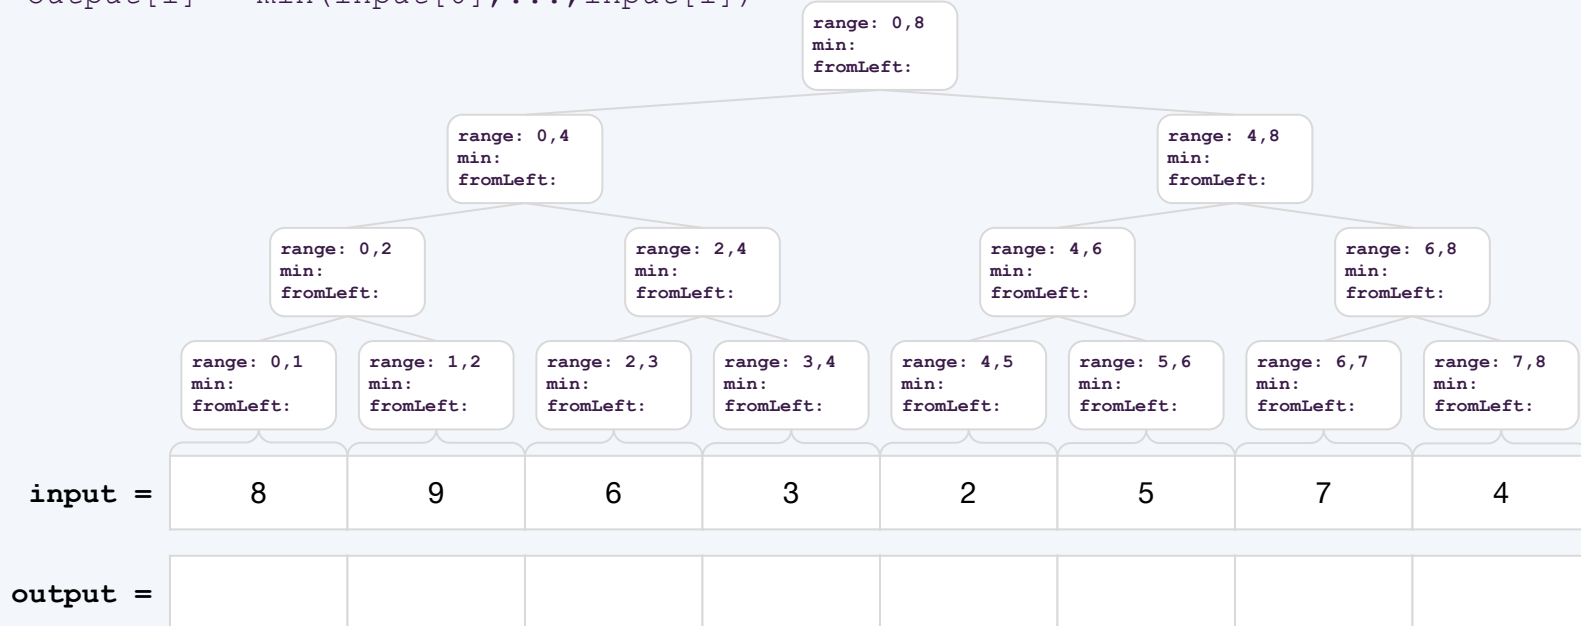
Problem 1

Parallel Prefix FindMin

Problem 1

1. Divide problem into parallel pieces
cutoff = 1

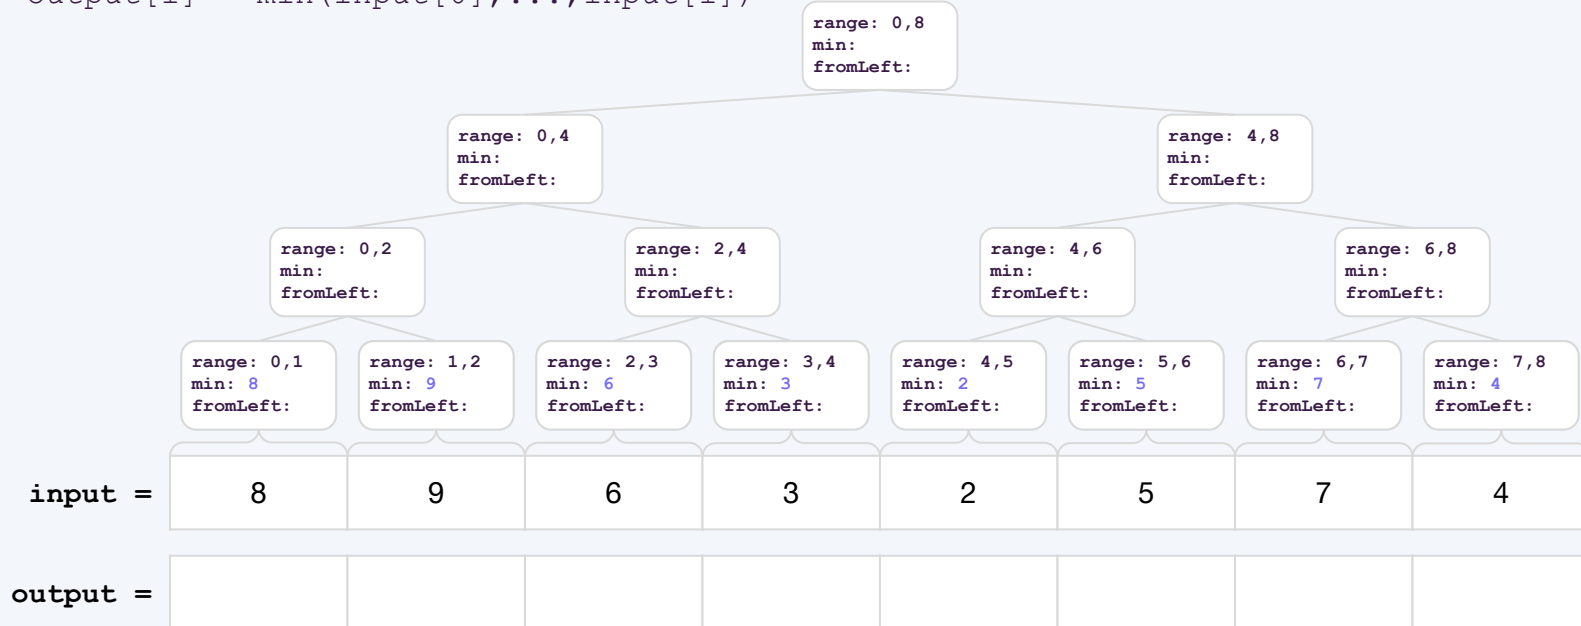
Given an array `input`, output an array such that
`output[i] = min(input[0], ..., input[i])`



Problem 1

1. Divide problem into parallel pieces
cutoff = 1
2. Find min at cutoff
leaves[i].min = input[i]

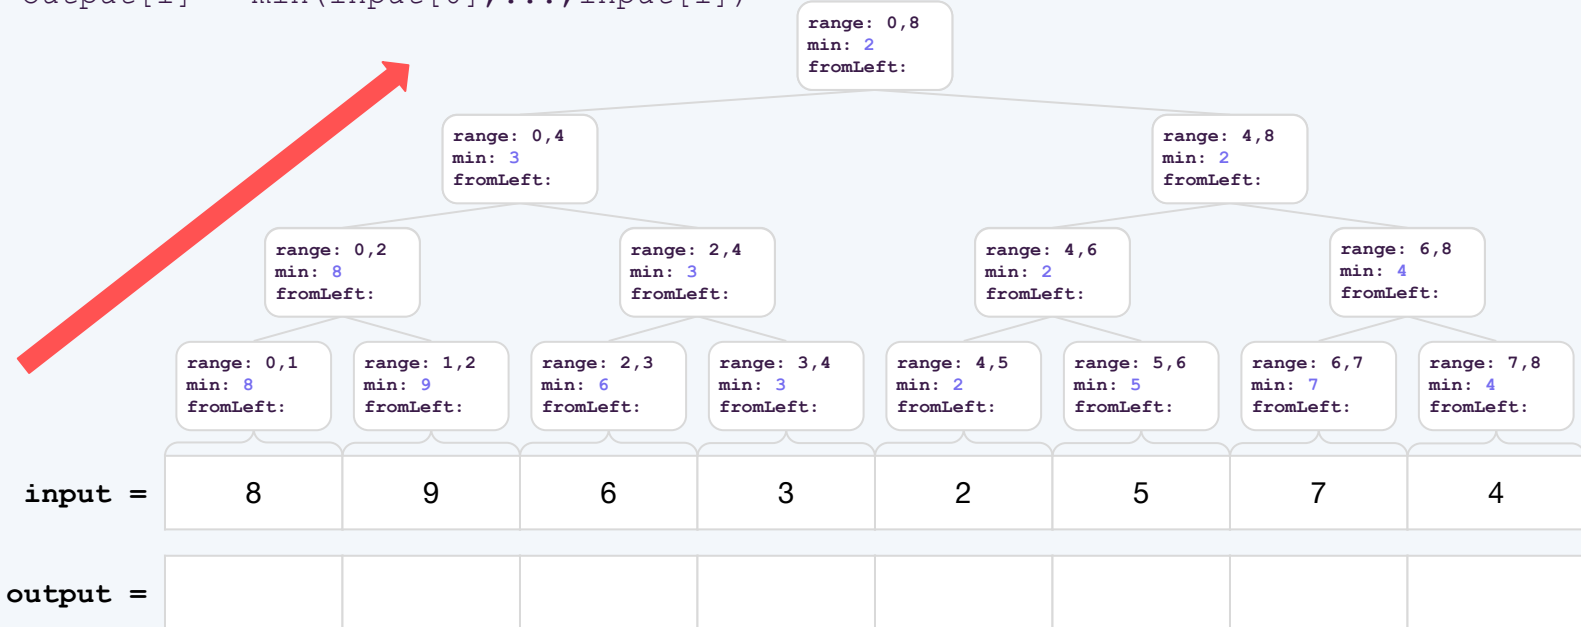
Given an array `input`, output an array such that
`output[i] = min(input[0], ..., input[i])`



Problem 1

1. Divide problem into parallel pieces
cutoff = 1
2. Find min at cutoff
leaves[i].min = input[i]
3. Propagate min up
parent.min = min(left.min, right.min)

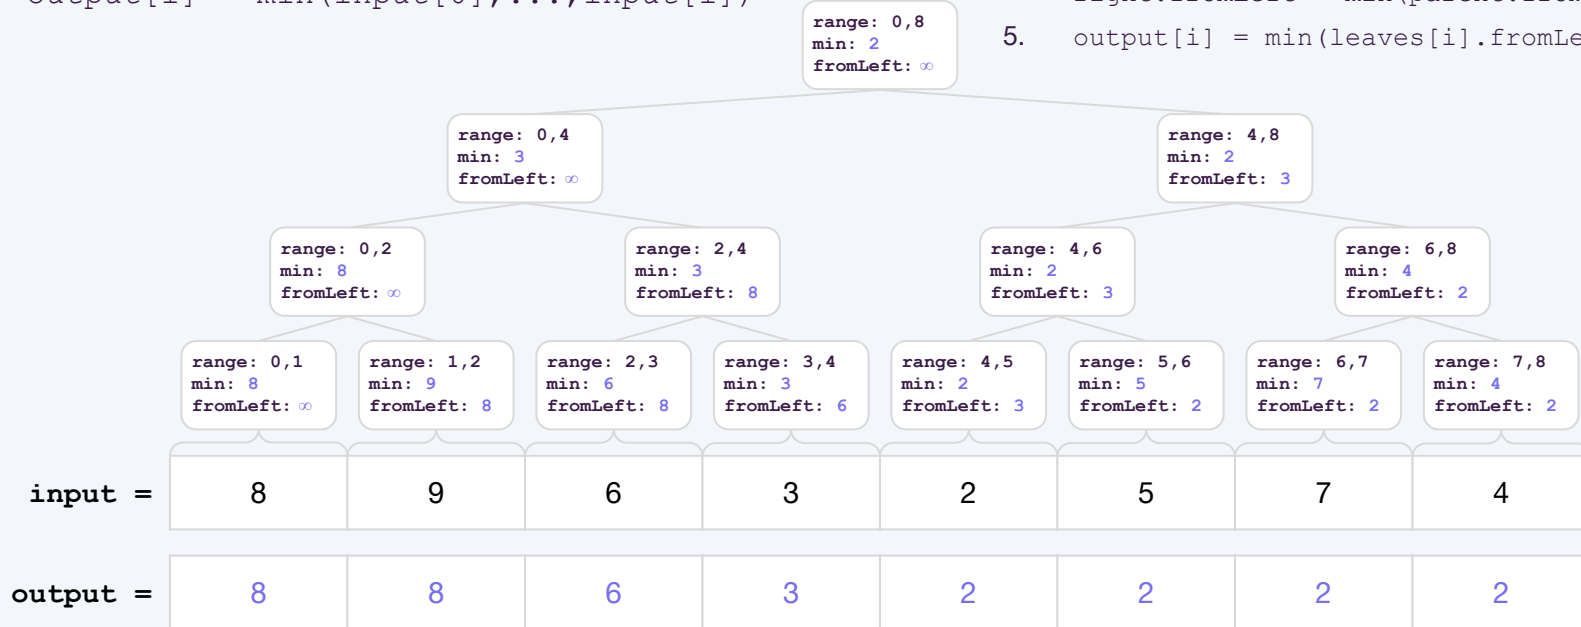
Given an array `input`, output an array such that
`output[i] = min(input[0], ..., input[i])`



Problem 1

Given an array `input`, output an array such that
`output[i] = min(input[0], ..., input[i])`

1. Divide problem into parallel pieces
`cutoff = 1`
2. Find min at cutoff
`leaves[i].min = input[i]`
3. Propagate min up
`parent.min = min(left.min, right.min)`
4. Propagate `fromLeft` down
`left.fromLeft = parent.fromLeft`
`right.fromLeft = min(parent.fromLeft, left.min)`
5. `output[i] = min(leaves[i].fromLeft, input[i])`



Problem 2

Parallel Pack

Problem 2 Overview

Given an array `input`, output an array that contains only the elements that are less than 10

1. **Parallel map** to compute `bits` array where `bits[i] = (input[i] < 10)`

<code>input =</code>	12	5	-8	34	6	10	2	7
<code>bits =</code>	0	1	1	0	1	0	1	1

2. **Parallel prefix sum** on `bits` array to compute `bitsum` array

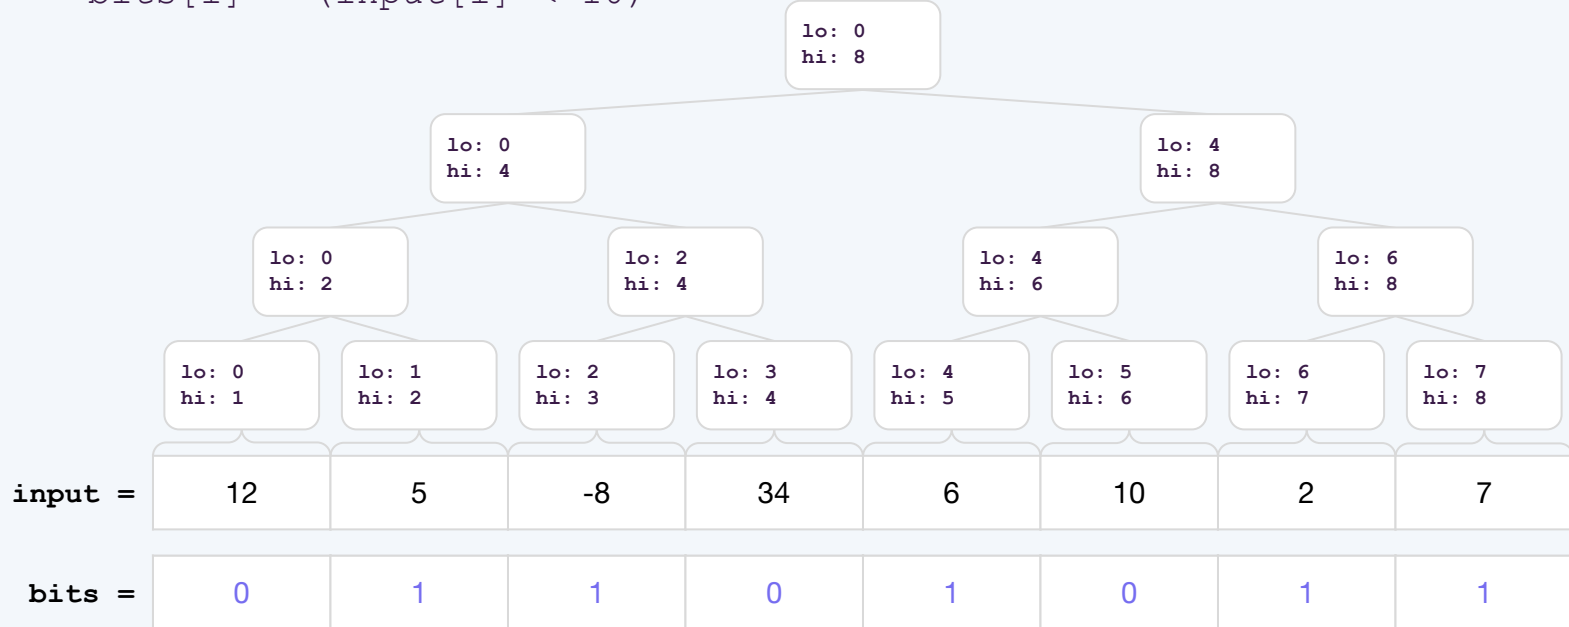
<code>bitsum =</code>	0	1	2	2	3	3	4	5
-----------------------	---	---	---	---	---	---	---	---

3. **Parallel map** to produce `output` array where `output[bitsum[i] - 1] = input[i]`
if `bits[i] == 1`

<code>output =</code>	5	-8	6	2	7
-----------------------	---	----	---	---	---

Problem 2

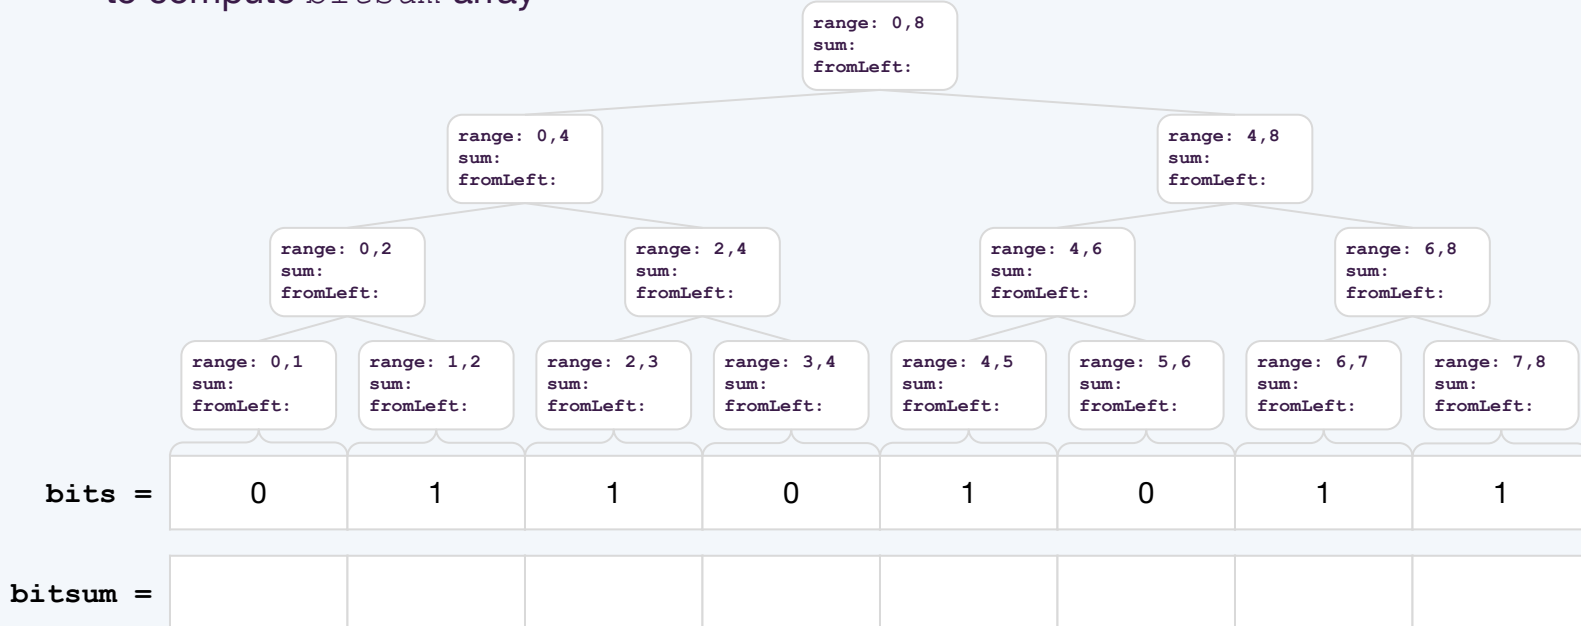
1. **Parallel map** to compute `bits` array where
`bits[i] = (input[i] < 10)`



Problem 2

1. Divide problem into parallel pieces
cutoff = 1

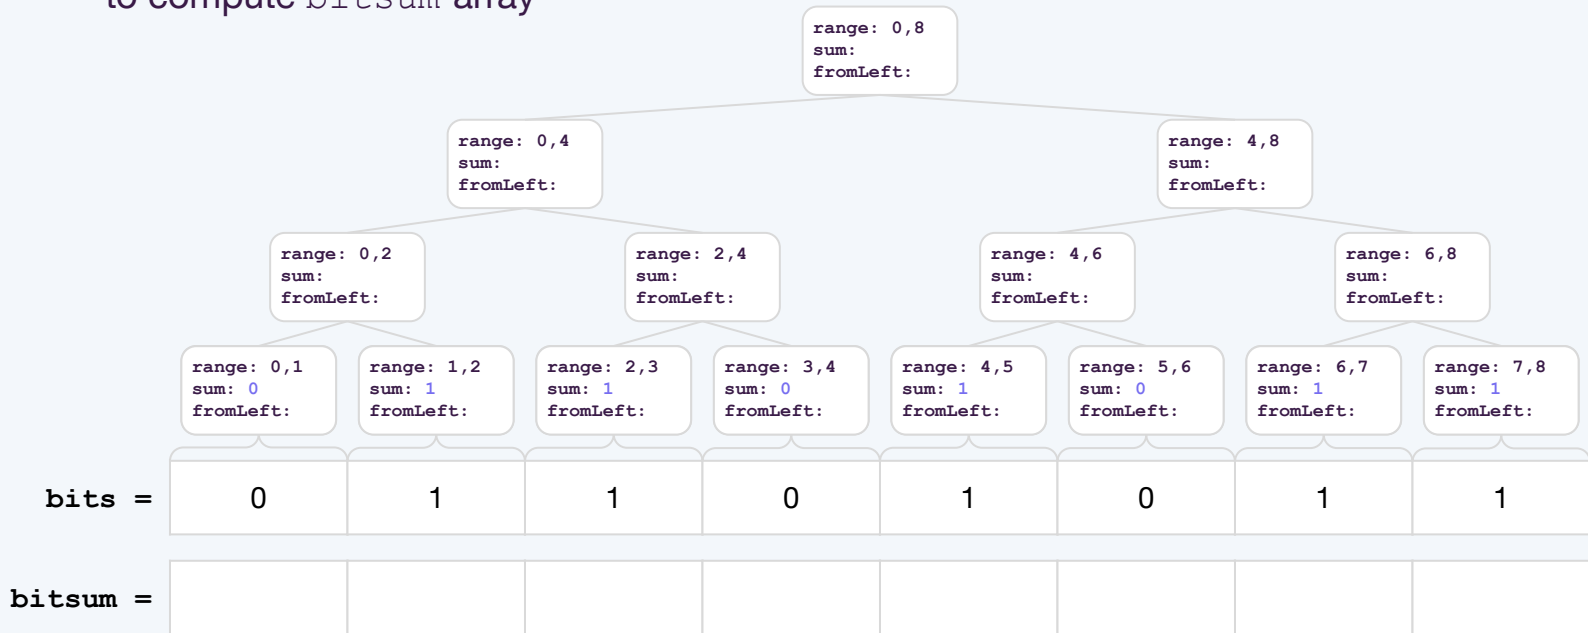
2. **Parallel prefix sum** on `bits` array
to compute `bitsum` array



Problem 2

1. Divide problem into parallel pieces
cutoff = 1
2. Find sum at cutoff
leaves[i].sum = bits[i]

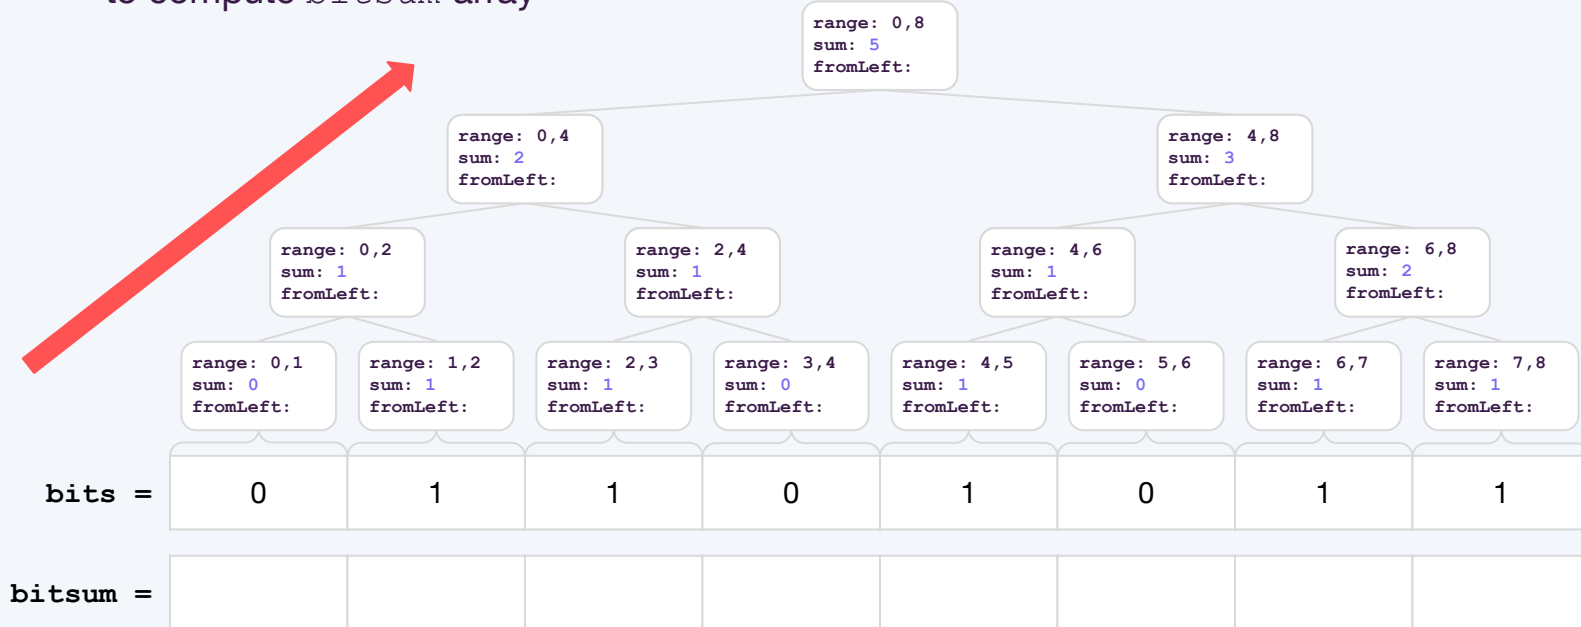
2. **Parallel prefix sum** on `bits` array
to compute `bitsum` array



Problem 2

1. Divide problem into parallel pieces
cutoff = 1
2. Find sum at cutoff
leaves[i].sum = bits[i]
3. Propagate sum up
parent.sum = left.sum + right.sum

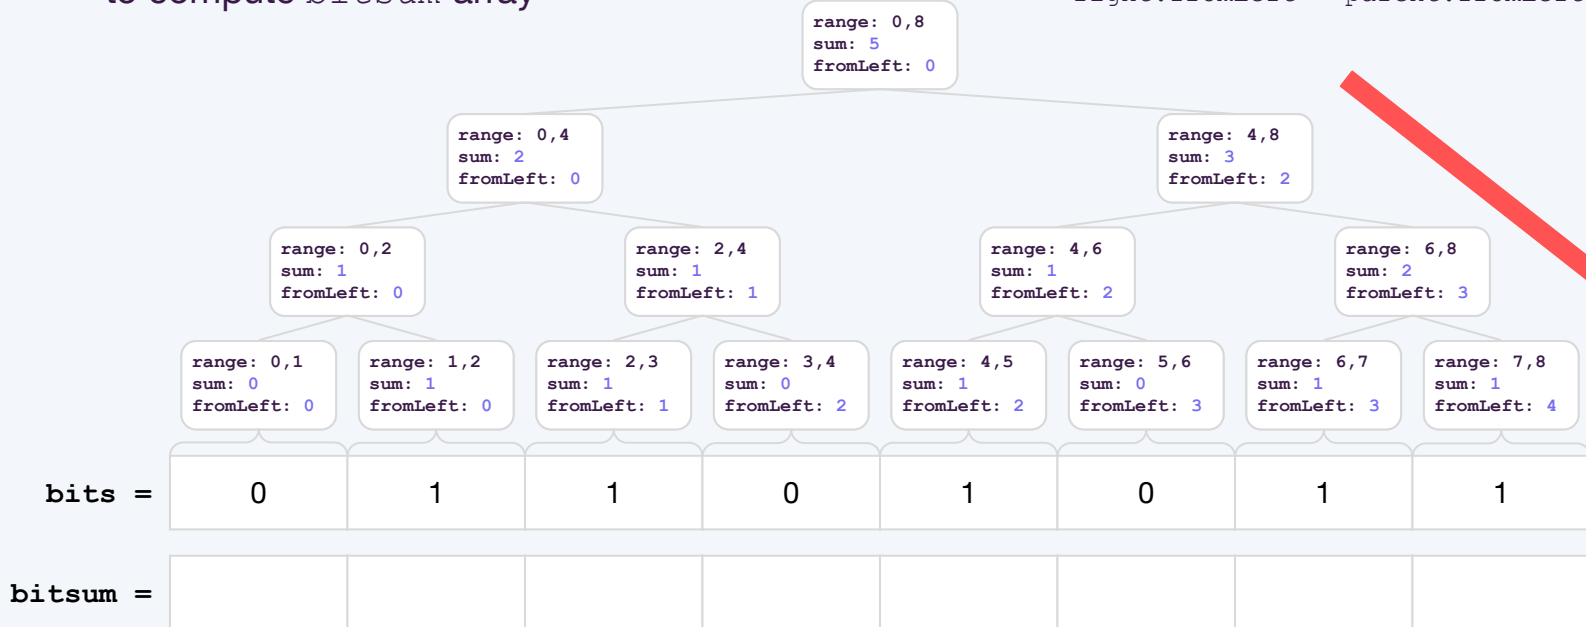
2. **Parallel prefix sum** on bits array
to compute bitsum array



Problem 2

2. Parallel prefix sum on bits array to compute bitsum array

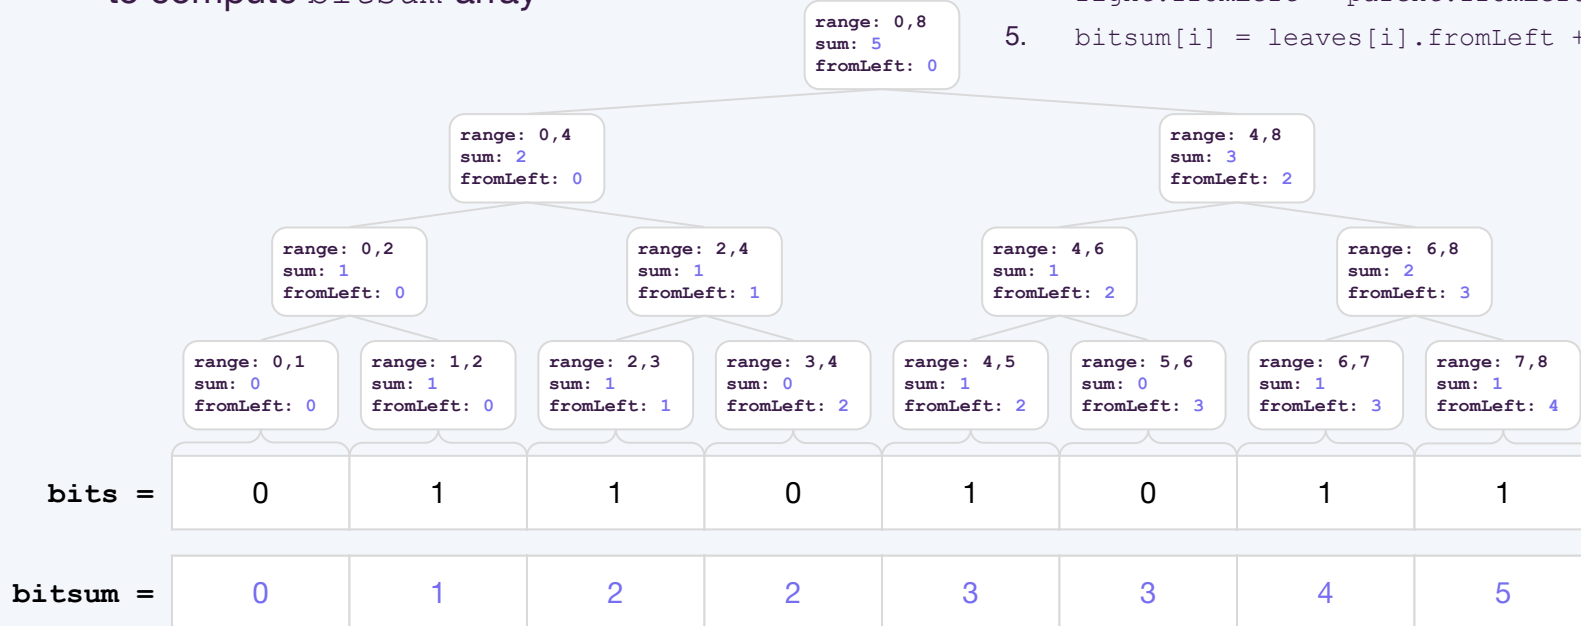
1. Divide problem into parallel pieces
cutoff = 1
2. Find sum at cutoff
leaves[i].sum = bits[i]
3. Propagate sum up
parent.sum = left.sum + right.sum
4. Propagate fromLeft down
left.fromLeft = parent.fromLeft
right.fromLeft = parent.fromLeft + left.sum



Problem 2

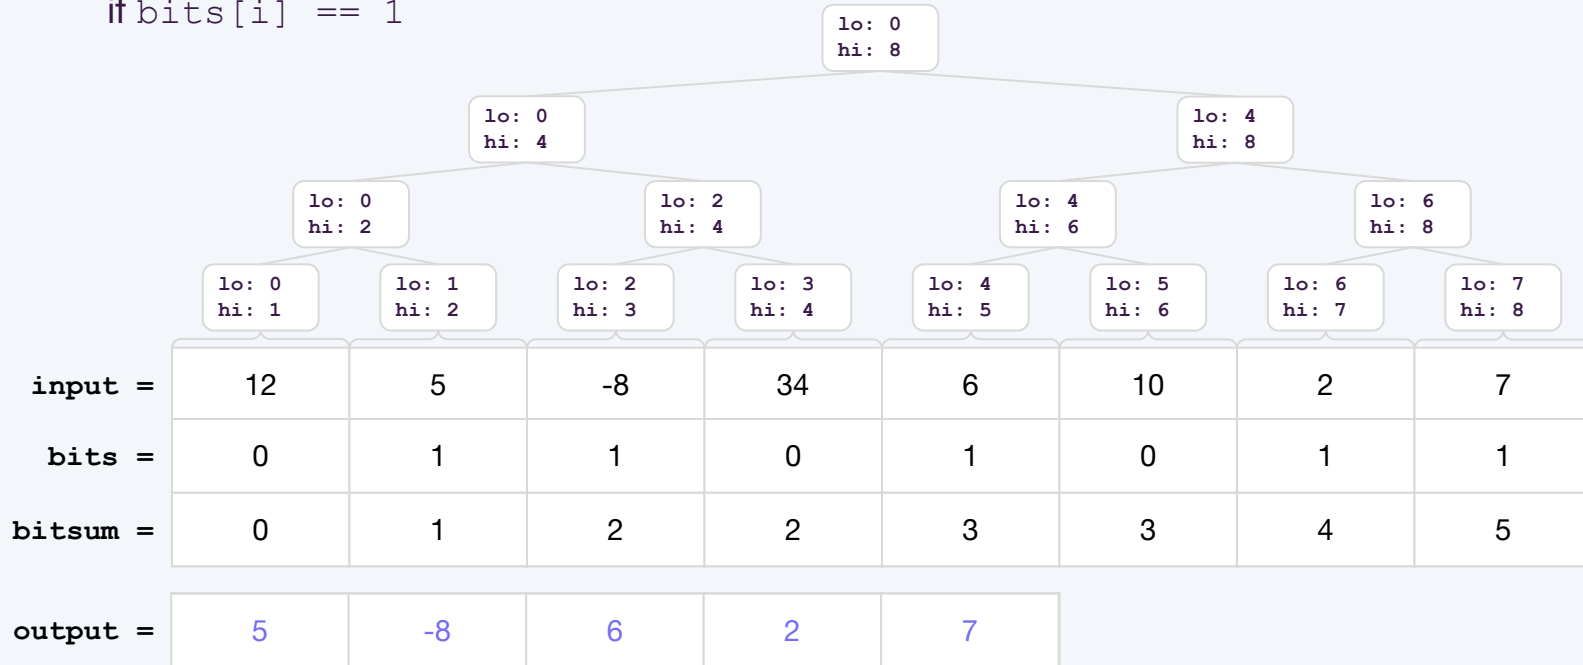
2. Parallel prefix sum on bits array to compute bitsum array

1. Divide problem into parallel pieces
cutoff = 1
2. Find sum at cutoff
`leaves[i].sum = bits[i]`
3. Propagate sum up
`parent.sum = left.sum + right.sum`
4. Propagate fromLeft down
`left.fromLeft = parent.fromLeft`
`right.fromLeft = parent.fromLeft + left.sum`
5. `bitsum[i] = leaves[i].fromLeft + bits[i]`



Problem 2

3. **Parallel map** to produce output array where $\text{output}[\text{bitsum}[i] - 1] = \text{input}[i]$ if $\text{bits}[i] == 1$



Problem 3

Problem 3

a) Define work and span.

Work is the runtime of a program on **one** processor.

Span is the runtime of a program on **infinitely** many processors.

b) How do we calculate work and span?

Work is the sum of all the work done by all processors

Span is the sum of all the work done by the longest dependence chain / parallel 'tree' branch

c) Does adding more processors affect the work or span?

Neither. Work is defined for one processor and span is defined for infinitely many processors.

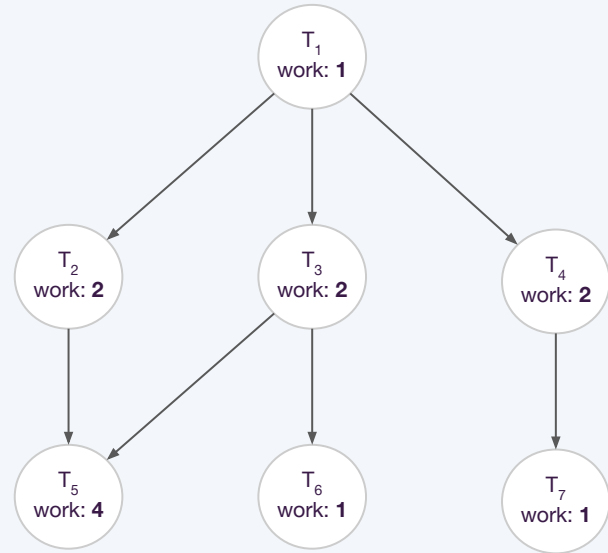
Adding more processors will not affect either work or span.

Problem 3

d) What is the work and span of the graph?

$$\begin{aligned} \text{Work} &= 1 + 2 + 2 + 2 + 4 + 1 + 1 \\ &= 13 \end{aligned}$$

$$\begin{aligned} \text{Span } (T_1 \rightarrow T_2 \rightarrow T_5) &= 1 + 2 + 4 \\ &= 7 \end{aligned}$$



Thank You!