

# Worksheet: 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 front of a vector, and packs the keys with a 1 to the end 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 (complete I-down) can be performed in parallel

	<b>[101 111 011 001 100 010 111 010]</b>	
<b>1. A =</b>	<b>[5 7 3 1 4 2 7 2]</b>	
<b>2. A&lt;0&gt; =</b>	<b>[1 1 1 1 0 0 1 0]</b>	<b>//lowest bit</b>
<b>3. A ← split(A, A&lt;0&gt;) =</b>	<b>[4 2 2 5 7 3 1 7]</b>	
<b>4. A&lt;1&gt; =</b>	<b>[0 1 1 0 1 1 0 1]</b>	<b>// middle bit</b>
<b>5. A ← split(A, A&lt;1&gt;) =</b>	<b>[4 5 1 2 2 7 3 7]</b>	
<b>6. A&lt;2&gt; =</b>	<b>[1 1 0 0 0 1 0 1]</b>	<b>// highest bit</b>
<b>7. A ← split(A, A&lt;2&gt;) =</b>	<b>[1 2 2 3 4 5 7 7]</b>	

<b>procedure split(A, Flags)</b>										
	I-down	←								
	I-up	←	<b>rev(n - i-scan(+, rev(Flags)) // rev = reverse</b>							
	<b>in parallel for each index i</b>									
	if (Flags[i])									
		Index[i]	←	I-up[i]						
	else									
		Index[i]	←	I-down[i]						
	result	←	permute(A, Index)							
A	=	[	5	7	3	1	4	2	7	2]
Flags	=	[	1	1	1	1	0	0	1	0]
I-down	=	[	0	0	0	0	0	1	2	2]
I-up	=	[	3	4	5	6	6	7	7	]
Index	=	[	3	4	5	6	0	1	7	2]
permute(A, Index)	=	[	4	2	2	5	7	3	1	7]

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.

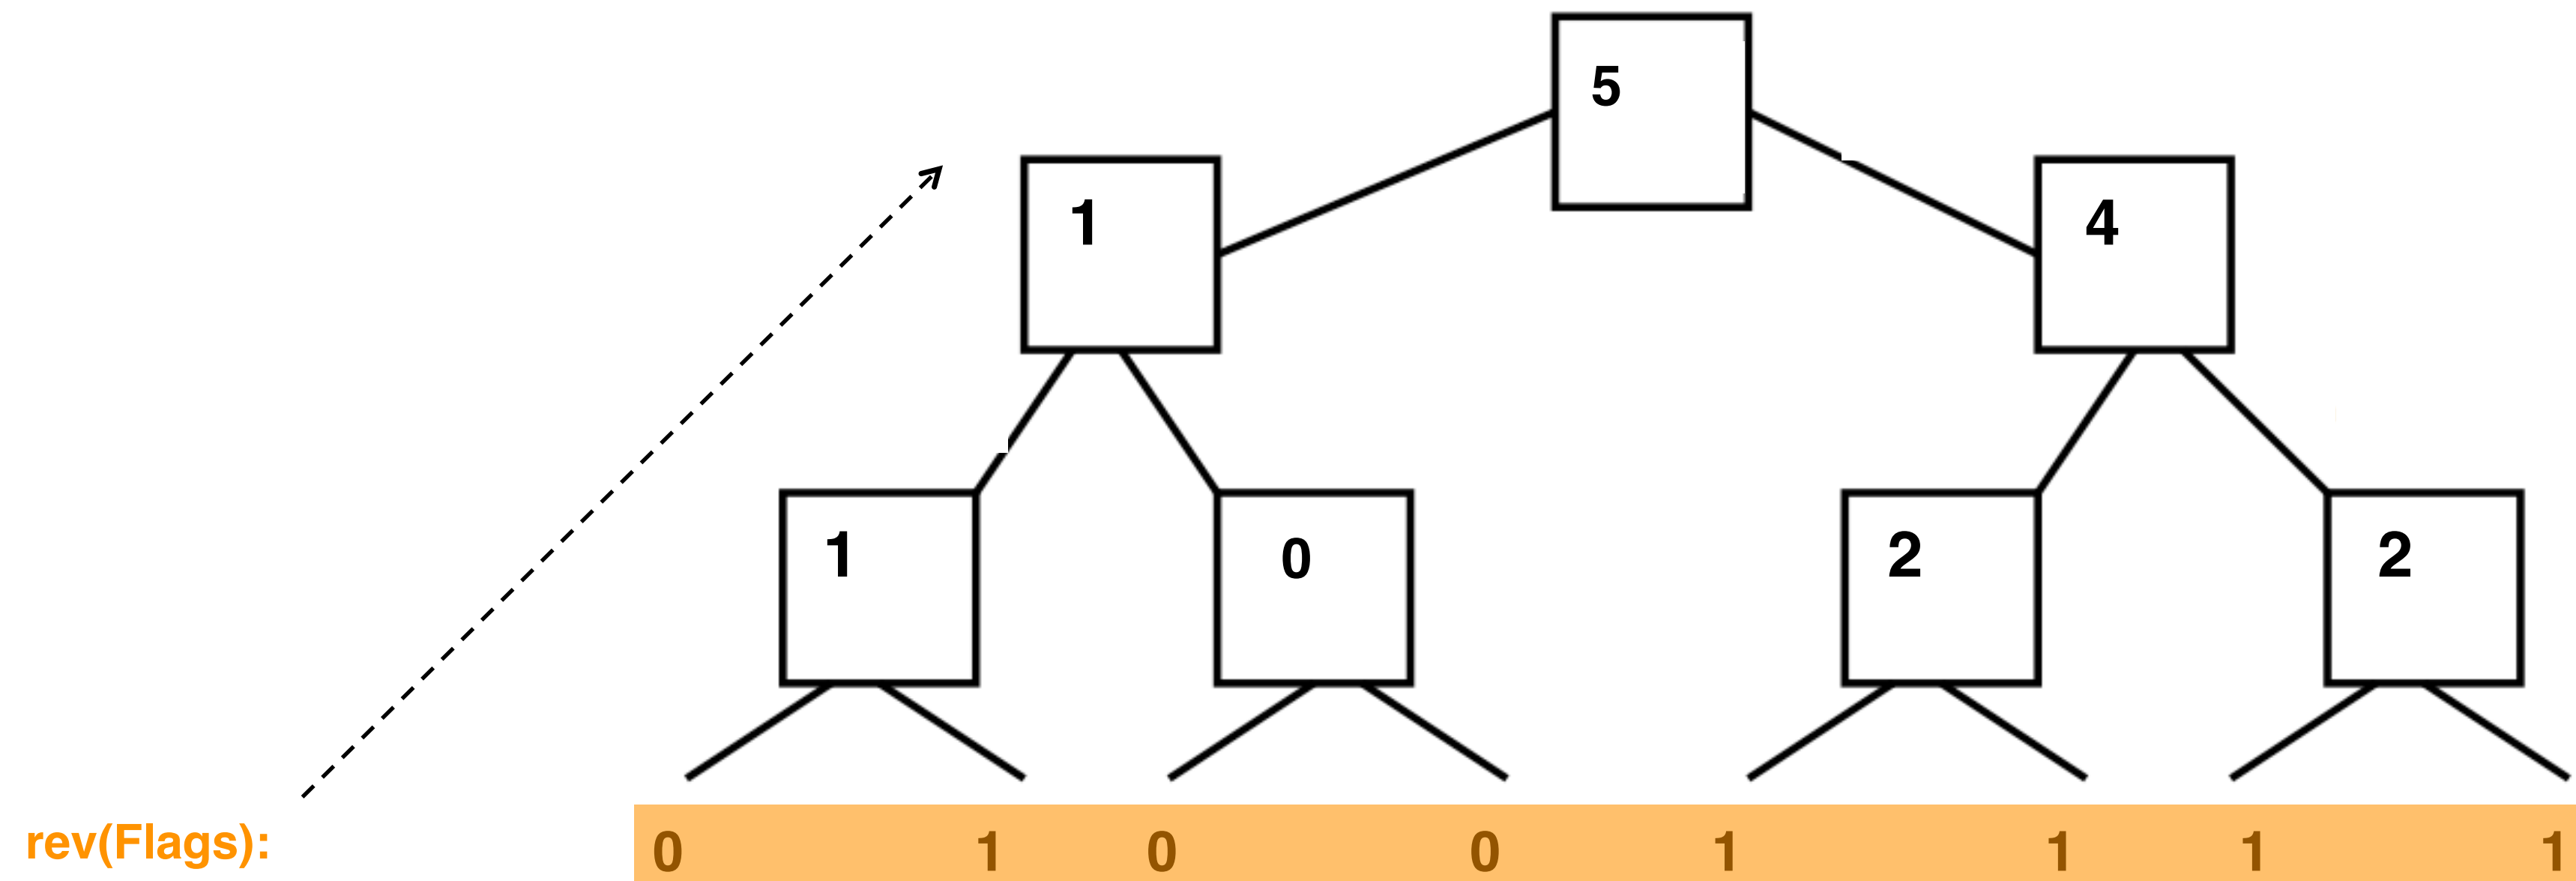


# Parallel Pre-scan Sum (I-Up): Upward Sweep

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

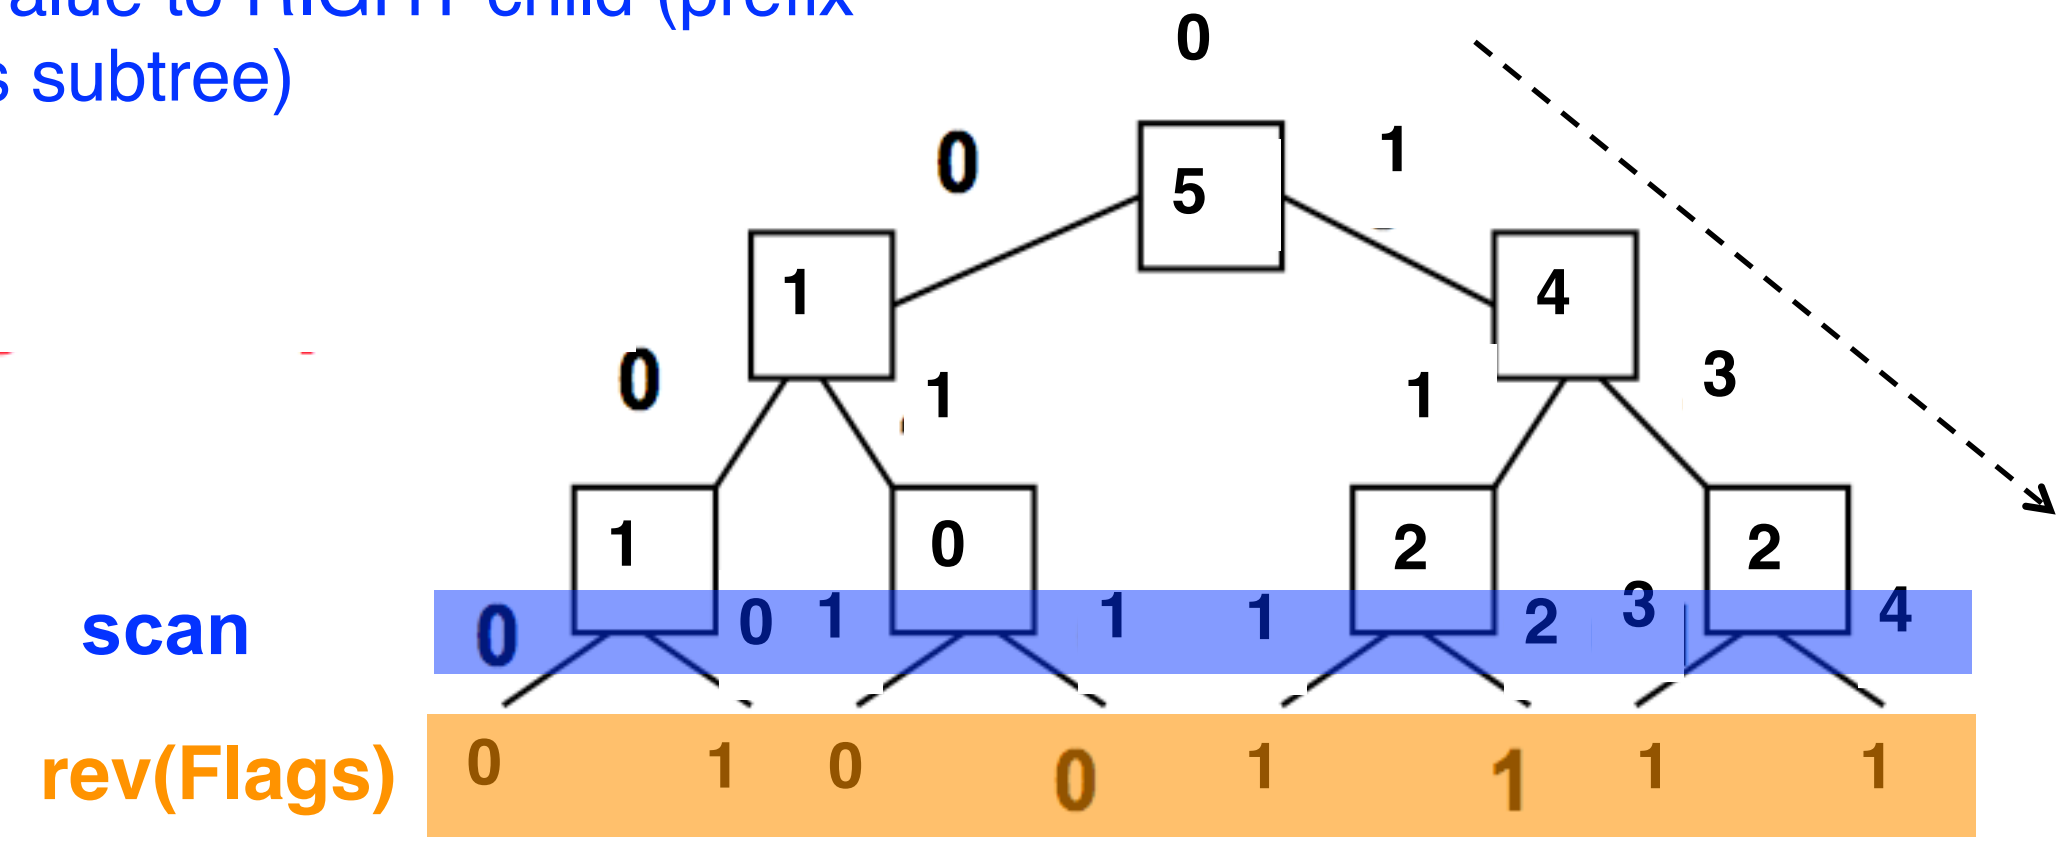
A	=	[ 5	7	3	1	4	2	7	2]
Flags	=	[ 1	1	1	1	0	0	1	0]
I-down	=	[ 0	0	0	0	0	1	2	2]
I-up	=	[ 3	4	5	6	6	6	7	7]
Index	=	[ 3	4	5	6	0	1	7	2]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7]



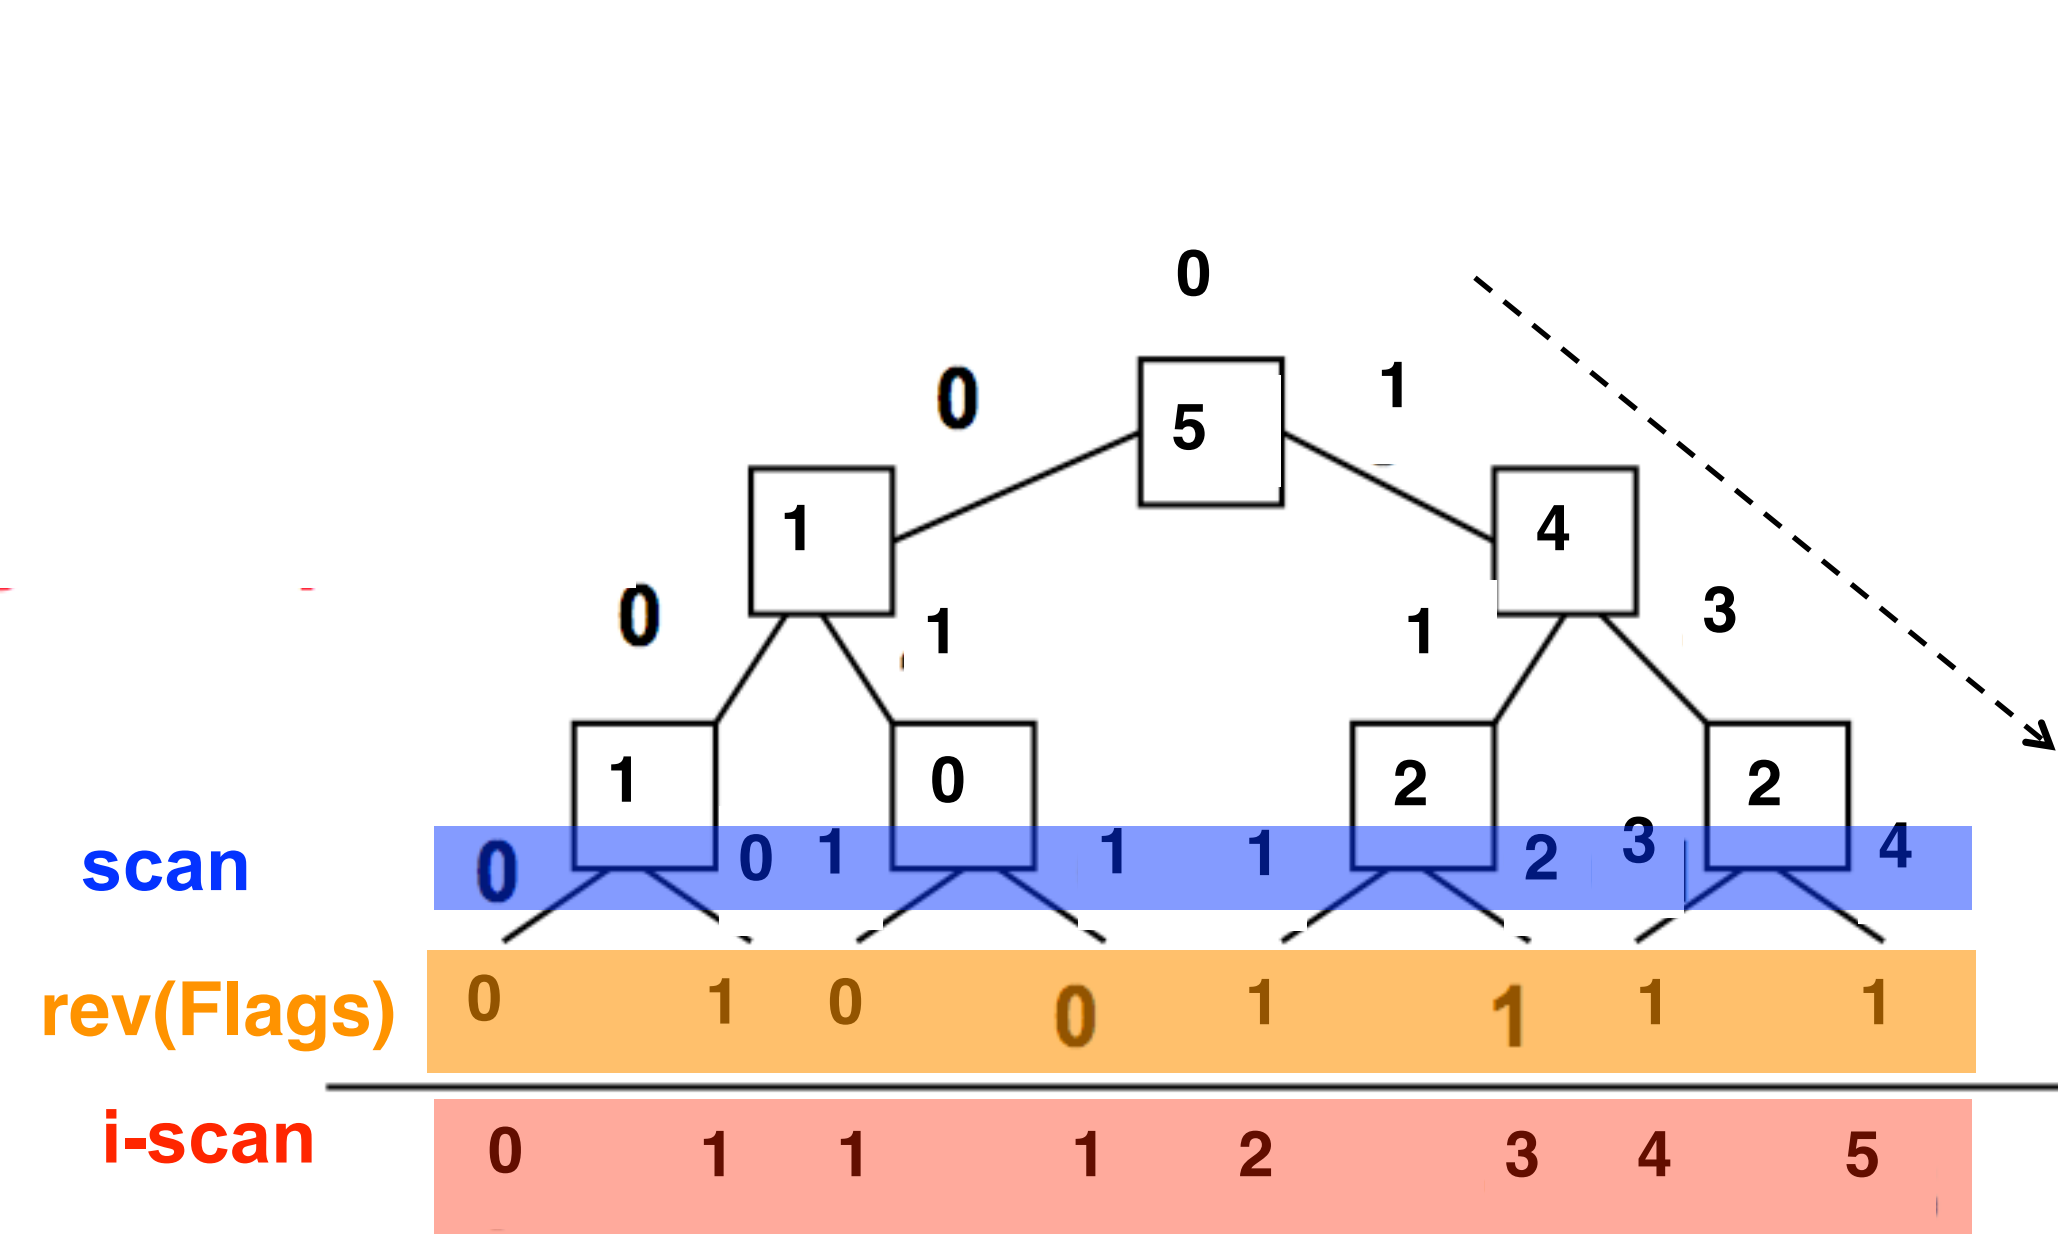
# Parallel Scan Sum (I-Up): Downward Sweep

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

A	=	[ 5	7	3	1	4	2	7	2 ]
Flags	=	[ 1	1	1	1	0	0	1	0 ]
I-down	=	[ 0	0	0	0	0	1	2	2 ]
I-up	=	[ 3	4	5	6	6	7	7	]
Index	=	[ 3	4	5	6	0	1	7	2 ]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7 ]



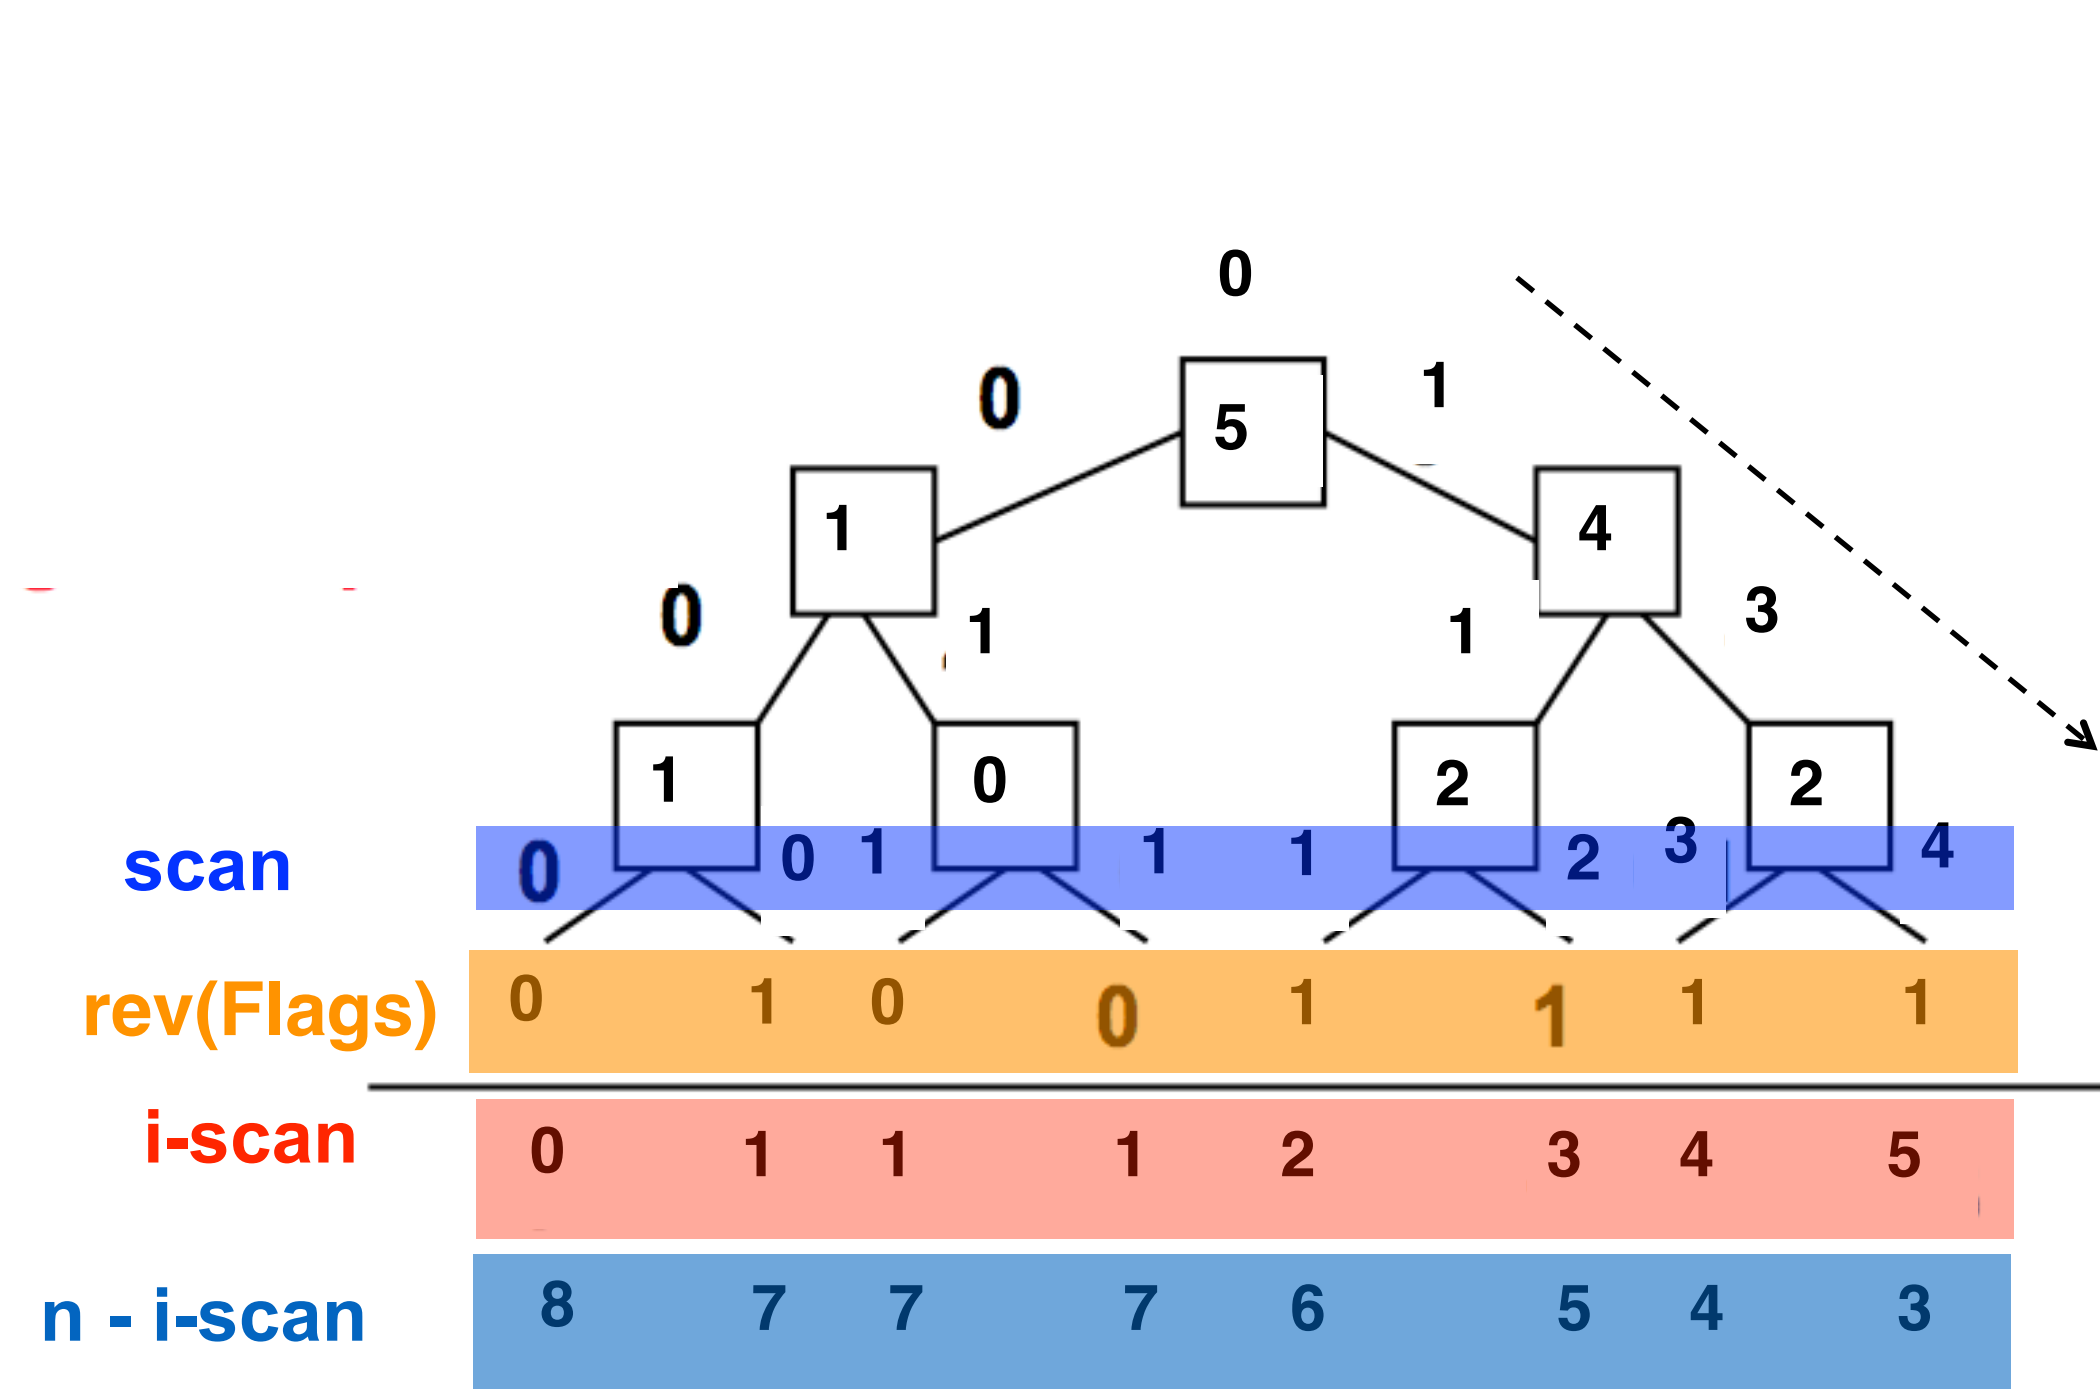
# Parallel Scan Sum (I-Up): Downward Sweep



A	=	[	5	7	3	1	4	2	7	2]
Flags	=	[	1	1	1	1	0	0	1	0]
I-down	=	[	0	0	0	0	0	1	2	2]
I-up	=	[	3	4	5	6	6	6	7	7]
Index	=	[	3	4	5	6	0	1	7	2]
permute(A, Index)	=	[	4	2	2	5	7	3	1	7]



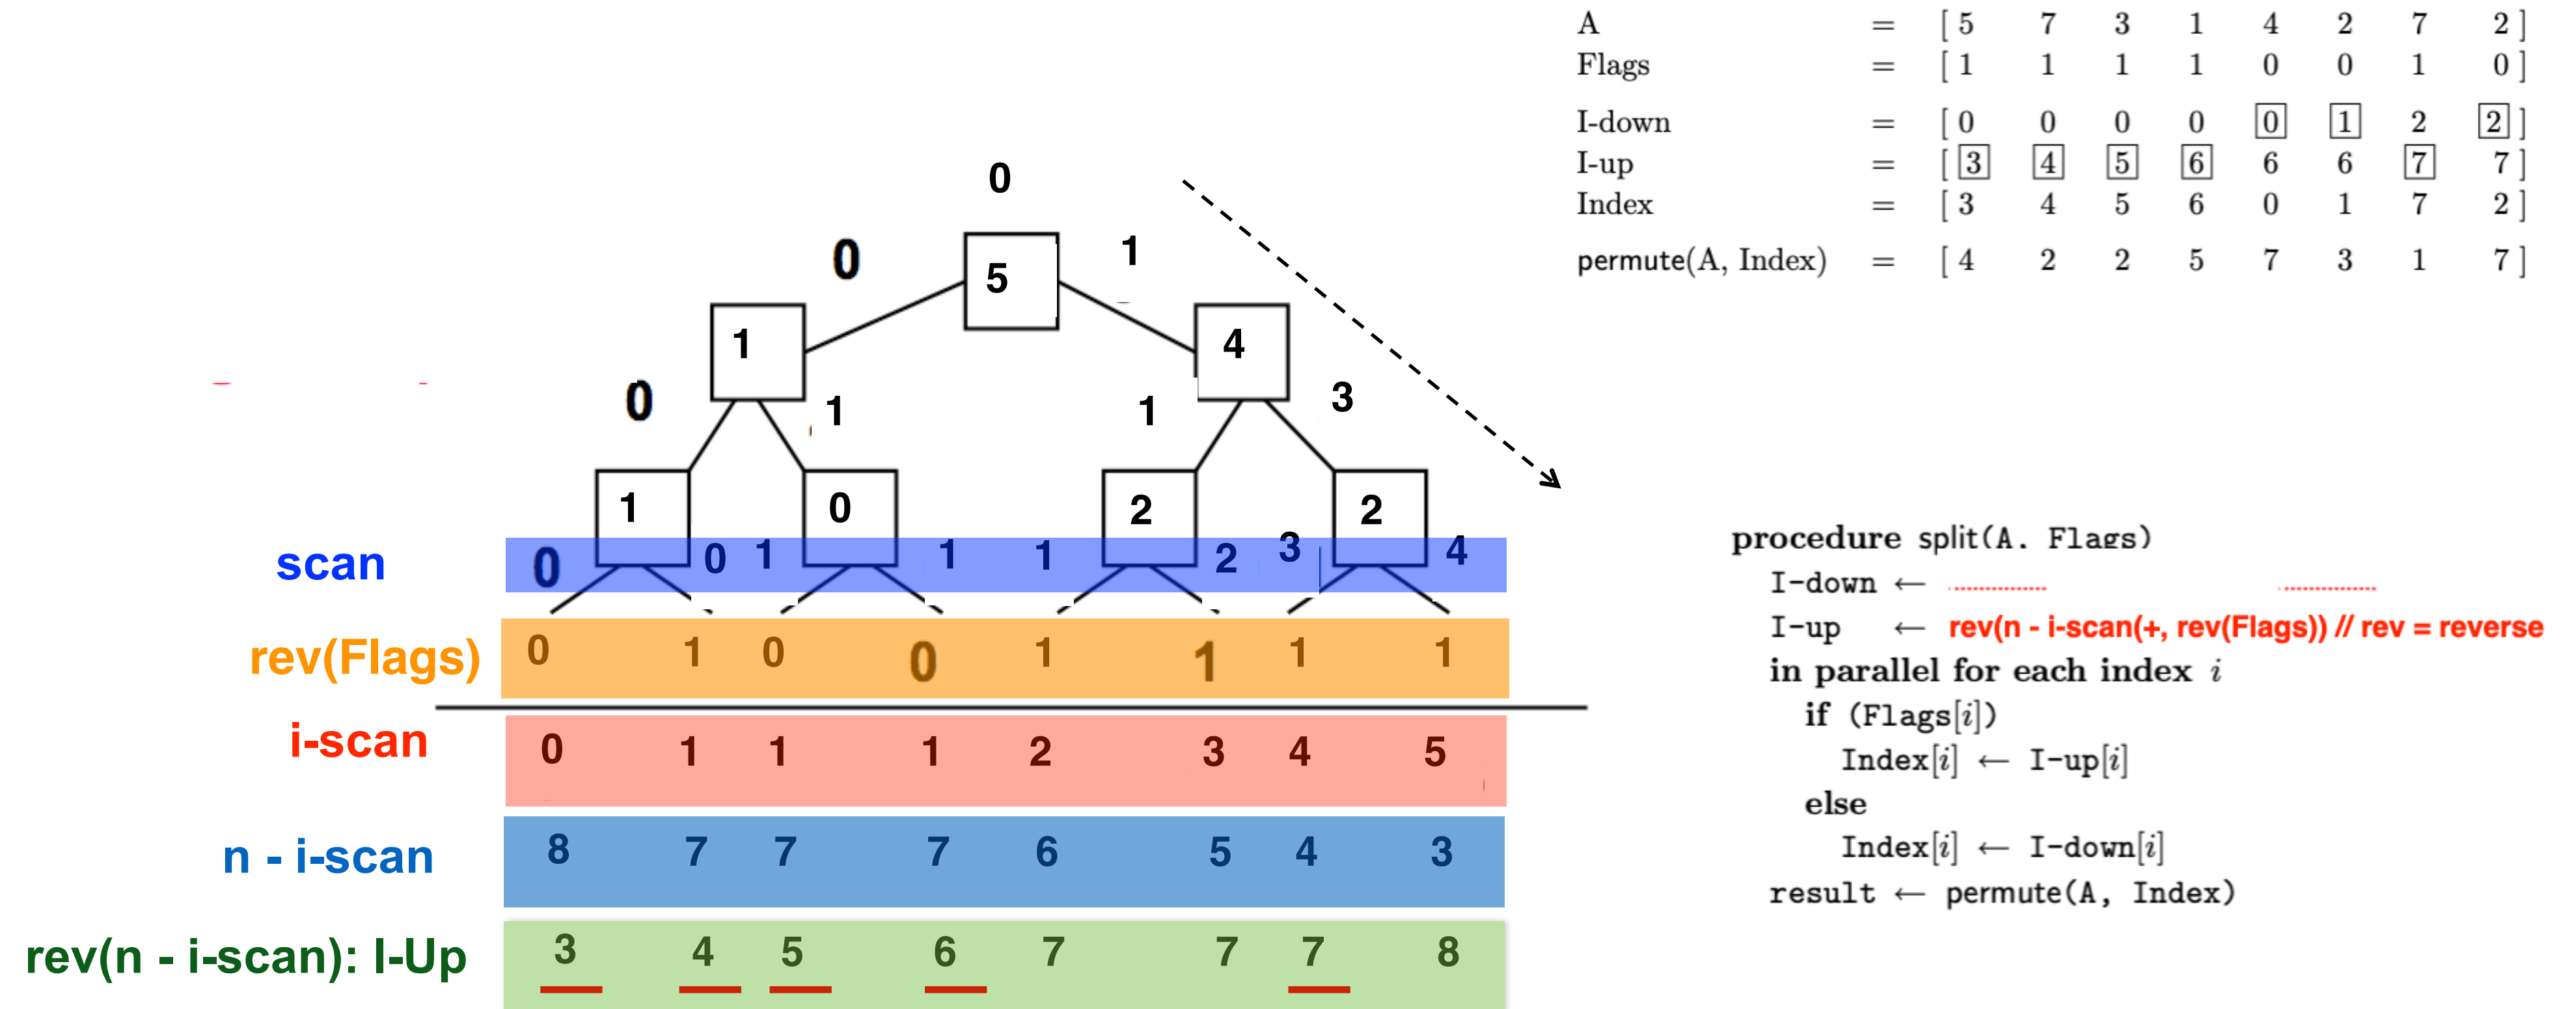
# Parallel Scan Sum (I-Up): Downward Sweep



A	=	[ 5	7	3	1	4	2	7	2 ]
Flags	=	[ 1	1	1	1	0	0	1	0 ]
I-down	=	[ 0	0	0	0	0	1	2	2 ]
I-up	=	[ 3	4	5	6	6	6	7	7 ]
Index	=	[ 3	4	5	6	0	1	7	2 ]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7 ]



# Parallel Scan Sum (I-Up): Downward Sweep



# Worksheet: 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 front of a vector, and packs the keys with a 1 to the end 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 (complete I-down) can be performed in parallel

		<b>[101 111 011 001 100 010 111 010]</b>	
1. <b>A</b> =		<b>[5 7 3 1 4 2 7 2]</b>	
2. <b>A{0}</b> =		<b>[1 1 1 1 0 0 1 0]</b>	<b>//lowest bit</b>
3. <b>A ← split(A, A{0})</b> =		<b>[4 2 2 5 7 3 1 7]</b>	
4. <b>A{1}</b> =		<b>[0 1 1 0 1 1 0 1]</b>	<b>// middle bit</b>
5. <b>A ← split(A, A{1})</b> =		<b>[4 5 1 2 2 7 3 7]</b>	
6. <b>A{2}</b> =		<b>[1 1 0 0 0 1 0 1]</b>	<b>// highest bit</b>
7. <b>A ← split(A, A{2})</b> =		<b>[1 2 2 3 4 5 7 7]</b>	

```

procedure split(A, Flags)
  I-down ←
  I-up ← rev(n - i-scan(+, rev(Flags)) // rev = reverse
  in parallel for each index i
    if (Flags[i])
      Index[i] ← I-up[i]
    else
      Index[i] ← I-down[i]
  result ← permute(A, Index)
  
```

A	=	[ 5	7	3	1	4	2	7	2]
Flags	=	[ 1	1	1	1	0	0	1	0]
I-down	=	[ 0	0	0	0	<b>[0]</b>	<b>[1]</b>	2	<b>[2]</b>
I-up	=	<b>[3]</b>	<b>[4]</b>	<b>[5]</b>	<b>[6]</b>	6	6	<b>[7]</b>	7]
Index	=	[ 3	4	5	6	0	1	7	2]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7]

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.



# Worksheet: 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 front of a vector, and packs the keys with a 1 to the end 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 (complete I-down) can be performed in parallel

	<span style="color: red;">[101 111 011 001 100 010 111 010]</span>	
1. $A =$	$[5 \ 7 \ 3 \ 1 \ 4 \ 2 \ 7 \ 2]$	
2. $A\langle 0 \rangle =$	$[1 \ 1 \ 1 \ 1 \ 0 \ 0 \ 1 \ 0]$	// lowest bit
3. $A \leftarrow \text{split}(A, A\langle 0 \rangle) =$	$[4 \ 2 \ 2 \ 5 \ 7 \ 3 \ 1 \ 7]$	
4. $A\langle 1 \rangle =$	$[0 \ 1 \ 1 \ 0 \ 1 \ 1 \ 0 \ 1]$	// middle bit
5. $A \leftarrow \text{split}(A, A\langle 1 \rangle) =$	$[4 \ 5 \ 1 \ 2 \ 2 \ 7 \ 3 \ 7]$	
6. $A\langle 2 \rangle =$	$[1 \ 1 \ 0 \ 0 \ 0 \ 1 \ 0 \ 1]$	// highest bit
7. $A \leftarrow \text{split}(A, A\langle 2 \rangle) =$	$[1 \ 2 \ 2 \ 3 \ 4 \ 5 \ 7 \ 7]$	

```

procedure split(A, Flags)
  I-down  $\leftarrow$  scan(+, not(Flags))
  I-up    $\leftarrow$  rev(n - i-scan(+, rev(Flags)) // rev = reverse
  in parallel for each index  $i$ 
    if (Flags[ $i$ ])
      Index[ $i$ ]  $\leftarrow$  I-up[ $i$ ]
    else
      Index[ $i$ ]  $\leftarrow$  I-down[ $i$ ]
  result  $\leftarrow$  permute(A, Index)
  
```

A	=	[ 5	7	3	1	4	2	7	2]
Flags	=	[ 1	1	1	1	0	0	1	0]
I-down	=	[ 0	0	0	0	<span style="border: 1px solid black; padding: 2px;">0</span>	<span style="border: 1px solid black; padding: 2px;">1</span>	2	<span style="border: 1px solid black; padding: 2px;">2</span> ]
I-up	=	[ <span style="border: 1px solid black; padding: 2px;">3</span>	<span style="border: 1px solid black; padding: 2px;">4</span>	<span style="border: 1px solid black; padding: 2px;">5</span>	<span style="border: 1px solid black; padding: 2px;">6</span>	6	6	<span style="border: 1px solid black; padding: 2px;">7</span>	7]
Index	=	[ 3	4	5	6	0	1	7	2]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7]

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.






# Parallel Scan Sum (I-Down): Upward Sweep

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

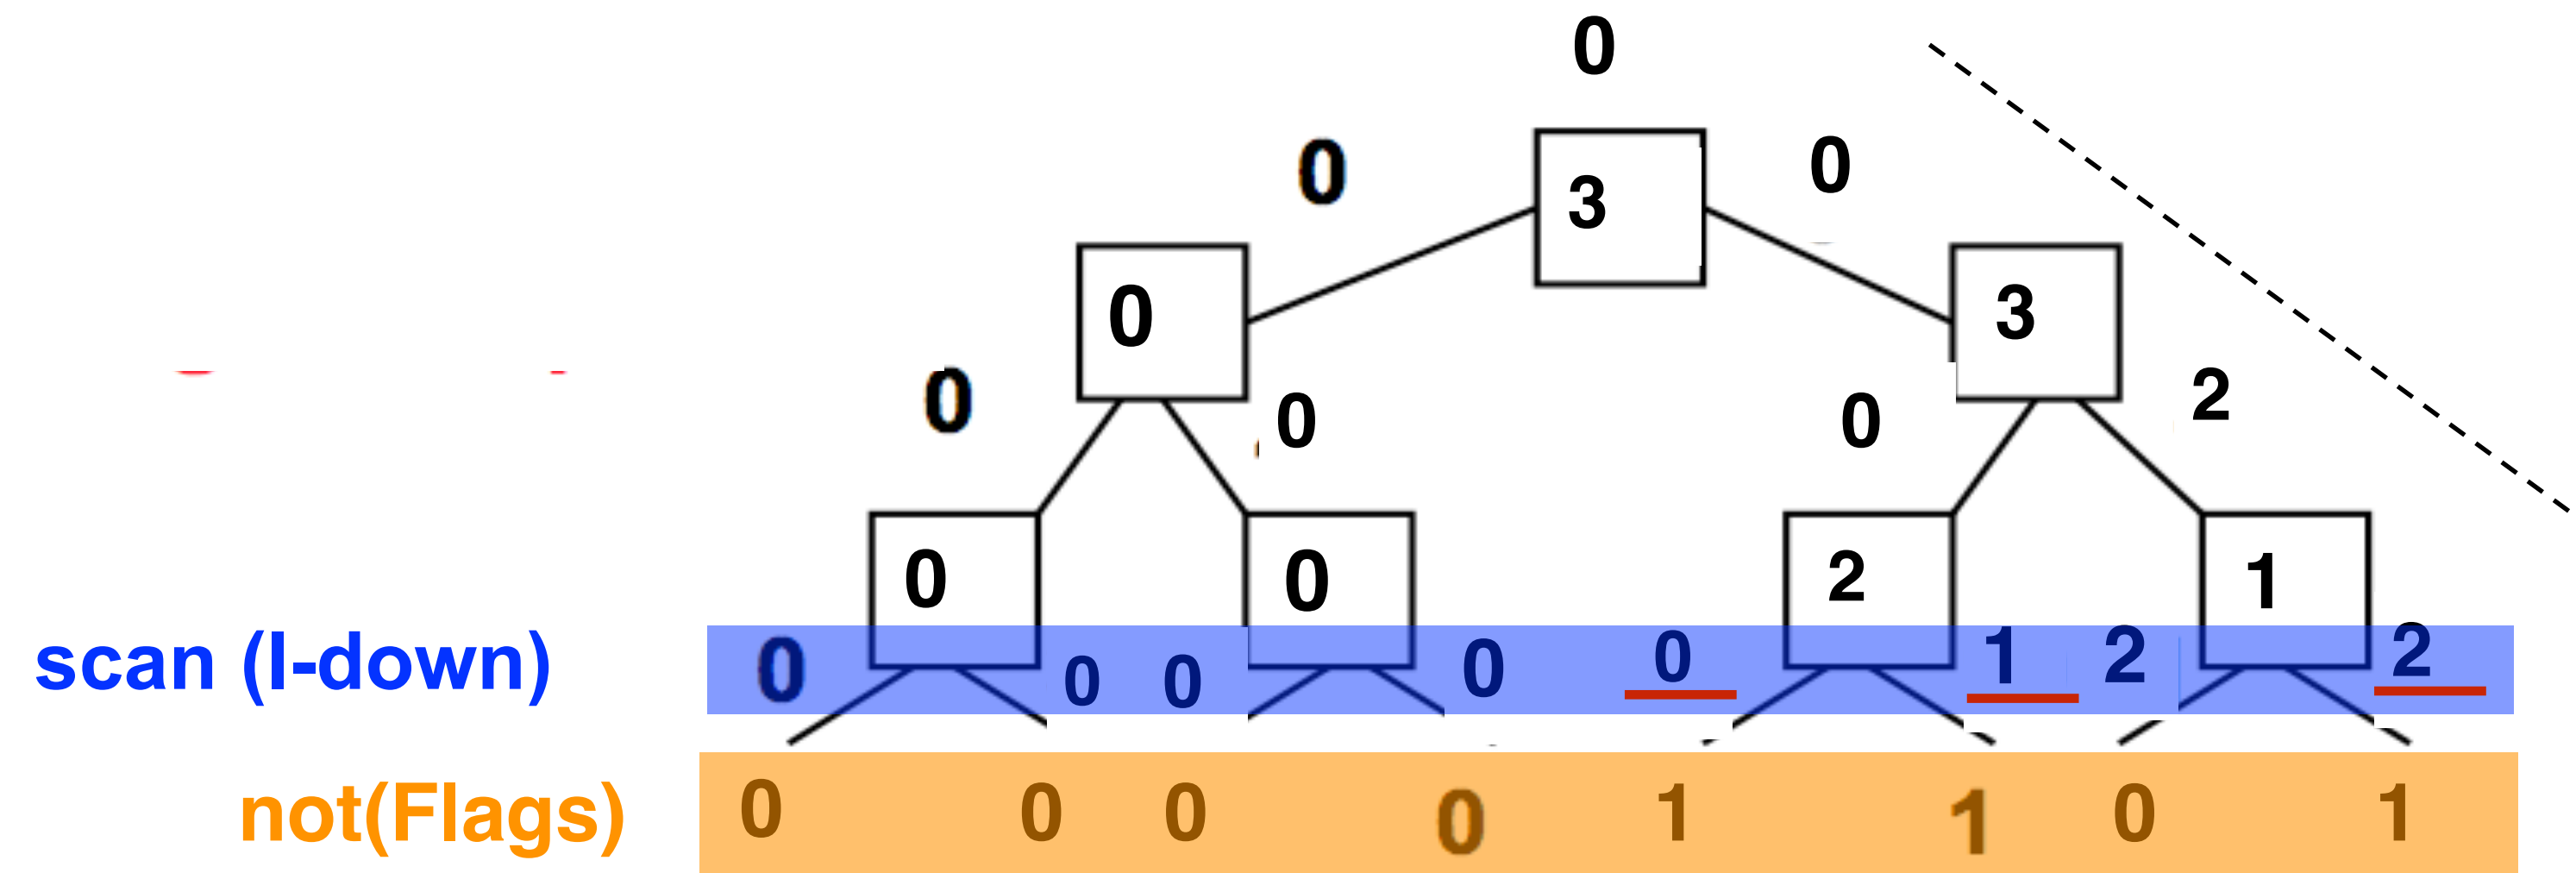
A	=	[ 5	7	3	1	4	2	7	2 ]
Flags	=	[ 1	1	1	1	0	0	1	0 ]
I-down	=	[ 0	0	0	0	0	1	2	2 ]
I-up	=	[ 3	4	5	6	6	6	7	7 ]
Index	=	[ 3	4	5	6	0	1	7	2 ]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7 ]



# Parallel Scan Sum (I-Down): Downward Sweep

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

A	=	[ 5	7	3	1	4	2	7	2 ]
Flags	=	[ 1	1	1	1	0	0	1	0 ]
I-down	=	[ 0	0	0	0	0	1	2	2 ]
I-up	=	[ 3	4	5	6	6	6	7	7 ]
Index	=	[ 3	4	5	6	0	1	7	2 ]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7 ]



```

procedure split(A, Flags)
  I-down ← .....
  I-up ← rev(n - i-scan(+, rev(Flags)) // rev = reverse
  in parallel for each index i
    if (Flags[i])
      Index[i] ← I-up[i]
    else
      Index[i] ← I-down[i]
  result ← permute(A, Index)
  
```



# Worksheet: 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 front of a vector, and packs the keys with a 1 to the end 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 (complete I-down) can be performed in parallel

	<span style="color: red;">[101 111 011 001 100 010 111 010]</span>	
1. $A =$	$[5 \ 7 \ 3 \ 1 \ 4 \ 2 \ 7 \ 2]$	
2. $A\langle 0 \rangle =$	$[1 \ 1 \ 1 \ 1 \ 0 \ 0 \ 1 \ 0]$	// lowest bit
3. $A \leftarrow \text{split}(A, A\langle 0 \rangle) =$	$[4 \ 2 \ 2 \ 5 \ 7 \ 3 \ 1 \ 7]$	
4. $A\langle 1 \rangle =$	$[0 \ 1 \ 1 \ 0 \ 1 \ 1 \ 0 \ 1]$	// middle bit
5. $A \leftarrow \text{split}(A, A\langle 1 \rangle) =$	$[4 \ 5 \ 1 \ 2 \ 2 \ 7 \ 3 \ 7]$	
6. $A\langle 2 \rangle =$	$[1 \ 1 \ 0 \ 0 \ 0 \ 1 \ 0 \ 1]$	// highest bit
7. $A \leftarrow \text{split}(A, A\langle 2 \rangle) =$	$[1 \ 2 \ 2 \ 3 \ 4 \ 5 \ 7 \ 7]$	

```

procedure split(A, Flags)
  I-down  $\leftarrow$  scan(+, not(Flags))
  I-up    $\leftarrow$  rev(n - i-scan(+, rev(Flags)) // rev = reverse
  in parallel for each index  $i$ 
    if (Flags[ $i$ ])
      Index[ $i$ ]  $\leftarrow$  I-up[ $i$ ]
    else
      Index[ $i$ ]  $\leftarrow$  I-down[ $i$ ]
  result  $\leftarrow$  permute(A, Index)
  
```

A	=	[ 5	7	3	1	4	2	7	2 ]
Flags	=	[ 1	1	1	1	0	0	1	0 ]
I-down	=	[ 0	0	0	0	0	1	2	2 ]
I-up	=	[ 3	4	5	6	6	6	7	7 ]
Index	=	[ 3	4	5	6	0	1	7	2 ]
permute(A, Index)	=	[ 4	2	2	5	7	3	1	7 ]

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.



# Another way of looking at the parallel algorithm

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

