

CS 577 - Union-Find & Binary Heaps

Manolis Vlatakis

Department of Computer Sciences
University of Wisconsin – Madison


Fall 2024



UNION-FIND


DISJOINT SETS

Imagine you have a set of objects, and you want to group them based on some criteria.




DISJOINT SETS

Imagine you have a set of objects, and you want to group them based on some criteria.



One possible grouping might look like this:



QUERIES ON DISJOINT SETS

While grouping the objects, you may ask:

- How can I merge more groups?
- How many distinct groups are there?
- Do two objects belong to the same group?

QUERIES ON DISJOINT SETS

While grouping the objects, you may ask:

- How can I merge more groups?
- How many distinct groups are there?
- Do two objects belong to the same group?

Union-Find Operations

The Union-Find data structure supports three operations:

QUERIES ON DISJOINT SETS

While grouping the objects, you may ask:

- How can I merge more groups?
- How many distinct groups are there?
- Do two objects belong to the same group?

Union-Find Operations

The Union-Find data structure supports three operations:

- **INIT:** Takes a set of elements $S = \{x_1, \dots, x_\ell\}$ and constructs ℓ singleton sets $S_i = \{x_i\}$ for all $i = 1, \dots, \ell$.

QUERIES ON DISJOINT SETS

While grouping the objects, you may ask:

- How can I merge more groups?
- How many distinct groups are there?
- Do two objects belong to the same group?

Union-Find Operations

The Union-Find data structure supports three operations:

- **INIT:** Takes a set of elements $S = \{x_1, \dots, x_\ell\}$ and constructs ℓ singleton sets $S_i = \{x_i\}$ for all $i = 1, \dots, \ell$.
- **UNION:** Takes two elements x and y and merges the groups containing x and y .

QUERIES ON DISJOINT SETS

While grouping the objects, you may ask:

- How can I merge more groups?
- How many distinct groups are there?
- Do two objects belong to the same group?


Union-Find Operations

The Union-Find data structure supports three operations:

- **INIT:** Takes a set of elements $S = \{x_1, \dots, x_\ell\}$ and constructs ℓ singleton sets $S_i = \{x_i\}$ for all $i = 1, \dots, \ell$.
- **UNION:** Takes two elements x and y and merges the groups containing x and y .
- **FIND:** Takes an element x and identifies the group containing x .

UNION-FIND DATA STRUCTURE

INIT-EXAMPLE




Explanation

Initially, each element starts in its own singleton set. This is the result of calling `Init`, where we have disjoint sets for each individual element.

UNION-FIND DATA STRUCTURE

UNION-EXAMPLE




Explanation

For example, initially, 0 and 2 belong to the same group, and 1 and 3 belong to another group. When we call the function `union(0, 1)`, it merges the groups containing 0 and 1.

UNION-FIND DATA STRUCTURE

FIND-EXAMPLE




Explanation

For example, when we call `find(2)`, it returns some signal for the set containing 2.

UNION-FIND DATA STRUCTURE

FIND-EXAMPLE



"Leader Element"

Each set should have a unique 'leader' element, which identifies the set.





Since the sets are always disjoint, the same object cannot be the leader of more than one set.

UNION-FIND DATA STRUCTURE

HIGH-LEVEL REPRESENTATION

One of the most easiest ways to represent sets is through graphs:

- Same set=Same connected component. Connect all objects that belong to the same set with edges, like this:







UNION-FIND DATA STRUCTURE

HIGH-LEVEL REPRESENTATION

One of the most easiest ways to represent sets is through graphs:

- Same set=Same tree/path:







UNION-FIND DATA STRUCTURE

HIGH-LEVEL REPRESENTATION

One of the most easiest ways to represent sets is through graphs:

- To determine the group an element belongs to, we designate a representative element for each group.

This representative acts as the identifier for the entire group:



Find Operations

- $\text{FIND}(0) = 0$
- $\text{FIND}(1) = 0$
- $\text{FIND}(2) = 0$
- $\text{FIND}(3) = 0$
- $\text{FIND}(4) = 4$
- $\text{FIND}(5) = 4$

UNION-FIND DATA STRUCTURE





HIGH-LEVEL REPRESENTATION

One of the most easiest ways to represent sets is through graphs:

- To determine if two elements belong to the same set:

 $\text{SAMESET}(x,y) = \text{True}$ if $\text{FIND}(x) = \text{FIND}(y)$ otherwise **False**

This representative acts as the identifier for the entire group:



Find Operations


- $\text{FIND}(0) = 0$
- $\text{FIND}(1) = 0$
- $\text{FIND}(2) = 0$
- $\text{FIND}(3) = 0$
- $\text{FIND}(4) = 4$
- $\text{FIND}(5) = 4$

UNION-FIND DATA STRUCTURE

HIGH-LEVEL REPRESENTATION

One of the most easiest ways to represent sets is through graphs:

🤖 What is the worst-case complexity of FIND if sets are represented by paths/chains for a data structure of n elements?




UNION-FIND DATA STRUCTURE


HIGH-LEVEL REPRESENTATION

One of the most easiest ways to represent sets is through graphs:

🤖 What is the worst-case complexity of FIND if sets are represented by paths/chains for a data structure of n elements?



Traversing the entire chain to find the representative results in FIND having a worst-case complexity of $\Omega(n)$



- $\text{FIND}(5) = 4$

REVERSED TREES

To make paths shorter, it's beneficial to represent them with trees.

- Each node points to another node, called its parent, except for the leader of each set, which points to itself and thus is the root of the tree.

INIT(x)


```
parent(x) ← x
```

FIND(x)

```
while (x ≠ parent(x))
  x ← parent(x)
return x
```

UNION(x, y)

```
x ← FIND(x)
y ← FIND(y)
parent(y) ← x
```



- INIT is trivial. $Cost[INIT] = \Theta(n)$.

REVERSED TREES

To make paths shorter, it's beneficial to represent them with trees.

- Each node points to another node, called its parent, except for the leader of each set, which points to itself and thus is the root of the tree.

INIT(x)


```
parent(x) ← x
```

FIND(x)

```
while (x ≠ parent(x))
  x ← parent(x)
return x
```

UNION(x, y)

```
x ← FIND(x)
y ← FIND(y)
parent(y) ← x
```



- FIND traverses the tree until the root is found. Worst-case cost equals the height of tree.

REVERSED TREES

To make paths shorter, it's beneficial to represent them with trees.

- Each node points to another node, called its parent, except for the leader of each set, which points to itself and thus is the root of the tree.

INIT(x)


```
parent(x) ← x
```

FIND(x)

```
while (x ≠ parent(x))
  x ← parent(x)
return x
```

UNION(x, y)

```
x ← FIND(x)
y ← FIND(y)
parent(y) ← x
```



- UNION just redirects the parent pointer of one leader to the other. $Cost[UNION] = \Theta(1)$.

REVERSED TREES

To make paths shorter, it's beneficial to represent them with trees.

- Each node points to another node, called its parent, except for the leader of each set, which points to itself and thus is the root of the tree.

INIT(x)

```
parent(x) ← x
```

FIND(x)


```
while (x ≠ parent(x))
  x ← parent(x)
return x
```

UNION(x, y)

```
x ← FIND(x)
y ← FIND(y)
parent(y) ← x
```

🤔 How do we avoid creating very tall sparse trees?

Union is straightforward, but if we're not careful, Find can become costly.




- UNION just redirects the parent pointer of one leader to the other. $Cost[UNION] = \Theta(1)$.

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

☹ If we always attach the taller tree to the shorter one, we might inadvertently create a long path-tree:

UNION(1, 2), UNION(1, 3), ..., UNION(1, n)



But what if we apply the inverse rule?

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Whenever we need to merge two trees, we always make the root of the shallower tree a child of the deeper one. This requires us to also maintain the depth of each tree, but this is quite easy.

MAKESET(x)

```
parent( $x$ )  $\leftarrow$   $x$   
depth( $x$ )  $\leftarrow$  0
```

FIND(x)

```
while  $x \neq$  parent( $x$ )  
   $x \leftarrow$  parent( $x$ )  
return  $x$ 
```

UNION(x, y)

```
 $\bar{x} \leftarrow$  Find( $x$ )  
 $\bar{y} \leftarrow$  Find( $y$ )  
if depth( $\bar{x}$ ) > depth( $\bar{y}$ )  
  parent( $\bar{y}$ )  $\leftarrow$   $\bar{x}$   
else  
  parent( $\bar{x}$ )  $\leftarrow$   $\bar{y}$   
  if depth( $\bar{x}$ ) = depth( $\bar{y}$ )  
    depth( $\bar{y}$ )  $\leftarrow$  depth( $\bar{y}$ ) + 1
```

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Whenever we need to merge two trees, we always make the root of the shallower tree a child of the deeper one. This requires us to also maintain the depth of each tree, but this is quite easy.

MAKESET(x)

```
parent( $x$ )  $\leftarrow x$   
depth( $x$ )  $\leftarrow 0$ 
```

FIND(x)

```
while  $x \neq \text{parent}(x)$   
   $x \leftarrow \text{parent}(x)$   
return  $x$ 
```

UNION(x, y)

```
 $\bar{x} \leftarrow \text{Find}(x)$   
 $\bar{y} \leftarrow \text{Find}(y)$   
if depth( $\bar{x}$ ) > depth( $\bar{y}$ )  
  parent( $\bar{y}$ )  $\leftarrow \bar{x}$   
else  
  parent( $\bar{x}$ )  $\leftarrow \bar{y}$   
  if depth( $\bar{x}$ ) = depth( $\bar{y}$ )  
    depth( $\bar{y}$ )  $\leftarrow \text{depth}(\bar{y}) + 1$ 
```

When depth(\bar{x}) reaches d for the first time, \bar{x} becomes the leader of two merged sets, each with leaders of depth $d - 1$.

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

- **Base Case:** If $\text{depth}(\bar{x}) = 0$, then \bar{x} is the leader of a singleton set, so the size of \bar{x} 's set is $2^0 = 1$.

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

- **Base Case:** If $\text{depth}(\bar{x}) = 0$, then \bar{x} is the leader of a singleton set, so the size of \bar{x} 's set is $2^0 = 1$.
- **Inductive Step:** Assume that for any set leader \bar{y} with depth $d - 1$, the size of \bar{y} 's set is at least 2^{d-1} .

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

- **Base Case:** If $\text{depth}(\bar{x}) = 0$, then \bar{x} is the leader of a singleton set, so the size of \bar{x} 's set is $2^0 = 1$.
- **Inductive Step:** Assume that for any set leader \bar{y} with depth $d - 1$, the size of \bar{y} 's set is at least 2^{d-1} .
 - When $\text{depth}(\bar{x})$ becomes d for the first time, \bar{x} is the leader of the union of two sets, both of whose leaders had depth $d - 1$.

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

- **Base Case:** If $\text{depth}(\bar{x}) = 0$, then \bar{x} is the leader of a singleton set, so the size of \bar{x} 's set is $2^0 = 1$.
- **Inductive Step:** Assume that for any set leader \bar{y} with depth $d - 1$, the size of \bar{y} 's set is at least 2^{d-1} .
 - When $\text{depth}(\bar{x})$ becomes d for the first time, \bar{x} is the leader of the union of two sets, both of whose leaders had depth $d - 1$.
 - By the inductive hypothesis, both component sets had at least 2^{d-1} elements.

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

- **Base Case:** If $\text{depth}(\bar{x}) = 0$, then \bar{x} is the leader of a singleton set, so the size of \bar{x} 's set is $2^0 = 1$.
- **Inductive Step:** Assume that for any set leader \bar{y} with depth $d - 1$, the size of \bar{y} 's set is at least 2^{d-1} .
 - When $\text{depth}(\bar{x})$ becomes d for the first time, \bar{x} is the leader of the union of two sets, both of whose leaders had depth $d - 1$.
 - By the inductive hypothesis, both component sets had at least 2^{d-1} elements.
 - Therefore, the new set has at least 2^d elements.

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

- **Base Case:** If $\text{depth}(\bar{x}) = 0$, then \bar{x} is the leader of a singleton set, so the size of \bar{x} 's set is $2^0 = 1$.
- **Inductive Step:** Assume that for any set leader \bar{y} with depth $d - 1$, the size of \bar{y} 's set is at least 2^{d-1} .
 - When $\text{depth}(\bar{x})$ becomes d for the first time, \bar{x} is the leader of the union of two sets, both of whose leaders had depth $d - 1$.
 - By the inductive hypothesis, both component sets had at least 2^{d-1} elements.
 - Therefore, the new set has at least 2^d elements.

Contrapositive Trick: If $A \Rightarrow B$, then $\neg B \Rightarrow \neg A$.

For any leader \bar{x} , if the size of the tree of \bar{x} is strictly less than 2^k , then $\text{depth}(\bar{x}) < k$.

Explanation: If $\text{depth}(\bar{x}) \geq k$, then by above theorem, the tree would have more than 2^k nodes.
Contradiction.

UNION BY DEPTH

MAKING FIND OPERATION EFFICIENT

Theorem. For any leader \bar{x} , the size of the tree of \bar{x} is at least $2^{\text{depth}(\bar{x})}$.

Proof by Induction:

- **Base Case:** If $\text{depth}(\bar{x}) = 0$, then \bar{x} is the leader of a singleton set, so the size of \bar{x} 's set is $2^0 = 1$.
- **Inductive Step:** Assume that for any set leader \bar{y} with depth $d - 1$, the size of \bar{y} 's set is at least 2^{d-1} .
 - When $\text{depth}(\bar{x})$ becomes d for the first time, \bar{x} is the leader of the union of two sets, both of whose leaders had depth $d - 1$.
 - By the inductive hypothesis, both component sets had at least 2^{d-1} elements.
 - Therefore, the new set has at least 2^d elements.

Conclusion

Since there are at most n elements in total, the maximum depth of any set is $\log n$. Therefore, both FIND and UNION run in $\Theta(\log n)$ time in the worst case.

PATH COMPRESSION

OPTIMIZING FIND OPERATION


- Path compression flattens the structure, making future FIND operations quicker.
- Path compression makes every node on the FIND path from x to the root point directly to the root.

FIND(x)

```

if  $x \neq \text{parent}(x)$ 
   $\text{parent}(x) \leftarrow$ 
   $\text{FIND}(\text{parent}(x))$ 
return  $\text{parent}(x)$ 

```



PATH COMPRESSION

OPTIMIZING FIND OPERATION


- Path compression flattens the structure, making future FIND operations quicker.
- Path compression makes every node on the FIND path from x to the root point directly to the root.

FIND(x)

```

if  $x \neq \text{parent}(x)$ 
   $\text{parent}(x) \leftarrow$ 
   $\text{FIND}(\text{parent}(x))$ 
return  $\text{parent}(x)$ 

```



Nodes on the FIND path (highlighted in orange) are directly attached to the root.

PATH COMPRESSION

OPTIMIZING FIND OPERATION


- Path compression flattens the structure, making future FIND operations quicker.
- Path compression makes every node on the FIND path from x to the root point directly to the root.

FIND(x)

```

if  $x \neq \text{parent}(x)$ 
   $\text{parent}(x) \leftarrow$ 
   $\text{FIND}(\text{parent}(x))$ 
return  $\text{parent}(x)$ 

```



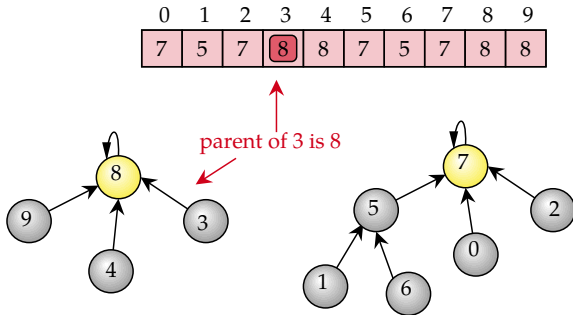
Teaser for Discussion Section - A $\log^*(n)$ Analysis

We will perform an amortized analysis showing how path compression can provide a super-exponential improvement, making the FIND operation almost constant time.

UNION-FIND

A FINAL THOUGHT ON ARRAY REPRESENTATION

- Union-Find can be viewed as a reversed search tree.
- Efficiency comes from the simplicity of tracking roots with only parent pointers.
 - Use an array `parent[]` of length n .
 - `parent[i] = j` means the parent of element i is element j .




DEMYSTIFY BINARY SEARCH TREES

INTRODUCTION TO BINARY SEARCH TREES

WHY BALANCED TREES?

- We can construct any tree we want, but in this course, our focus is on efficiency.
- From the union-find example, we learned that it's better to have a balanced tree rather than a skewed one.

🤖 Can you find the number 7 if you start searching from the root of the tree?



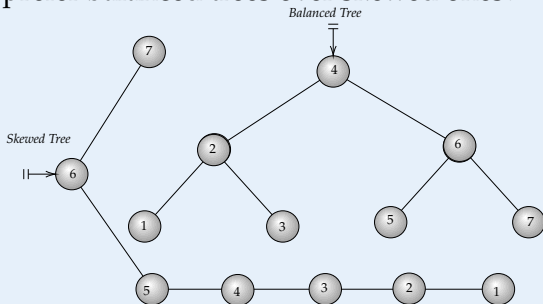
INTRODUCTION TO BINARY SEARCH TREES

WHY BALANCED TREES?

- We can construct any tree we want, but in this course, our focus is on efficiency.
- From the union-find example, we learned that it's better to have a balanced tree rather than a skewed one.



Why do we prefer balanced trees over skewed ones?



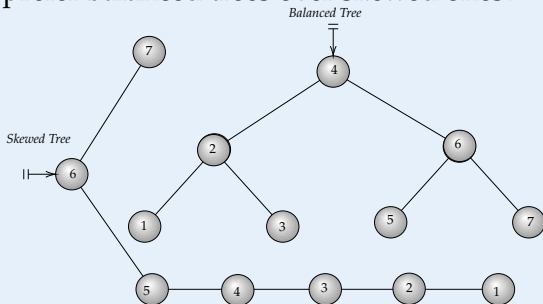
INTRODUCTION TO BINARY SEARCH TREES

WHY BALANCED TREES?

- We can construct any tree we want, but in this course, our focus is on efficiency.

The complexity of an efficient traversal is proportional to the height!

🤖 Why do we prefer balanced trees over skewed ones?




BINARY BALANCED TREE

Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

🤖 Which level is not completely filled in this tree?



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

Which level is not completely filled in this tree?



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

🤖 Which rule did we violate when filling the last level of this tree?



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

Which rule did we violate when filling the last level of this tree?



A NOTE FOR CODERS

- Constructing trees in programming often requires manipulating pointers and structures.


```
typedef struct treenode {           /* tree node struct */
    int data;                       /* integer field */
    struct treenode *left;         /* pointer to the left child */
    struct treenode *right;       /* pointer to the right child */
} node;
```


A NOTE FOR CODERS

- Constructing trees in programming often requires manipulating pointers and structures.

```
typedef struct treenode {           /* tree node struct */
    int data;                       /* integer field */
    struct treenode *left;         /* pointer to the left child */
    struct treenode *right;       /* pointer to the right child */
} node;
```


- For balanced trees, we can represent them simply using an array:
 - The children of node i are at positions $2i + 1$ and $2i + 2$.
 - The parent of node i is at position $(i - 1) \div 2$.



EXAMPLE OF ARRAY REPRESENTATION

Rule of Thumb

- The children of node i are at positions $2i + 1$ and $2i + 2$.
- The parent of node i is at position $(i - 1) \div 2$.




- The root is at position 0, with children at positions 1 and 2.

$$1 = 2 * 0 + 1 \quad 2 = 2 * 0 + 2 \quad \text{No parent since } (0 - 1) \div 2 < 0$$

EXAMPLE OF ARRAY REPRESENTATION

Rule of Thumb

- The children of node i are at positions $2i + 1$ and $2i + 2$.
- The parent of node i is at position $(i - 1) \div 2$.



- The node at position 1 has its parent at position 0 and children at positions 3 and 4.

EXAMPLE OF ARRAY REPRESENTATION

Rule of Thumb

- The children of node i are at positions $2i + 1$ and $2i + 2$.
- The parent of node i is at position $(i - 1) \div 2$.

0		2			5	6
P		★			C	C

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Height \rightarrow Size

Question 1: If I give you a complete tree of height h , how many nodes does it have?


Definition of Height

- The height of a tree is the length of the path from the root to its farthest leaf node.
- A tree with only a root node has a height of 0, while an empty tree has a height of -1.

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Height \rightarrow Size

Question 1: If I give you a complete tree of height h , how many nodes does it have?



The total number of nodes is:

$$1 + 2 + 4 + \dots + 2^h = 2^{h+1} - 1$$

A NOTE FOR MATH ENTHUSIASTS


Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?



- Reverse engineering the previous equation from bottom to top:


$$1 + 2 + 4 + \dots + 2^{h-1} + 2^h = 2^{h+1} - 1$$

- **Remark:** 50% of nodes are leaves 😊

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?



- Reverse engineering the previous equation from bottom to top:


$$1 + 2 + 4 + \dots + 2^{h-1} + 2^h = \underbrace{2^{h+1} - 1}_{n-1}$$

- **Remark:** 50% of nodes are leaves 😊

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?



- Reverse engineering the previous equation from bottom to top:


$$1 + 2 + 4 + \dots + 2^{h-1} + \underbrace{2^h}_{\frac{n}{2}} = \underbrace{2^{h+1} - 1}_{n-1}$$

- **Remark:** 50% of nodes are leaves 😊

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?



- Reverse engineering the previous equation from bottom to top:


$$1 + 2 + 4 + \dots + \underbrace{2^{h-1}}_{\frac{n}{4}} + \underbrace{2^h}_{\frac{n}{2}} = \underbrace{2^{h+1} - 1}_{n-1}$$

- Remark:** 50% of nodes are leaves 😊

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?



- Reverse engineering the previous equation from bottom to top:


$$1 + 2 + \underbrace{4}_{\frac{n}{2^{h-1}}} + \dots + \underbrace{2^{h-1}}_{\frac{n}{4}} + \underbrace{2^h}_{\frac{n}{2}} = \underbrace{2^{h+1} - 1}_{n-1}$$

- **Remark:** 50% of nodes are leaves ☺

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?



- Reverse engineering the previous equation from bottom to top:


$$1 + \underbrace{2}_{\frac{n}{2^h}} + \underbrace{4}_{\frac{n}{2^{h-1}}} + \dots + \underbrace{2^{h-1}}_{\frac{n}{4}} + \underbrace{2^h}_{\frac{n}{2}} = \underbrace{2^{h+1} - 1}_{n-1}$$

- Remark:** 50% of nodes are leaves ☺

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Number of Nodes Per Level

Question 2: If I give you a complete tree with $n - 1$ total nodes, how many nodes are at each level?



- Reverse engineering the previous equation from bottom to top:

$$\underbrace{1}_{\frac{n}{2^{h+1}}} + \underbrace{2}_{\frac{n}{2^h}} + \underbrace{4}_{\frac{n}{2^{h-1}}} + \dots + \underbrace{2^{h-1}}_{\frac{n}{4}} + \underbrace{2^h}_{\frac{n}{2}} = \underbrace{2^{h+1} - 1}_{n-1}$$

- Remark:** 50% of nodes are leaves ☺

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Size \rightarrow Height

Question 3: If I give you a complete tree with n nodes, what is its height?

A NOTE FOR MATH ENTHUSIASTS

Binary Balanced Trees: Size \rightarrow Height

Question 3: If I give you a complete tree with n nodes, what is its height?


- We know that a tree of height h has $2^{h+1} - 1$ nodes.
- Therefore:

$$n = 2^{h+1} - 1 \quad \Rightarrow \quad h = \log_2 n - 1 = \Theta(\log n)$$

A NOTE FOR MATH ENTHUSIASTS

Search Efficiency in a Balanced Binary Search Tree

Question 4: In a well-balanced and sorted binary search tree, how efficient is your algorithm in finding a number?



- Given the logarithmic height of the tree, the search operation will take $\Theta(\log n)$ time.

BINARY HEAPS

EARLIEST DEADLINE FIRST

A CLASSICAL PROBLEM FROM OPERATING SYSTEMS

- Imagine you have multiple processes to be executed.
- Each process comes with a deadline.
- New processes keep arriving at the operating system.

EARLIEST DEADLINE FIRST

A CLASSICAL PROBLEM FROM OPERATING SYSTEMS

- Imagine you have multiple processes to be executed.
- Each process comes with a deadline.
- New processes keep arriving at the operating system.

EDF Task

The CPU or server always executes the process with the earliest deadline.



EARLIEST DEADLINE FIRST

TRANSFORMING A DYNAMIC PROBLEM INTO A DATA STRUCTURE PROBLEM

Intuitive Goal

- We need to efficiently manage processes as they arrive and meet their deadlines.
- The challenge is to design a data structure that supports all required operations in the best possible time.

EARLIEST DEADLINE FIRST

TRANSFORMING A DYNAMIC PROBLEM INTO A DATA STRUCTURE PROBLEM

Intuitive Goal

- We need to efficiently manage processes as they arrive and meet their deadlines.
- The challenge is to design a data structure that supports all required operations in the best possible time.

Objective. Design an efficient data structure for managing processes.

Let's consider a set of processes $\mathcal{P} := \{p \mid \text{set of } n \text{ processes}\}$.

- **FINDMIN**(\mathcal{P}): Identify the process with the earliest deadline.
- **DELETEMIN**(\mathcal{P}): Remove the process with the earliest deadline.

EARLIEST DEADLINE FIRST

TRANSFORMING A DYNAMIC PROBLEM INTO A DATA STRUCTURE PROBLEM

Intuitive Goal

- We need to efficiently manage processes as they arrive and meet their deadlines.
- The challenge is to design a data structure that supports all required operations in the best possible time.

Objective. Design an efficient data structure for managing processes.

Let's consider a set of processes $\mathcal{P} := \{p \mid \text{set of } n \text{ processes}\}$.

- **FINDMIN**(\mathcal{P}): Identify the process with the earliest deadline.
- **DELETEMIN**(\mathcal{P}): Remove the process with the earliest deadline.
- **DELETEPROCESS**(\mathcal{P}, p): Remove a specific process p .

EARLIEST DEADLINE FIRST

TRANSFORMING A DYNAMIC PROBLEM INTO A DATA STRUCTURE PROBLEM

Intuitive Goal

- We need to efficiently manage processes as they arrive and meet their deadlines.
- The challenge is to design a data structure that supports all required operations in the best possible time.

Objective. Design an efficient data structure for managing processes.

Let's consider a set of processes $\mathcal{P} := \{p \mid \text{set of } n \text{ processes}\}$.

- **FINDMIN**(\mathcal{P}): Identify the process with the earliest deadline.
- **DELETEMIN**(\mathcal{P}): Remove the process with the earliest deadline.
- **DELETEPROCESS**(\mathcal{P}, p): Remove a specific process p .
- **UPDATE**($\mathcal{P}, p, \text{NewDeadline}$): Update the deadline of process p to a new value.

EARLIEST DEADLINE FIRST

TRANSFORMING A DYNAMIC PROBLEM INTO A DATA STRUCTURE PROBLEM

Intuitive Goal

- We need to efficiently manage processes as they arrive and meet their deadlines.
- The challenge is to design a data structure that supports all required operations in the best possible time.

Objective. Design an efficient data structure for managing processes.


Let's consider a set of processes $\mathcal{P} := \{p \mid \text{set of } n \text{ processes}\}$.

- **FINDMIN**(\mathcal{P}): Identify the process with the earliest deadline.
- **DELETEMIN**(\mathcal{P}): Remove the process with the earliest deadline.
- **DELETEPROCESS**(\mathcal{P}, p): Remove a specific process p .
- **UPDATE**($\mathcal{P}, p, \text{NewDeadline}$): Update the deadline of process p to a new value.
- **INSERT**($\mathcal{P}, p_{\text{new}}, \text{Deadline}$): Insert a new process with a specified deadline.
- **BUILD**(\mathcal{P}): Construct the data structure with the set of processes \mathcal{P} .

EARLIEST DEADLINE FIRST

BAD SOLUTION #1

🤖 If we store processes in an unsorted array, what is the cost of $\text{FINDMIN}(\mathcal{P})$?



EARLIEST DEADLINE FIRST

BAD SOLUTION #1

🤖 If we store processes in an unsorted array, what is the cost of $\text{FINDMIN}(\mathcal{P})$?



Computing the minimum in an unsorted array costs $\Omega(n)$.

EARLIEST DEADLINE FIRST

BAD SOLUTION #2

🤖 If we store processes in a sorted array, what is the cost of $\text{FINDMIN}(\mathcal{P})$?



EARLIEST DEADLINE FIRST

BAD SOLUTION #2

🤖 If we store processes in a sorted array, what is the cost of $\text{FINDMIN}(\mathcal{P})$?



Computing the minimum in a sorted array costs $\Theta(1)$.

EARLIEST DEADLINE FIRST

BAD SOLUTION #2

🤖 If we store processes in a sorted array, what is the cost of $\text{INSERT}(\mathcal{P}, p_{\text{new}}, \text{Deadline})$?



EARLIEST DEADLINE FIRST

BAD SOLUTION #2


🤖 If we store processes in a sorted array, what is the cost of $\text{INSERT}(\mathcal{P}, p_{\text{new}}, \text{Deadline})$?



Inserting into and expanding a sorted array costs $\Omega(n)$.


Remember: Insertion Sort

EARLIEST DEADLINE FIRST

 Does anyone have an idea for a better solution?

EARLIEST DEADLINE FIRST


PRIORITY QUEUE AS A TREE

 Does anyone have an idea for a better solution?


Eureka: What if we use a tree structure with the minimum element at the root?

EARLIEST DEADLINE FIRST

PRIORITY QUEUE AS A TREE

 Does anyone have an idea for a better solution?

Eureka: What if we use a tree structure with the minimum element at the root?

 How should we order the remaining elements?

EARLIEST DEADLINE FIRST

PRIORITY QUEUE AS A TREE

🤖 Does anyone have an idea for a better solution?

Eureka: What if we use a tree structure with the minimum element at the root?

🤖 How should we order the remaining elements?

- 1) Divide in the middle to create a balanced tree.
- 2) Apply the same logic recursively to each partition.


BINARY BALANCED TREE

Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

🤖 Which level is not completely filled in this tree?



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

Which level is not completely filled in this tree?



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

🤖 Which rule did we violate when filling the last level of this tree?



BINARY BALANCED TREE


Definition

A binary tree is fully balanced if and only if all levels, except possibly the last, are completely filled from left to right.

Convention: The last level is filled with priority from left to right.

Imagine that in place of any missing nodes, we insert a dummy value, such as $+\infty$ or $-\infty$.

Which rule did we violate when filling the last level of this tree?



A NOTE FOR CODERS

- Constructing trees in programming often requires manipulating pointers and structures.

```
typedef struct treenode {           /* tree node struct */
    int data;                       /* integer field */
    struct treenode *left;          /* pointer to the left child */
    struct treenode *right;         /* pointer to the right child */
} node;
```

A NOTE FOR CODERS

- Constructing trees in programming often requires manipulating pointers and structures.

```
typedef struct treenode {           /* tree node struct */
    int data;                       /* integer field */
    struct treenode *left;         /* pointer to the left child */
    struct treenode *right;        /* pointer to the right child */
} node;
```


- For balanced trees, we can represent them simply using an array:
 - The children of node i are at positions $2i + 1$ and $2i + 2$.
 - The parent of node i is at position $(i - 1) \div 2$.

Node: i

Left child: $2i + 1$

Right child: $2i + 2$

Parent: $(i - 1) \div 2$

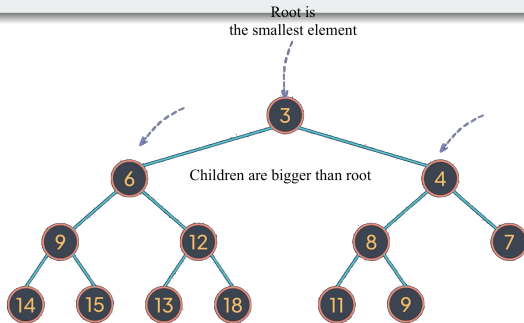


BINARY HEAP

A PICTURE IS WORTH A THOUSAND WORDS

Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.




BINARY HEAP

A PICTURE IS WORTH A THOUSAND WORDS

Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.




BINARY HEAP

A PICTURE IS WORTH A THOUSAND WORDS

Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.




BINARY HEAP

A PICTURE IS WORTH A THOUSAND WORDS

Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.



🤖 How do we store such a tree efficiently?

BINARY HEAP


A PICTURE IS WORTH A THOUSAND WORDS

Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.

🤖 How do we store such a tree efficiently?

Since it is a balanced binary tree, we can use an array representation.



BINARY HEAP


A PICTURE IS WORTH A THOUSAND WORDS

Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.

🤖 But inserting into an array always costs at least $\Omega(n)$, right?

Since it is a balanced binary tree, we can use an array representation.



BINARY HEAP

A PICTURE IS WORTH A THOUSAND WORDS


Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.

We can always insert an element at the end of array at $\Theta(1)$.

Remember REALLOC in C, C++, etc.

Since it is a balanced binary tree, we can use an array representation.



BINARY HEAP


A PICTURE IS WORTH A THOUSAND WORDS

Definition of a Binary Heap

A binary heap is a balanced binary tree where the root of every (sub)tree is its minimum element.

🤖 How do we maintain the tree as a binary heap?

Since it is a balanced binary tree, we can use an array representation.




SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




🤖 What operation do we perform after adding a new element to maintain the heap property?

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




We perform a "shift up" operation to place the new element and correct the subtree [10, 5, null].

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




🤖 Should we continue? · Do we need to check the entire heap?

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




Yes! · All the subtrees that do not include the new element are already correct.

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




🤖 If we change an element in the min-heap (w.l.o.g., increase), what do we have to do?

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."



4 8 7 9 12 6 10 14 15 13 18 11 9


Shift down until the heap property is restored.

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."



4 8 7 9 12 6 10 14 15 13 18 11 9


Shift down until the heap property is restored.

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."



4	8	6	9	12	7	10	14	15	13	18	11	9
---	---	---	---	----	---	----	----	----	----	----	----	---


Shift down until the heap property is restored.

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."




- 🧐 * What is the complexity of any shift up & shift down?
- * Can the update process create a cycle?
- * What if we change an element in the middle?

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."



4 8 6 9 12 7 10 14 15 13 18 11 9

Shift-Up	Shift-Down
$O(\log n)$	$O(\log n)$


* In the worst case, we move from the root to the leaf.

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."



4 8 6 9 12 7 10 14 15 13 18 11 9

Rule of Thumb:


{ If priority $\uparrow \Rightarrow$ Shift down
 { If priority $\downarrow \Rightarrow$ Shift up

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."



The subtree structure remains valid except for the affected branch after an update.

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."

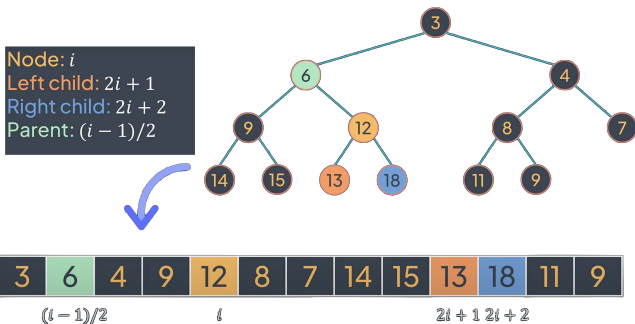
🤖 How do we know the parent & children indices are in $O(1)$?

SHIFT UP & SHIFT DOWN

IT'S ALL ABOUT SWAPS

Rule of Thumb

Whether you **INSERT** or **UPDATE** in the heap, always perform swaps to maintain:
The heap rule: "In a binary heap, the root of every (sub)tree is its minimum element."



DELETE FROM HEAP

SWAP WITH LAST ELEMENT

🤖 Deleting an element from an array typically costs at least $\Omega(n)$, right?

DELETE FROM HEAP

SWAP WITH LAST ELEMENT

🤖 Deleting an element from an array typically costs at least $\Omega(n)$, right?

We can delete the last element of the array in $\Theta(1)$ time.

*Remember **REALLOC** in C, C++, etc.*

DELETE FROM HEAP

SWAP WITH LAST ELEMENT

🤖 Deleting an element from an array typically costs at least $\Omega(n)$, right?

We can delete the last element of the array in $\Theta(1)$ time.

*Remember **REALLOC** in C, C++, etc.*

Deletion Strategy


- 1 Swap the element to be deleted with the last element in the heap.
- 2 Remove the last element from the array (reallocate).
- 3 Perform a **shift down** to restore the heap property.

DELETE FROM HEAP

EXAMPLE

Short Strategy

- 1 Swap the element to be deleted with the last element in the heap.
- 2 Remove the last element from the array (reallocate).
- 3 Perform a **shift down** to restore the heap property.




DELETE FROM HEAP

EXAMPLE

Short Strategy

- 1 Swap the element to be deleted with the last element in the heap.
- 2 Remove the last element from the array (reallocate).
- 3 Perform a **shift down** to restore the heap property.




DELETE FROM HEAP

EXAMPLE

Short Strategy

- 1 Swap the element to be deleted with the last element in the heap.
- 2 Remove the last element from the array (reallocate).
- 3 Perform a **shift down** to restore the heap property.




DELETE FROM HEAP

EXAMPLE

Short Strategy

- 1 Swap the element to be deleted with the last element in the heap.
- 2 Remove the last element from the array (reallocate).
- 3 Perform a **shift down** to restore the heap property.




DELETE FROM HEAP

EXAMPLE

Short Strategy

- 1 Swap the element to be deleted with the last element in the heap.
- 2 Remove the last element from the array (reallocate).
- 3 Perform a **shift down** to restore the heap property.



BUILDING A HEAP

BUILD(\mathcal{P}):Trivial Way

- Assume you have a set of processes $\mathcal{P} = \{p_1, \dots, p_n\}$ with priorities/deadlines (v_1, \dots, v_n) .

BUILDING A HEAP

BUILD(\mathcal{P}):Trivial Way

- Assume you have a set of processes $\mathcal{P} = \{p_1, \dots, p_n\}$ with priorities/deadlines (v_1, \dots, v_n) .
- Start with an empty heap $H = \emptyset$ and insert each element one by one using **INSERT**(H, p_i, v_i).

Runtime: $n \times \mathcal{O}(\log n) = \mathcal{O}(n \log n)$

BUILDING A HEAP


CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

Instead of inserting elements one by one, we can take the entire array and turn it into a heap in a more efficient way.

This process is called **heapify**.

- Start by treating the array as a binary tree.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.




BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.



BUILDING A HEAP

CLEVER WAY

BUILD(\mathcal{P}): Heapify Method

- Begin from the last non-leaf node, and perform a "shift down" operation to ensure each subtree satisfies the heap property.
- Continue this process moving upward to the root.

Worst-case Scenario

Every correction takes a shift-down $\mathcal{O}(\log n) \Rightarrow$
Runtime: $\mathcal{O}(n \log n)$.

But...

Optimistic Observation

Expensive corrections are rare & cheap corrections are frequent*

** Remember: 50% of nodes are leaves and 25% of nodes are parents of leaves.*

TOTAL SWAPS

PROOF SKETCH

Observation: Every level has different height. Shift-down costs actually $\mathcal{O}(\text{height of subtree})$.

$$\begin{aligned}
 & \frac{n}{2} : (\text{the leaves}) \times 0 \\
 & \quad + \\
 & \frac{n}{4} : (\text{parents of leaves}) \times 1 \\
 & \quad + \\
 & \frac{n}{8} : (\text{grandparents of leaves}) \times 2 = \mathcal{O}(n) \\
 & \quad + \\
 & \quad \dots + \\
 & 1 : \text{root} \times \mathcal{O}(\log n)
 \end{aligned}$$

HEAPSORT

HOW TO SORT AN ARRAY USING A HEAP

🤖 How we can sort an array using a min-heap?

HEAPSORT

HOW TO SORT AN ARRAY USING A HEAP

🤖 How we can sort an array using a min-heap?

Heapsort Explained

- First, build a max heap from the array using **heapify**.
- Then, repeatedly remove the root (the minimum element) and place it at the end of the array.
- After each removal, perform a "shift down" to restore the heap property.

$$\text{Runtime: } \underbrace{\text{heapify}}_{\mathcal{O}(n)} + n \times \underbrace{\text{delete}}_{\mathcal{O}(\log n)} + n \times \underbrace{\text{swaps}}_{\mathcal{O}(1)}$$

FIND K-TH ELEMENT

HOW TO SORT AN ARRAY USING A HEAP

How we can compute the k -th smallest element of an array using a min-heap?

FIND K-TH ELEMENT

HOW TO SORT AN ARRAY USING A HEAP

How we can compute the k -th smallest element of an array using a min-heap?

Two solutions (I will give you the complexity-describe the algorithm & application)

- 1 $\Theta(n) + O(k \log n)$
- 2 $\Theta(k) + O(n \log k)$

FIND K-TH ELEMENT

HOW TO SORT AN ARRAY USING A HEAP

How we can compute the k -th smallest element of an array using a min-heap?

Two solutions (I will give you the complexity-describe the algorithm & application)

- 1 $\Theta(n) + O(k \log n)$
- 2 $\Theta(k) + O(n \log k)$

Application: First version is offline. The second is online.

APPENDIX


REFERENCES

IMAGE SOURCES I



WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON

<https://brand.wisc.edu/web/logos/>



An injective non-surjective function (injective, not a bijection)

An injective surjective function (bijection)

A non-injective surjective function (surjective, not a bijection)

A non-injective non-surjective function (neither, not a bijection)

<https://en.wikipedia.org/wiki/Bijection>