

Sorting

SFWRENG 2CO3: Data Structures and Algorithms

Jelle Hellings

Department of Computing and Software
McMaster University



Winter 2024

A final sort algorithm: HEAPSORT

MERGESORT Worst-case $\Theta(N \log_2(N))$ runtime complexity.
But also: $\Theta(N)$ memory usage, high constants.

QUICKSORT Expected $\Theta(N \log_2(N))$ runtime complexity.
But also: $\Theta(\log_2(N))$ memory usage, finicky pivot choices,
worst-case $\Theta(N^2)$.

A final sort algorithm: HEAPSORT

MERGESORT Worst-case $\Theta(N \log_2(N))$ runtime complexity.
But also: $\Theta(N)$ memory usage, high constants.

QUICKSORT Expected $\Theta(N \log_2(N))$ runtime complexity.
But also: $\Theta(\log_2(N))$ memory usage, finicky pivot choices,
worst-case $\Theta(N^2)$.

Next: HEAPSORT

Worst-case $\Theta(N \log_2(N))$ runtime complexity and $\Theta(1)$ memory usage!

HEAPSORT: High-level overview

Algorithm SELECTIONLIKESORT(L):

Input: List $L[0 \dots N)$ of N values.

- 2: **for** $pos := N$ **to** 2 **do**
- 3: Find the position p of the
 maximum value in $L[0 \dots pos)$.
- 4: Exchange $L[pos - 1]$ and $L[p]$.

} Comparisons: $\sum_{pos=2}^N pos = \Theta(N^2)$.

HEAPSORT: High-level overview

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N)$ of N values.

- 1: Restructure L so that it is easy to find the *maximum*.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: Use structure to find the *maximum* and
 efficiently remove the *maximum*.
- 4: Place the *maximum* at $L[pos - 1]$.

HEAPSORT: High-level overview

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N]$ of N values.

- 1: Restructure L so that it is easy to find the *maximum* \leftarrow a *binary max-heap*.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: Use structure to find the *maximum* and
 efficiently remove the *maximum*.
- 4: Place the *maximum* at $L[pos - 1]$.

Max-heaps

A max-heap H is a collection of values

$ADD(H, v)$ add value v to a max-heap H ;

$DELMAX(H)$ removes the maximum value $w \in H$ and return w .

$SIZE(H)$ returns the number of values in H .

HEAPSORT: High-level overview

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N)$ of N values.

- 1: Restructure L so that it is easy to find the *maximum* \leftarrow a *binary max-heap*.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: Use structure to find the *maximum* and
 efficiently remove the *maximum*.
- 4: Place the *maximum* at $L[pos - 1]$.

Max-heaps

A max-heap H is a collection of values

$ADD(H, v)$ add value v to a max-heap H ; \leftarrow in $\Theta(\log_2(|H|))$.

$DELMAX(H)$ removes the maximum value $w \in H$ and return w . \leftarrow in $\Theta(\log_2(|H|))$.

$SIZE(H)$ returns the number of values in H .

HEAPSORT: High-level overview

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N]$ of N values.

- 1: Restructure L so that it is easy to find the *maximum* \leftarrow a *binary max-heap*.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: Use structure to find the *maximum* and
 efficiently remove the *maximum*.
- 4: Place the *maximum* at $L[pos - 1]$.

Max-heaps

A max-heap H is a collection of values

$ADD(H, v)$ add value v to a max-heap H ; \leftarrow in $\Theta(\log_2(|H|))$.

$DELMAX(H)$ removes the maximum value $w \in H$ and return w . \leftarrow in $\Theta(\log_2(|H|))$.

$SIZE(H)$ returns the number of values in H .

We can store a max-heap of $|H|$ values in an array of $|H|$ values.

HEAPSORT: High-level overview

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N]$ of N values.

- 1: Restructure L so that it is easy to find the *maximum*. $\approx N$ ADDS $\rightarrow \Theta(N \log_2(N))$.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: Use structure to find the *maximum* and
 efficiently remove the *maximum*. $\approx N$ DELMAXS $\rightarrow \Theta(N \log_2(N))$.
- 4: Place the *maximum* at $L[pos - 1]$.

Max-heaps

A max-heap H is a collection of values

$ADD(H, v)$ add value v to a max-heap H ; \leftarrow in $\Theta(\log_2(|H|))$.

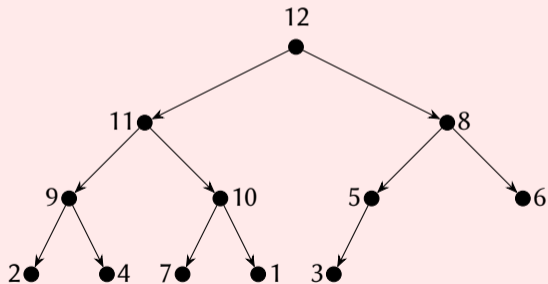
$DELMAX(H)$ removes the maximum value $w \in H$ and return w . \leftarrow in $\Theta(\log_2(|H|))$.

$SIZE(H)$ returns the number of values in H .

We can store a max-heap of $|H|$ values in an array of $|H|$ values.

The binary max-heap data structure

A binary max-heap is a *binary tree* in which each node n has a key $k(n)$ such that:

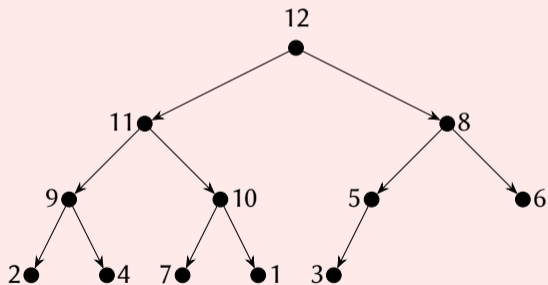


The binary max-heap data structure

A binary max-heap is a *binary tree* in which each node n has a key $k(n)$ such that:

- ▶ the tree is *nearly complete*.

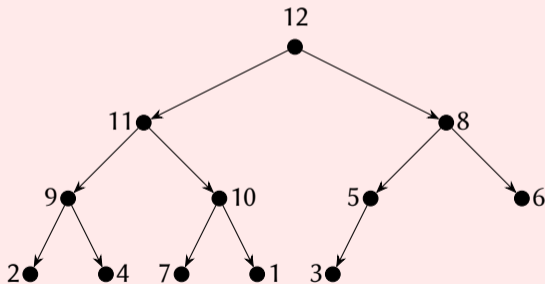
Or: the tree is filled from top-to-bottom, left-to-right.



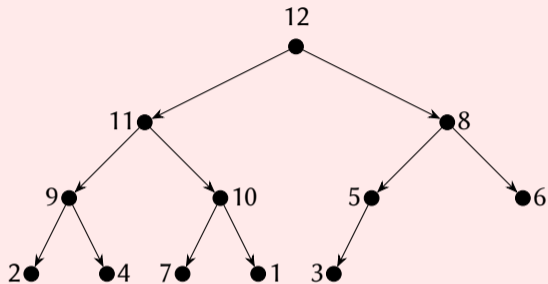
The binary max-heap data structure

A binary max-heap is a *binary tree* in which each node n has a key $k(n)$ such that:

- ▶ the tree is *nearly complete*.
Or: the tree is filled from top-to-bottom, left-to-right.
- ▶ the tree satisfies the *heap property*: if node n has child c , then $k(n) \geq k(c)$.
Or: the key in each node is larger-or-equal to the keys in the children of n .

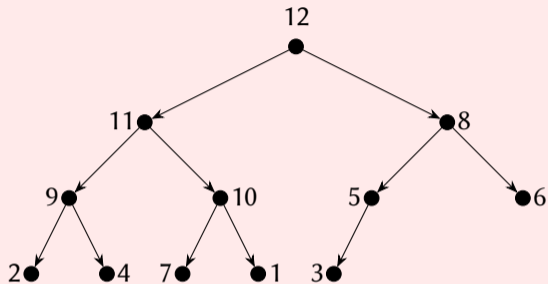


The binary max-heap data structure



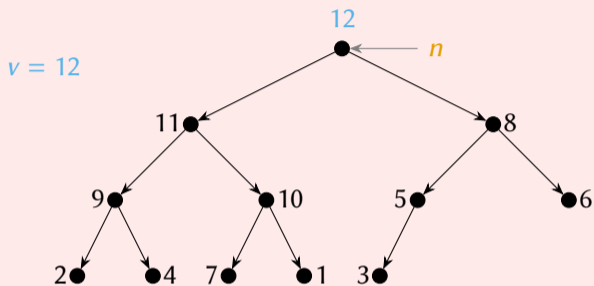
The *maximum* is straightforward to find: root of the tree.

The binary max-heap data structure



Algorithm DELMAX(H) (high-level overview):

The binary max-heap data structure

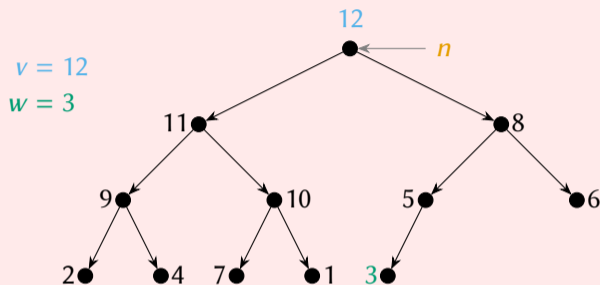


Algorithm $\text{DELMAX}(H)$ (high-level overview):

1: Let v be the value at the root n of H .

5: **return** v .

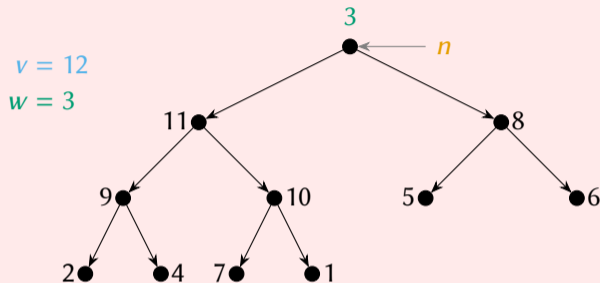
The binary max-heap data structure



Algorithm DELMAX(H) (high-level overview):

- 1: Let v be the value at the root n of H .
- 2: Let w be the last value in H .
- 3: Swap v and w .
- 4: Repeat steps 1 and 2 until v is a leaf.
- 5: **return** v .

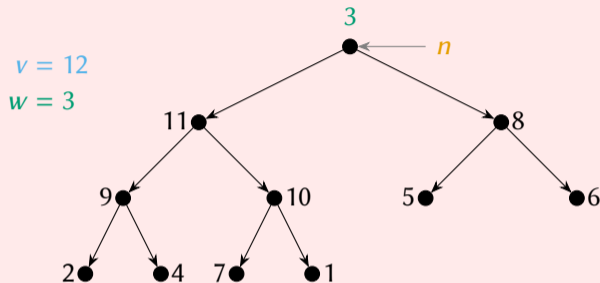
The binary max-heap data structure



Algorithm DELMAX(H) (high-level overview):

- 1: Let v be the value at the root n of H .
- 2: Let w be the last value in H .
- 3: Remove the last node in H and set $k(n) := w$.
- 5: **return** v .

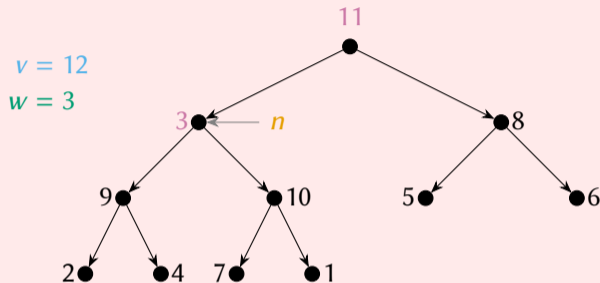
The binary max-heap data structure



Algorithm DELMAX(H) (high-level overview):

- 1: Let v be the value at the root n of H .
- 2: Let w be the last value in H .
- 3: Remove the last node in H and set $k(n) := w$.
- 4: Sink n to a valid position (reestablish the heap property).
- 5: **return** v .

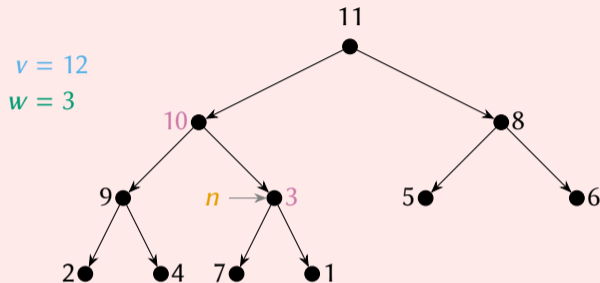
The binary max-heap data structure



Algorithm DELMAX(H) (high-level overview):

- 1: Let v be the value at the root n of H .
- 2: Let w be the last value in H .
- 3: Remove the last node in H and set $k(n) := w$.
- 4: Sink n to a valid position (reestablish the heap property).
- 5: **return** v .

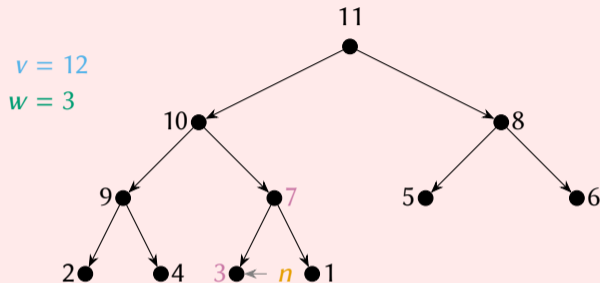
The binary max-heap data structure



Algorithm DELMAX(H) (high-level overview):

- 1: Let v be the value at the root n of H .
- 2: Let w be the last value in H .
- 3: Remove the last node in H and set $k(n) := w$.
- 4: Sink n to a valid position (reestablish the heap property).
- 5: **return** v .

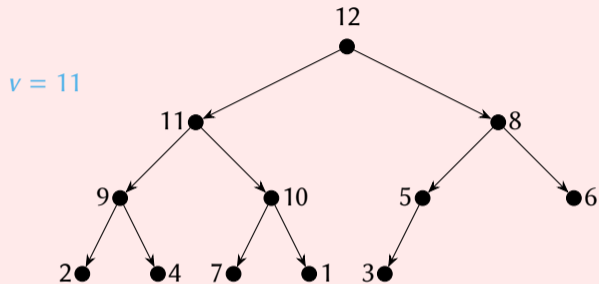
The binary max-heap data structure



Algorithm DELMAX(H) (high-level overview):

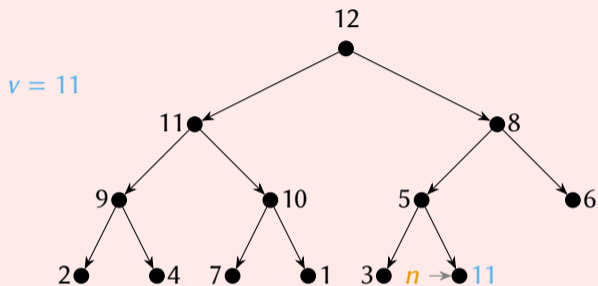
- 1: Let v be the value at the root n of H .
- 2: Let w be the last value in H .
- 3: Remove the last node in H and set $k(n) := w$.
- 4: Sink n to a valid position (reestablish the heap property).
- 5: **return** v .

The binary max-heap data structure



Algorithm $\text{ADD}(H, v)$ (high-level overview):

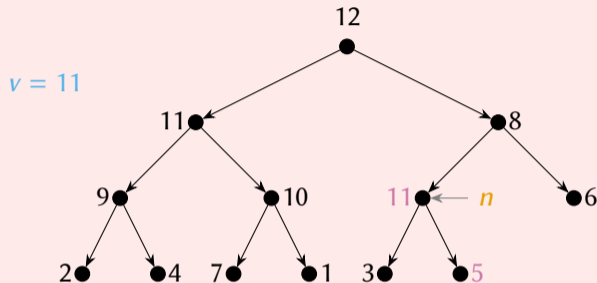
The binary max-heap data structure



Algorithm $\text{ADD}(H, v)$ (high-level overview):

- 1: Add a node n to the end of H with $k(n) := v$.

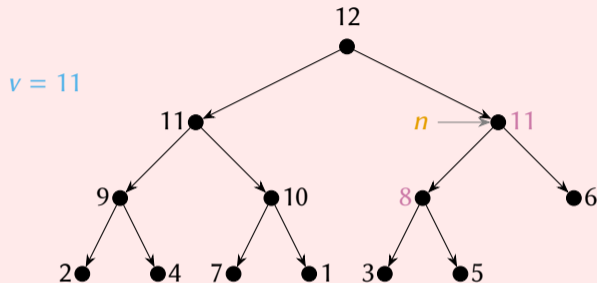
The binary max-heap data structure



Algorithm $\text{ADD}(H, v)$ (high-level overview):

- 1: Add a node n to the end of H with $k(n) := v$.
- 2: Swim n upward to a valid position (reestablish the heap property).

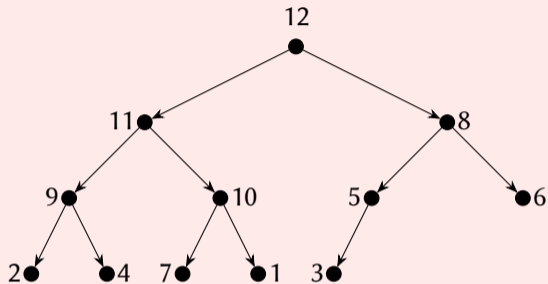
The binary max-heap data structure



Algorithm $\text{ADD}(H, v)$ (high-level overview):

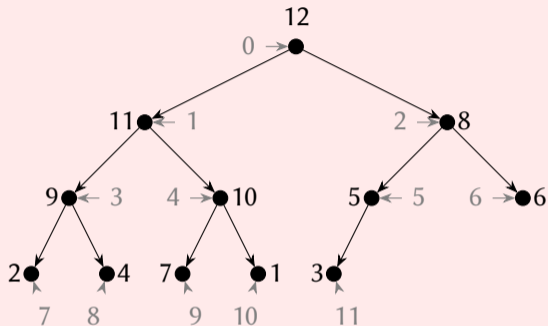
- 1: Add a node n to the end of H with $k(n) := v$.
- 2: Swim n upward to a valid position (reestablish the heap property).

The binary max-heap data structure



Storing a max-heap in an array

The binary max-heap data structure

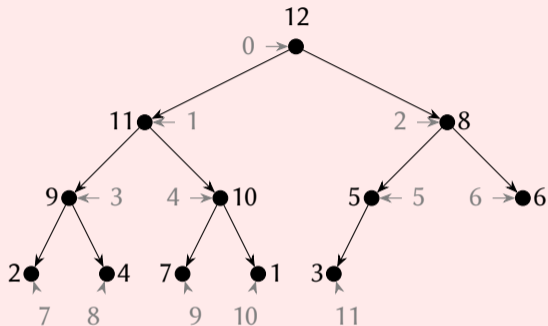


Storing a max-heap in an array

Number the nodes from top-to-bottom, left-to-right.

Warning: The book numbers values in max-heaps starting at 1 instead of 0!

The binary max-heap data structure

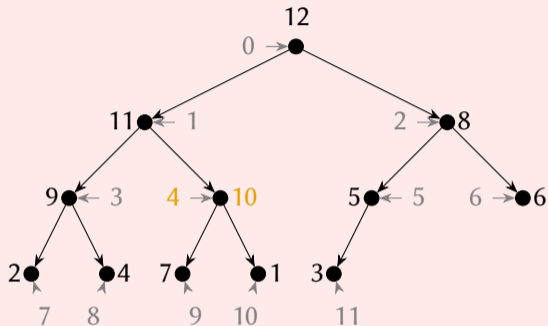


12	11	8	9	10	5	6	2	4	7	1	3
----	----	---	---	----	---	---	---	---	---	---	---

Storing a max-heap in an array

Number the nodes from top-to-bottom, left-to-right → positions in the array.

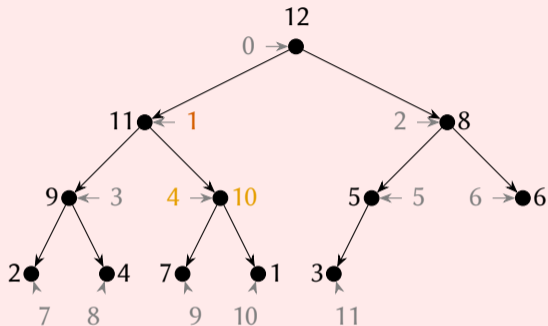
The binary max-heap data structure



12	11	8	9	10	5	6	2	4	7	1	3
----	----	---	---	----	---	---	---	---	---	---	---

If a node is at position p ,

The binary max-heap data structure

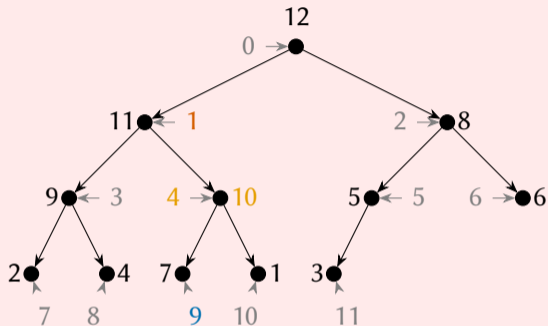


12	11	8	9	10	5	6	2	4	7	1	3
----	----	---	---	----	---	---	---	---	---	---	---

If a node is at position p ,

- ▶ then the **parent** is at position $\text{parent}(p) = (p - 1) \text{div } 2$.

The binary max-heap data structure

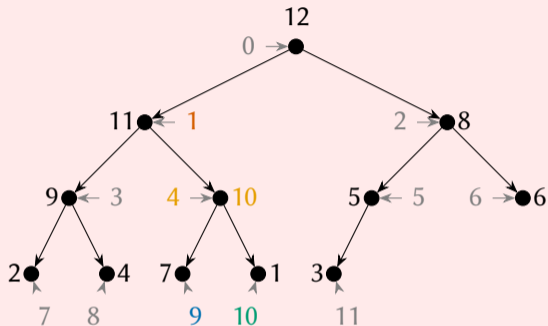


12	11	8	9	10	5	6	2	4	7	1	3
----	----	---	---	----	---	---	---	---	---	---	---

If a node is at position p ,

- ▶ then the **parent** is at position $\text{parent}(p) = (p - 1) \text{ div } 2$.
- ▶ then the **left child** is at position $\text{lchild}(p) = 2 \cdot p + 1$.

The binary max-heap data structure



12	11	8	9	10	5	6	2	4	7	1	3
----	----	---	---	----	---	---	---	---	---	---	---

If a node is at position p ,

- ▶ then the **parent** is at position $\text{parent}(p) = (p - 1) \text{ div } 2$.
- ▶ then the **left child** is at position $\text{lchild}(p) = 2 \cdot p + 1$.
- ▶ then the **right child** is at position $\text{rchild}(p) = 2 \cdot p + 2$.

The binary max-heap data structure

If a node is at position p ,

- ▶ then the **parent** is at position $\text{parent}(p) = (p - 1) \text{ div } 2$.
- ▶ then the **left child** is at position $\text{lchild}(p) = 2 \cdot p + 1$.
- ▶ then the **right child** is at position $\text{rchild}(p) = 2 \cdot p + 2$.

Algorithm SINK($L[0 \dots N], p$):

Sink $L[p]$ to a valid position.

Algorithm SWIM($L[0 \dots N], p$):

Swim $L[p]$ upward to a valid position.

The binary max-heap data structure

If a node is at position p ,

- ▶ then the **parent** is at position $\text{parent}(p) = (p - 1) \text{ div } 2$.
- ▶ then the **left child** is at position $\text{lchild}(p) = 2 \cdot p + 1$.
- ▶ then the **right child** is at position $\text{rchild}(p) = 2 \cdot p + 2$.

Algorithm SINK($L[0 \dots N], p$):

- 1: **while** true **do**
- 2: $np := p$.
- 3: **if** $\text{lchild}(p) < N$ **and**
 $L[np] < L[\text{lchild}(p)]$ **then**
- 4: $np := \text{lchild}(p)$.
- 5: **if** $\text{rchild}(p) < N$ **and**
 $L[np] < L[\text{rchild}(p)]$ **then**
- 6: $np := \text{rchild}(p)$.
- 7: **if** $np = p$ **then return**.
- 8: Exchange $L[p]$ and $L[np]$.
- 9: $p := np$.

Algorithm SWIM($L[0 \dots N], p$):

- 1: **while** $p \neq 0$ **and**
 $L[p] > L[\text{parent}(p)]$ **do**
- 2: Exchange $L[p]$ and $L[\text{parent}(p)]$.
- 3: $p := \text{parent}(p)$.

HEAPSORT: Filling in the details

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N]$ of N values.

- 1: Turn L into a max-heap.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: $max := \text{DELMAX}(L[0 \dots pos])$
- 4: $L[pos - 1] := max$.

HEAPSORT: Filling in the details

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N]$ of N values.

- 1: Turn L into a max-heap.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: $max := \text{DELMAX}(L[0 \dots pos])$
- 4: $L[pos - 1] := max$.

Algorithm MAKEHEAP($L[0 \dots N]$):

- 1: $len := 1$.
 /* inv: $L[0 \dots len]$ is a max-heap, bf: $N - len$ */
- 2: **while** $len \neq N$ **do**
- 3: $\text{ADD}(L[0 \dots len), L[len])$.
- 4: $len := len + 1$.

HEAPSORT: Filling in the details

Algorithm HEAPSORT(L):

Input: List $L[0 \dots N]$ of N values.

- 1: Turn L into a max-heap.
- 2: **for** $pos := N$ **to** 2 **do**
- 3: $max := \text{DELMAX}(L[0 \dots pos])$
- 4: $L[pos - 1] := max$.

Algorithm MAKEHEAP($L[0 \dots N]$):

- 1: $len := 1$.
 /* inv: $L[0 \dots len]$ is a max-heap, bf: $N - len$ */
 - 2: **while** $len \neq N$ **do**
 - 3: $\text{ADD}(L[0 \dots len), L[len])$.
 - 4: $len := len + 1$.
- } $N - 1$ SWIM operations.

HEAPSORT: Filling in the details

A faster MAKEHEAP

L :

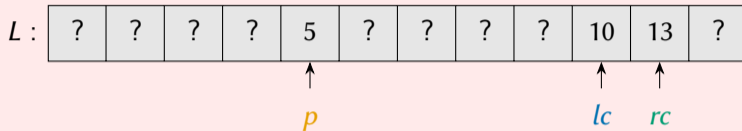
?	?	?	?	5	?	?	?	?	10	13	?
---	---	---	---	---	---	---	---	---	----	----	---

Consider a position p in L such that

- ▶ the **left child** at position $\text{lchild}(p)$ already forms a valid max-heap in L ; and
- ▶ the **right child** at position $\text{rchild}(p)$ already forms a valid max-heap in L .

HEAPSORT: Filling in the details

A faster MAKEHEAP

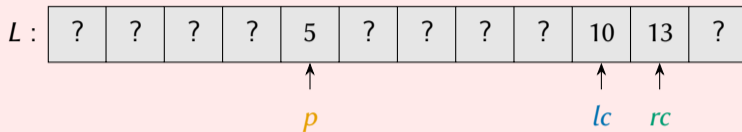


Consider a position p in L such that

- ▶ the **left child** at position $lchild(p)$ already forms a valid max-heap in L ; and
- ▶ the **right child** at position $rchild(p)$ already forms a valid max-heap in L .

HEAPSORT: Filling in the details

A faster MAKEHEAP



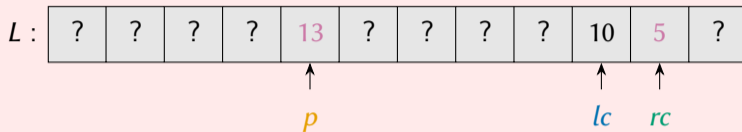
Consider a position p in L such that

- ▶ the **left child** at position $lchild(p)$ already forms a valid max-heap in L ; and
- ▶ the **right child** at position $rchild(p)$ already forms a valid max-heap in L .

SINK($L[0..N]$, p) assures that p also forms a valid max-heap in L .

HEAPSORT: Filling in the details

A faster MAKEHEAP



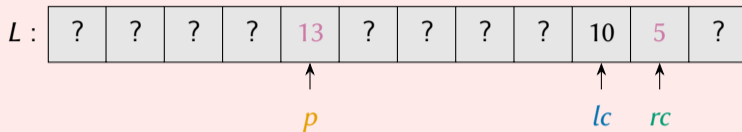
Consider a position p in L such that

- ▶ the **left child** at position $lchild(p)$ already forms a valid max-heap in L ; and
- ▶ the **right child** at position $rchild(p)$ already forms a valid max-heap in L .

SINK($L[0..N]$, p) assures that p also forms a valid max-heap in L .

HEAPSORT: Filling in the details

A faster MAKEHEAP



Consider a position p in L such that

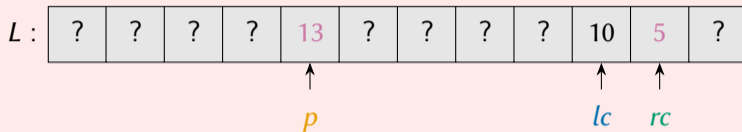
- ▶ the **left child** at position $lchild(p)$ already forms a valid max-heap in L ; and
- ▶ the **right child** at position $rchild(p)$ already forms a valid max-heap in L .

$SINK(L[0..N], p)$ assures that p also forms a valid max-heap in L .

Leaves *always* form valid max-heaps.

HEAPSORT: Filling in the details

A faster MAKEHEAP



Consider a position p in L such that

- ▶ the **left child** at position $lchild(p)$ already forms a valid max-heap in L ; and
- ▶ the **right child** at position $rchild(p)$ already forms a valid max-heap in L .

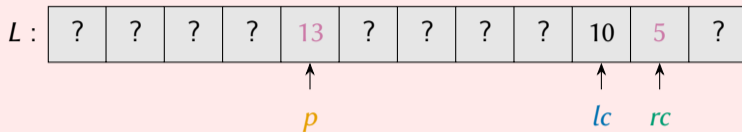
$SINK(L[0..N], p)$ assures that p also forms a valid max-heap in L .

Leaves *always* form valid max-heaps.

Positions $m \leq N$ with $lchild(m) = m \cdot 2 + 1 \geq N$ represent the leaves in a max-heap of L .

HEAPSORT: Filling in the details

A faster MAKEHEAP



Consider a position p in L such that

- ▶ the **left child** at position $lchild(p)$ already forms a valid max-heap in L ; and
- ▶ the **right child** at position $rchild(p)$ already forms a valid max-heap in L .

SINK($L[0..N]$, p) assures that p also forms a valid max-heap in L .

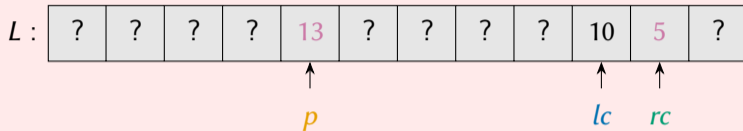
Leaves *always* form valid max-heaps.

Positions $m \leq N$ with $lchild(m) = m \cdot 2 + 1 \geq N$ represent the leaves in a max-heap of L

→ The last non-child in a max-heap of L values is at position $\lfloor \frac{N-1}{2} \rfloor$.

HEAPSORT: Filling in the details

A faster MAKEHEAP

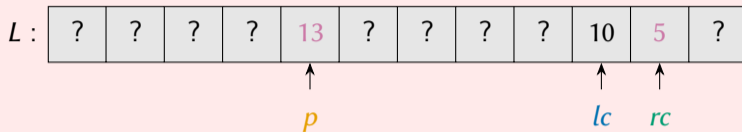


Algorithm FASTMAKEHEAP(L):

- 1: $k := N - (N \text{ div } 2)$.
- 2: **while** $k \neq 0$ **do**
- 3: $\text{SINK}(L[0 \dots N], k)$.
- 4: $k := k - 1$.

HEAPSORT: Filling in the details

A faster MAKEHEAP



Algorithm FASTMAKEHEAP(L):

- 1: $k := N - (N \text{ div } 2)$.
- 2: **while** $k \neq 0$ **do**
- 3: SINK($L[0 \dots N], k$).
- 4: $k := k - 1$.

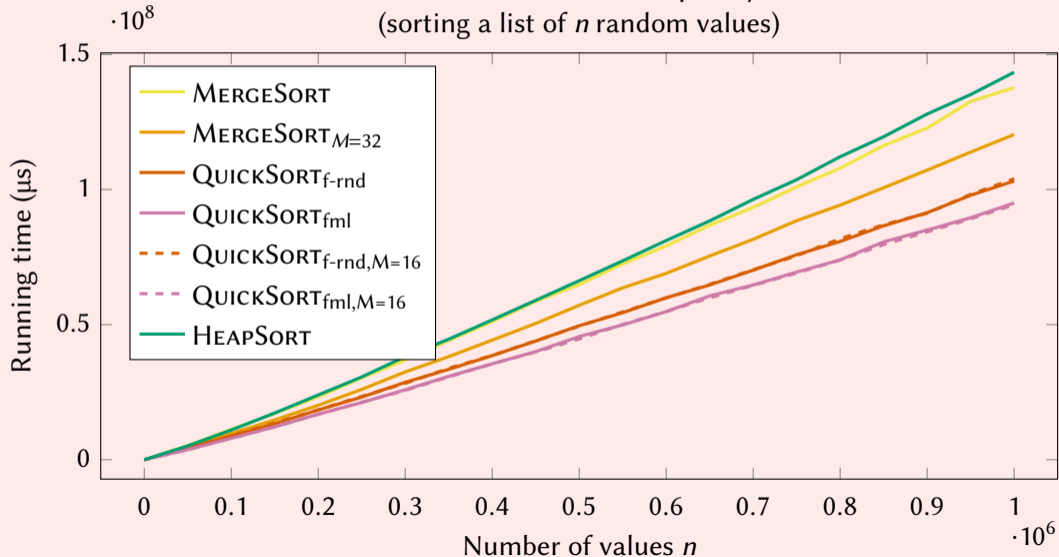
} $\lfloor \frac{N-1}{2} \rfloor$ SINK operations.

Comparing HEAPSORT with MERGESORT and QUICKSORT

	Comparisons	Changes	Memory	Stable
MERGESORT	$\Theta(N \log_2(N))$	$N \log_2(N)$	$\Theta(N)$	yes
QUICKSORT	$\Theta(N \log_2(N))$ (expected)	$\Theta(N \log_2(N))$ (expected)	$\Theta(\log_2(N))$ (expected)	no
HEAPSORT	$\Theta(N \log_2(N))$	$N \log_2(N)$	1	no

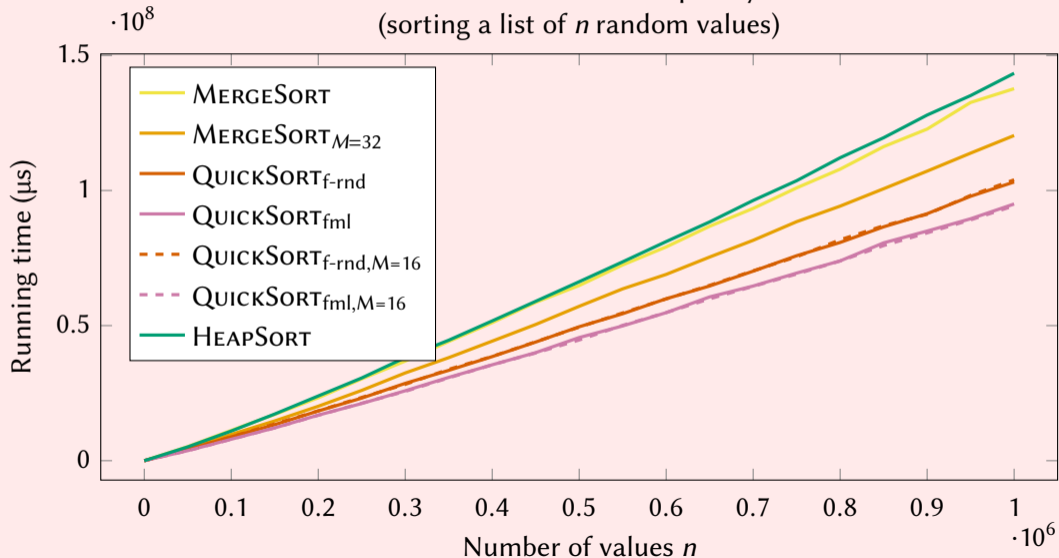
Comparing HEAPSORT with MERGESORT and QUICKSORT

Measured runtime complexity
(sorting a list of n random values)



Comparing HEAPSORT with MERGESORT and QUICKSORT

Measured runtime complexity
(sorting a list of n random values)



Final notes on HEAPSORT

A *max-heap* is often referred to as a *Priority Queue*.

There are also *min-heaps* that provide fast access to *minimum values*.

Final notes on HEAPSORT

A *max-heap* is often referred to as a *Priority Queue*.

There are also *min-heaps* that provide fast access to *minimum values*.

	C++	Java
Priority Queues	<code>std::priority_queue</code>	<code>java.util.PriorityQueue</code>
ADD	<code>std::push_heap</code>	
DELMAX	<code>std::pop_heap</code>	
(related)	<code>std::make_heap</code>	
	<code>std::is_heap</code>	
	<code>std::is_heap_until</code>	
	<code>std::sort_heap</code>	

INTROSORT: Putting all sorts together

Algorithm INTROSORT($L[start \dots end]$, $potential$):

$potential$ is the number of values we *could have* sorted with perfect pivot choices.

- 1: **if** $end - start \leq M$ **then**
- 2: Sort $L[start \dots end]$ using INSERTIONSORT.
- 3: **else if** $potential < 1.5 \cdot |L|$ **then**
- 4: Sort $L[start \dots end]$ using HEAPSORT.
- 5: **else**
- 6: Choose the position $p \in [start, end)$ of the *pivot value* $v := L[pos]$.
- 7: $pos := \text{PARTITION}(L, start, end, p)$.
- 8: INTROSORT($L[start \dots pos]$, $2 \cdot potential$).
- 9: INTROSORT($L[pos + 1 \dots end]$, $2 \cdot potential$).

Algorithm INTROSORT($L[0 \dots N]$):

- 10: INTROSORT($L[0 \dots N]$, 1).

INTROSORT: Putting all sorts together

Measured runtime complexity
(sorting a list of n random values)

