

Heaps
○○

Binary heaps
○○○○○○○○○○○○○○○○○○○○

d-heaps
○

Binomial heaps
○○○○○○○○○○○○○○○○○○

Fibonacci heaps
○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○

Heaps

Combinatorial Optimization

Giovanni Righini

Heaps

Let indicate the maximum number of elements in the data-structure with n .

- **List or array:** *Insert* in $O(1)$, *ExtractMin* in $O(n)$.
- **Sorted list:** *Insert* in $O(n)$, *ExtractMin* in $O(1)$.
- **Sorted array:** *Insert* in $O(n)^{(*)}$, *ExtractMin* in $O(1)^{(**)}$.
- **Balanced binary tree:** *Insert* in $O(\log n)$, *ExtractMin* in $O(\log n)$.

(*) Finding the correct insertion position by dichotomic search takes $O(\log n)$ time; shifting the last elements one position to the right to create the empty space in which the new element is to be inserted takes $O(n)$ time.

(**) The elements must be sorted in non-increasing order, so that the minimum is the last one.

Binary heaps

A binary heap is a binary tree in which a *key* is associated with each node and the following two properties hold:

- all levels of the tree are filled from the left to the right and only the last level is allowed to be incomplete;
- the key associated with a parent node is not larger than the keys associated with its two children nodes.

In a binary heap the operations *Insert*, *ExtractMin* and *DecreaseKey* can be implemented to run in $O(\log n)$, being n the number of nodes in the tree.

Insert(i, v)

An element i and its key value v are in input:

1. a new element is placed in the last position of the heap: $O(1)$
2. it is pushed upwards by swapping it with its parent, until the property of the heap is restored: $O(\log n)$

$$L \leftarrow L + 1$$
$$\text{heap}[L].\text{node} \leftarrow i$$
$$\text{heap}[L].\text{cost} \leftarrow v$$
$$\text{pos}[i] \leftarrow L$$
$$\text{MoveUp}(L)$$

ExtractMin(i, v)

The root node i and its key value v are in output:

1. the root of the heap is extracted and replaced by the last element: $O(1)$;
2. the new root is moved downwards by swapping it with its minimum cost child, until the property of the heap is restored: $O(\log n)$.

```
i ← heap[1].node  
v ← heap[1].cost  
pos[i] ← nil  
heap[1] ← heap[L]  
L ← L - 1  
MoveDn(1)
```

MoveUp(*p*)

```
stop ← false
while (p ≠ 1) ∧ (¬stop) do
  q ← parent(p)
  if heap[q].cost > heap[p].cost then
    Swap(p, q)
  else
    stop ← true
  end if
end while
```

MoveDn(*p*)

```
stop ← false
while (left(p) ≤ L) ∧ (¬stop) do
  if (right(p) > L) ∨ (heap[left(p)].cost < heap[right(p)].cost) then
    q ← left(p)
  else
    q ← right(p)
  end if
  if heap[p].cost > heap[q].cost then
    Swap(p, q)
  else
    stop ← true
  end if
end while
```

Swap(*p*, *q*)

Both *MoveUp* and *MoveDn* use *Swap* as a sub-routine. Its complexity is $O(1)$.

It uses a temporary record *r* and a temporary integer *k*.

```
r ← heap[p]  
heap[p] ← heap[q]  
heap[q] ← r  
k ← pos[p]  
pos[p] ← pos[q]  
pos[q] ← k  
k ← p  
p ← q  
q ← k
```

Building a heap

When n values to be partially sorted in a heap are known since the beginning, it is not needed to build the heap with n successive *Insert* operations, which would take $O(n \log n)$ time.

```
// An array  $v$  of elements is in input //
for  $i = 1, \dots, n$  do
     $heap[i].node \leftarrow i$ 
     $heap[i].value \leftarrow v[i]$ 
     $pos[i] \leftarrow i$ 
end for
for  $k = \lfloor n/2 \rfloor, \dots, 1$  do
     $MoveDn(k)$ 
end for
```

Complexity

Assume the root is at level 1 and the leaves at level h .

For a node at level k , the maximum number of swap operations when it is checked as a root is $h - k$.

At each level k there are at most 2^{k-1} nodes.

Hence, in the worst case the total number of swap operations is

$$S = \sum_{k=1}^{h-1} (h - k)2^{k-1}.$$

Complexity

Hence

$$S = \sum_{k=1}^{h-1} (h-k)2^{k-1} = h \sum_{k=1}^{h-1} 2^{k-1} - \sum_{k=1}^{h-1} k2^{k-1}.$$

The second term is further split as follows:

$$\sum_{k=1}^{h-1} k2^{k-1} = \sum_{k=0}^{h-2} (k+1)2^k = \sum_{k=0}^{h-2} k2^k + \sum_{k=0}^{h-2} 2^k.$$

Therefore $S = A - B - C$, with

- $A = h \sum_{k=1}^{h-1} 2^{k-1}$;
- $B = \sum_{k=0}^{h-2} k2^k$;
- $C = \sum_{k=0}^{h-2} 2^k$.

Complexity

$$A = h \sum_{k=1}^{h-1} 2^{k-1} = h \sum_{k=0}^{h-2} 2^k = h(2^{h-1} - 1).$$

$$C = \sum_{k=0}^{h-2} 2^k = 2^{h-1} - 1.$$

Complexity

$$B = \sum_{k=0}^{h-2} k2^k = 1 \cdot 2^1 + 2 \cdot 2^2 + 3 \cdot 2^3 + \dots + (h-2)2^{h-2}.$$

$$2B = \sum_{k=0}^{h-2} k2^{k+1} = 1 \cdot 2^2 + 2 \cdot 2^3 + 3 \cdot 2^4 + \dots + (h-2)2^{h-1}.$$

$$\begin{aligned} B &= 2B - B = -2^1 - 2^2 - 2^3 - \dots - 2^{h-2} + (h-2)2^{h-1} = \\ &= (h-2)2^{h-1} - \sum_{k=1}^{h-2} 2^k = (h-2)2^{h-1} - \left(\sum_{k=0}^{h-2} 2^k - 2^0 \right) = \\ &= (h-2)2^{h-1} - (2^{h-1} - 1 - 1) = (h-3)2^{h-1} + 2. \end{aligned}$$

Complexity

Finally

$$\begin{aligned} S &= A - B - C = h(2^{h-1} - 1) - [(h - 3)2^{h-1} + 2] - [2^{h-1} - 1] = \\ &= 2 \cdot 2^{h-1} - h - 2 + 1 = 2^h - h - 1. \end{aligned}$$

Since h grows as $O(\log n)$, then S grows as $O(n)$.

The complexity of *BuildHeap* when n elements are available is $O(n)$.

Find k^{th} smallest element

By-product of the linear complexity of *BuildHeap*:
find the k^{th} smallest element in a set of n elements.

Algorithm 1 (with an array):

- Sort the elements: $O(n \log n)$;
- Access the element in position k : $O(1)$.

Algorithm 2 (with a binary heap):

- *BuildHeap*: $O(n)$;
- For k times *ExtractMin*: $O(k \log n)$.

If k is “small” (i.e. $O(\frac{n}{\log n})$), algorithm 2 is likely to be faster.

Insert vs DecreaseKey

When the key value of an element of the heap decreases:

- *DecreaseKey* of the existing element in the heap;
- *Insert* of a new copy of the element with the new key value, leaving the old copy unchanged.

Both procedures take $O(\log n)$ in the worst case, but....

- *DecreaseKey* keeps the size of the heap constant, but it requires *pos* to access the element;
- *Insert* makes the heap growing, but it does not require *pos* (and the time to update it in *Swap*).

There is a trade-off between the additional time spent because the heap grows and the additional time spent to update *pos*.

d-heaps

A *d*-heap is a tree in which each element may have up to *d* children, with $d > 2$.

Number of levels: $\lceil \log_d n + 1 \rceil$ instead of $\lceil \log_2 n + 1 \rceil$.

Pro: fewer *Swap* operations are required.

Con: in *MoveDn* the selection of the minimum cost child costs $O(d)$ instead of $O(1)$.

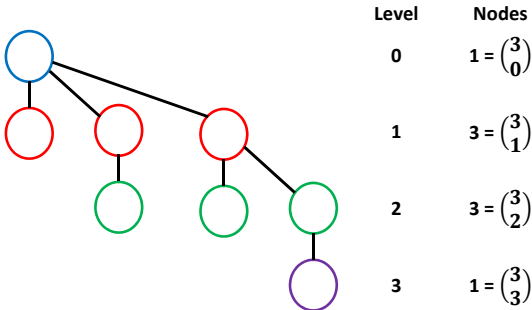
Binomial heaps

A binomial heap is a forest of heap-ordered trees of different size.

Property 1.

A tree of height k has exactly 2^k elements (height = n. edges between the root and the last leaf).

The number of elements at each level $l = 1, \dots, k$ is $\binom{k}{l}$ (binomial coefficient).



Binomial heaps

Property 2.

Elements in all trees are partially sorted as in a binary heap: the key of the predecessor is always less than or equal to the keys of its successors.

Property 3.

For each number n of nodes in the heap, there exists a unique corresponding set of heights of the trees:
one-to-one correspondence with the binary encoding of n .

Therefore the number of trees in the heap is $\lceil \log_2 n + 1 \rceil$.

Data structure

Elements in a binomial heap do not have a predefined number of successors.

A possible implementation uses three pointers for each record:

- a pointer to the **predecessor**;
- a pointer to the **first successor**;
- a pointer to the **next successor** of the predecessor (“sibling”).

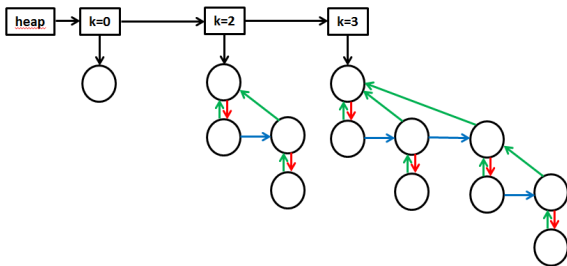


Figure: A sample binomial heap with 13 elements (binary encoding = 1101).

ExtractMin

There are $O(\log n)$ trees in the heap: comparing the keys of their roots and selecting the best one takes $O(\log n)$.

The deletion of the root of a tree T of height k produces:

- a tree T_1 of height $k - 2$;
- a tree T_2 of height $k - 1$.

Then, trees of the same height are merged as binary numbers are added up.

ExtractMin

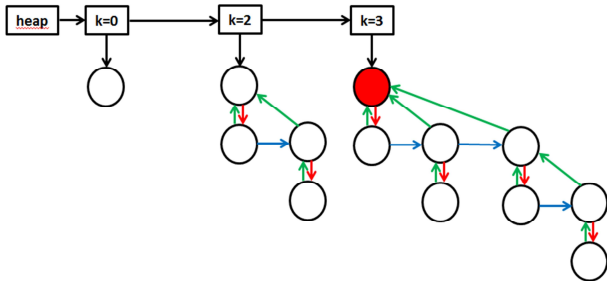


Figure: In this example we assume that the red element has the minimum key and it is extracted. Its tree has height $k^* = 3$.

Fibonacci heaps

With Fibonacci heaps we can do:

- *Insert* in $O(1)$;
- *DecreaseKey* in $O(1)$;
- *ExtractMin* in $O(\log n)$.

The idea is “lazy update”: the data structure is updated (spending time for this) only under special conditions.

The $O(1)$ complexity above is achieved as *amortized complexity*, i.e. considering sequences of operations, not single operations.

Amortized complexity is more realistic because not all operations in a sequence require the worst case time.

Amortized complexity

The **amortized complexity** of an operation is the **average** worst-case time complexity that is achieved when a sequence of k operations are executed for k large enough.

The amortized complexity of an operation is $O(f(n))$ if it exists a positive integer k such that for a sequence of at least k operations the total time required by the sequence is $O(kf(n))$.

The most common technique to establish the amortized complexity of an operation uses a *potential function*.

A potential function is increased by some operations and decreased by others, so that the maximum number of times an operation can occur can be bounded by the number of times other operations occur.

Structure and implementation

A Fibonacci heap is made by a set of trees.
Their roots are linked in a doubly linked list.

An array *Pred* stores the predecessor of each element.
Successors of the same element are linked in a doubly-linked list,
pointed by the predecessor.

Each element has a rank, indicating the number of its successors.

A pointer π^* indicates the root with the minimum key value.
An array *Lost* counts how many successors a non-root node has lost.
An array *Bucket* contains pointers to roots of rank $1, 2, \dots, \lfloor \log_\phi n \rfloor$.

Potential function

Consider a potential function τ defined as the number of trees in the heap.

Every *Link* operation decreases τ by 1 unit.

Every *Cut* operation increases τ by 1 unit.

Therefore, the number of executions of *Link* is bounded above by the number of executions of *Cut* plus the starting value of τ , which is at most n .

Lower bounding the size of trees

Property 4.

Every (sub-)tree whose root has rank k contains at least $F(k + 2)$ elements.

The proof is based on invariant Property 2 and invariant Property 3 and on the Fibonacci numbers sequence.

Proof.

Consider node w with rank k .

Sort its k successors according to the order in which they have been appended to w , from the earliest to the latest.

Consider node y , the i^{th} successor of w , with $1 \leq i \leq k$.

Bounding the rank

Property 5. The rank of all nodes in a Fibonacci heap is bounded by $\lfloor \log_{\phi} n \rfloor$.

Lemma. $F(k + 2) \geq \phi^k$, where $\phi = \frac{1 + \sqrt{5}}{2}$.
(The proof is omitted here. It will be given later.)

Since the size of any tree is bounded by n , we have

$$n \geq \text{size}(T) \geq F(k + 2) \geq \phi^k$$

where k is rank of the root of tree T . Therefore,

$$k \leq \log_{\phi} n.$$

Bounding the number of trees

Property 6. A Fibonacci heap contains at most $\lfloor \log_\phi n \rfloor$ trees.

Proof. This property follows from:

- Property 3: No two roots have the same rank.
- Property 5: No rank can be larger than $\lfloor \log_\phi n \rfloor$.

Primary and secondary *Cuts*

To keep invariant property 2 valid, we must ensure that every non-root node does not lose more than one successor after becoming a non-root node.

$Lost(i)$ indicates how many successors node i has already lost.

For the complexity analysis we use a potential function

$$\mu = \sum_{i=1}^n Lost(i).$$

Initially $\mu = 0$ and μ can never become negative.

Primary and secondary *Cuts*

When a $Cut(i)$ operation is done, two cases may occur. Note that $Pred(i)$ certainly exists, because i is not a root.

- $Lost(Pred(i)) = 0$: then $Lost(Pred(i))$ is set to 1 and the procedure stops.
- $Lost(Pred(i)) = 1$: then $Lost(Pred(i))$ is set to 2 but $Cut(Pred(i))$ is executed, to make $Pred(i)$ the root of a new tree. In turn $Cut(Pred(i))$ can trigger the execution of another Cut and so on.

We say that $Cut(i)$ is a **primary *Cut*** operation, while the Cut operations triggered by it are **secondary *Cut*** operations.

A cascade of secondary Cut terminates

- when a non-root j is reached with $Lost(j) = 0$;
- when the root of the tree is reached.

Complexity analysis

Consider a primary $Cut(i)$ operation.

- $Lost(i)$ is set to 0: it decreases by 1 unit or remains unchanged.
- $Lost(Pred(i))$ is increased by 1.

Therefore a primary cut does not increase the potential μ by more than 1 unit.

Consider a secondary $Cut(i)$ operation.

- $Lost(i)$ is updated from 2 to 0.
- $Lost(Pred(i))$ is increased by 1.

Therefore a secondary cut decreases the potential μ by 1.

Therefore, **the number of secondary cuts cannot exceed the number of primary cuts.**

FindMin

This simply uses the pointer π^* to the root with minimum value of the key.

No pointer is modified.

Complexity: $O(1)$.

Insert(i)

A new tree with a single element i is created and it is inserted in the list of roots in $O(1)$.

The pointer to the best root is possibly updated in $O(1)$.

A sequence of *Link* operations can be triggered to restore invariant property 3.

Complexity: $O(1)$ plus the contribution due to the cascade of *Link* operations.

Proving the lemma for even n

Proof for n even.

We use the following equations:

$$1 + \frac{1}{\phi} = \phi$$

$$F(n+2) = F(n+1) + F(n)$$

$$F(n+2) > \phi^n \quad \forall n \text{ odd.}$$

Assume $F(n) \geq \phi^{n-2}$ for an even n .

$$\begin{aligned} F(n+2) &= F(n+1) + F(n) \geq \phi^{n-1} + \phi^{n-2} = \\ &= \phi^{n-1} \left(1 + \frac{1}{\phi}\right) = \phi^{n-1} \phi = \phi^n \end{aligned}$$

Inductive step: $F(n) \geq \phi^{n-2}$ implies $F(n+2) \geq \phi^n$.

Induction basis: $F(4) = 3 > \phi^2$.

Therefore $F(n+2) > \phi^n$ also for all n even.