

CS 261: Data Structures

Week 5: Priority queues

Lecture 5a: Binary and k -ary heaps; Fibonacci heaps

David Eppstein

University of California, Irvine

Spring Quarter, 2025



This work is licensed under a Creative Commons Attribution 4.0 International License

Priority queues

Example: Dijkstra's algorithm

To find the shortest path distance from a starting vertex s to each other vertex t in a directed graph with positive edge lengths:

1. Initialize dictionary D of distances; $D[s] = 0$ and, for all other vertices t , $D[t] = +\infty$
2. Initialize collection Q of not-yet-processed vertices (initially all vertices)
3. While Q is non-empty:
 - ▶ Find the vertex $v \in Q$ with the smallest value of $D[v]$
 - ▶ Remove v from Q
 - ▶ For each edge $v \rightarrow w$, set $D[w] = \min(D[w], D[v] + \text{length}(v \rightarrow w))$

Same algorithm can find paths themselves (not just distances) by storing, for each w , its predecessor on its shortest path: the vertex v that gave the minimum value in the calculation of $D[w]$

Priority queue operations used by Dijkstra, I

For Dijkstra to be efficient, we need to organize Q into a **priority queue** data structure allowing vertex with minimum D to be found quickly.

First operation used by Dijkstra: **Create a new priority queue**

1. Initialize dictionary D of distances; $D[s] = 0$ and, for all other vertices t , $D[t] = +\infty$
2. **Initialize collection Q of not-yet-processed vertices (initially all vertices)**
3. While Q is non-empty:
 - ▶ Find the vertex $v \in Q$ with the smallest value of $D[v]$
 - ▶ Remove v from Q
 - ▶ For each edge $v \rightarrow w$, set $D[w] = \min(D[w], D[v] + \text{length}(v \rightarrow w))$

Priority queue operations used by Dijkstra, II

Second operation used by Dijkstra:

Find and remove item with minimum value

(Some applications use maximum value; some descriptions of this data structure separate find and remove into two operations.)

1. Initialize dictionary D of distances; $D[s] = 0$ and, for all other vertices t , $D[t] = +\infty$
2. Initialize collection Q of not-yet-processed vertices (initially all vertices)
3. While Q is non-empty:
 - ▶ Find the vertex $v \in Q$ with the smallest value of $D[v]$
 - ▶ Remove v from Q
 - ▶ For each edge $v \rightarrow w$, set $D[w] = \min(D[w], D[v] + \text{length}(v \rightarrow w))$

Priority queue operations used by Dijkstra, III

Third operation used by Dijkstra: **Change item's priority**

In Dijkstra the priority always gets smaller (earlier in the priority ordering); this will be important for efficiency.

1. Initialize dictionary D of distances; $D[s] = 0$ and, for all other vertices t , $D[t] = +\infty$
2. Initialize collection Q of not-yet-processed vertices (initially all vertices)
3. While Q is non-empty:
 - ▶ Find the vertex $v \in Q$ with the smallest value of $D[v]$
 - ▶ Remove v from Q
 - ▶ For each edge $v \rightarrow w$, set $D[w] = \min(D[w], D[v] + \text{length}(v \rightarrow w))$

Summary of priority queue operations

Used by Dijkstra:

- ▶ Create new priority queue for elements with priorities
- ▶ Find and remove minimum-priority element
- ▶ Decrease the priority of an element

Also sometimes useful:

- ▶ Add an element
- ▶ Remove arbitrary element
- ▶ Increase priority
- ▶ Merge two queues

Priority queues in Python

Main built-in type: `heapq`

Implements binary heap (this lecture), represented as dynamic array

Operations:

- ▶ `heapify`: Create new priority queue for collection of elements (no separate priorities: ordered by standard comparisons)
- ▶ `heappop`: Find and remove minimum priority element
- ▶ `heappush`: Add an element

No ability to associate numeric priorities to arbitrary vertex objects. No ability to change priorities of elements already in queue.

Dijkstra in Python

Cannot use decrease-priority; instead, store bigger priority queue of triples (d, u, v) where $u \rightarrow v$ is an edge and d is the length of the shortest path to u plus the length of the edge

First triple (d, u, v) found for v is the one giving the shortest path

```
def dijkstra(G,s,length):  
    D = {}  
    Q = [(0,None,s)]  
    while Q:  
        dist,u,v = heapq.heappop(Q)  
        if v not in D:  
            D[v] = dist  
            for w in G[v]:  
                heapq.heappush(Q,(D[v]+length(v,w),v,w))  
    return D
```

Binary heaps

Binary heaps

Binary heaps are a type of priority queue

(Min-heap: smaller priorities are earlier; max-heap: opposite)

Based on **heap-ordered trees** (min is at root) and an efficient representation of binary trees by arrays

(No need for separate tree nodes and pointers)

Allows priorities to change

(But to do this you need a separate data structure to keep track of the location of each object in the array – this extra complication is why Python doesn't implement this operation)

Find-and-remove, add, change-priority all take $O(\log n)$ time

Creating a binary heap from n objects takes $O(n)$ time


This is what Python `heapq` uses

Representing trees using arrays

Number the nodes of infinite binary tree row-by-row, left-to-right

Use node number as index for array cell representing that node

For finite trees with n nodes, use the first n indexes



$$\text{parent}(x) = \lfloor (x - 1) / 2 \rfloor$$

$$\text{children}(x) = (2x + 1, 2x + 2)$$

Heap order


A tree is **heap-ordered** when, for every non-root node x ,

$$\text{priority}(\text{parent}(x)) \leq \text{priority}(x)$$


A **binary heap** on n items is an array H of the items, so that the corresponding binary tree is heap-ordered: For all $0 < i < n$,

$$\text{priority}(H[\lfloor \frac{i-1}{2} \rfloor]) \leq \text{priority}(H[i])$$


Examples of heap order



2	4	7	99
---	---	---	----




2	4	99	7
---	---	----	---



2	7	4	99
---	---	---	----

Heap-ordered



2	7	99	4
---	---	----	---

Not heap-ordered

Sorted arrays are heap-ordered (left example)
but heap-ordered arrays might not be sorted (middle two examples)

Find and remove minimum-priority element

Where to find it: always in array cell 0

How to remove it:

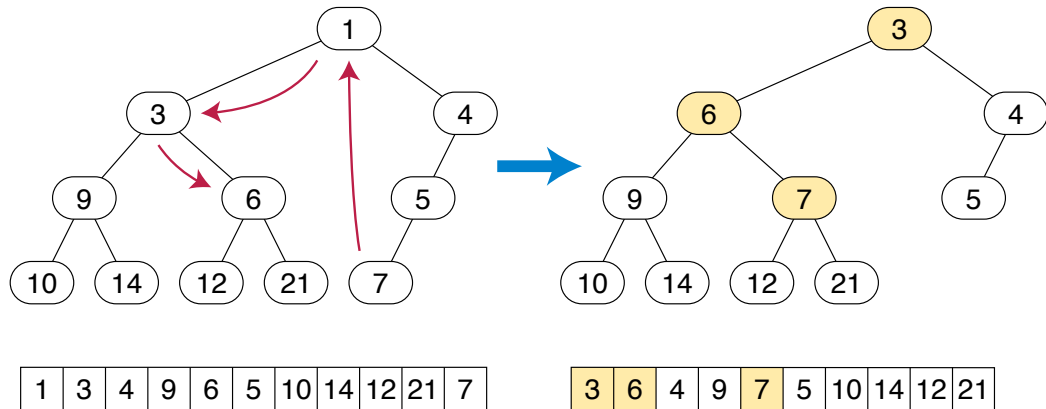
1. Swap element zero with element $n - 1$
2. Reduce array length by one
3. Fix heap at position 0 (might now be larger than children)

To fix heap at position p , repeat:

1. If p has no children ($2p + 1 \geq n$), return
2. Find child q with minimum priority
3. If $\text{priority}(p) \leq \text{priority}(q)$, return
4. Swap elements at positions p and q
5. Set $p = q$ and continue repeating

Time = $O(1)$ per level of the tree = $O(\log n)$ total

Example of delete-min operation



To delete the minimum element (element 1):

Swap the final element (7) into root position and shrink array

Swap element 7 with its smallest child (3) to fix heap order

Swap element 7 with its smallest child (6) to fix heap order

Both children of 7 are larger, stop swapping

Changing element priorities

First, find the position p of the element you want to change

(Simplest: keep a separate dictionary mapping elements to positions, update it whenever we change positions of elements)

To increase priority: fix heap at position p , same as for find-and-remove

To decrease priority, repeat:

1. If $p = 0$ or parent of p has smaller priority, return
2. Swap elements at positions p and parent
3. Set $p = \text{parent}$ and continue repeating

...swapping on a path upward from p instead of downward

For both increase and decrease, time is $O(\log n)$

Insert new element

To insert a new element x with priority p into a heap that already has n elements:

- ▶ Add x to array cell $H[n]$ with priority $+\infty$
- ▶ The heap property is satisfied!
(because of the big but fake priority)
- ▶ Change the priority of x to p
(swapping along path towards root)

Time: $O(\log n)$

Heapify

Given an unsorted array, put it into heap order

Main idea: Fix subtrees that have only one element out of order
(by swapping with children on downward path in tree)

Recursive version:

To recursively heapify the subtree rooted at i :

1. Recursively heapify left child of i
2. Recursively heapify right child of i
3. Fix heap at i

Non-recursive version:

For $i = n - 1, n - 2, \dots, 0$:
Fix heap at i

Heapify analysis (either version)

In a binary heap with n tree nodes:

- ▶ $\lceil n/2 \rceil$ are leaves, fix heap takes one step
- ▶ $\lceil n/4 \rceil$ are parents of leaves, fix heap takes two steps
- ...
- ▶ $\lceil n/2^i \rceil$ are i levels up, fix heap takes i steps

Total steps:

$$\sum_{i=1}^{\lceil \log_2(n+1) \rceil} i \left\lceil \frac{n}{2^i} \right\rceil \leq O(\log^2 n) + n \sum_{i=1}^{\infty} \frac{i}{2^i} = 2n + O(\log^2 n)$$

The $O(\log^2 n)$ term is how much the rounding-up part of the formula can contribute, compared to not rounding.)

So total time is $O(n)$, as stated earlier

k-ary heaps

Can we speed this up?

In a comparison model of computation, some priority queue operations must take $\Omega(\log n)$ time

- ▶ We can use a priority queue to sort n items by repeatedly finding and removing the smallest one (heapsort)
- ▶ But comparison-based sorting requires $\Omega(n \log n)$ time

But in Dijkstra, different operations have different frequencies

- ▶ In a graph with n vertices and m edges...
- ▶ Dijkstra does n find-and-remove-min operations,
- ▶ but up to m decrease-priority operations
- ▶ In some graphs, m can be much larger than n

Goal: speed up the more frequent operation (decrease-priority) without hurting the other operations too much

k -ary heap

Same as binary heap but with $k > 2$ children per node

$$\text{Parent}(x) = \lfloor (x - 1) / k \rfloor$$

$$\text{Children}(x) = (kx + 1, kx + 2, \dots, kx + k)$$

$$\text{Height of tree: } \log_k n = \frac{\log n}{\log k}$$

$$\text{Time to decrease priority} = O(\text{height}) = O\left(\frac{\log n}{\log k}\right)$$

(because path to root is shorter)

$$\text{Time to find-and-remove-min} = O\left(\frac{k \log n}{\log k}\right)$$

(because each step finds smallest of k children)

Dijkstra using k -ary heap

Time for m decrease-priority operations: $O\left(m \frac{\log n}{\log k}\right)$

Time for n find-and-remove-min operations: $O\left(nk \frac{\log n}{\log k}\right)$

To minimize total time, choose k to balance these two bounds

$$k = \max(2, \lceil m/n \rceil)$$

$$\text{Total time} = O\left(m \frac{\log n}{\log m/n}\right)$$

This becomes $O(m)$ whenever $m = \Omega(n^{1+\varepsilon})$ for any constant $\varepsilon > 0$

k -ary heaps in practice

4-ary heaps may be better than 2-ary heaps for all operations

For heaps that are too large to fit into cache memory, even larger choices of k may be better

(Lower tree height beats added complexity of more children, especially when each level costs a cache miss)

Fibonacci heaps

Overview

A different priority queue structure optimized for Dijkstra:

- ▶ Create n -item heap takes time $O(n)$ (same as binary heap)
- ▶ Find-and-remove-min is $O(\log n)$ (same as binary tree)
- ▶ Insert takes time $O(1)$
- ▶ Decrease-priority takes time $O(1)$

⇒ Dijkstra + Fibonacci heaps takes time $O(m + n \log n)$

Can merge heaps (destructively) in time $O(1)$; occasionally useful

All time bounds are amortized, not worst case

In practice, high constant factors make it worse than k -ary heaps

Structure of a Fibonacci heap

It's just a heap-ordered forest!

Each node = an object that stores:

- ▶ An element in the priority queue, and its priority
- ▶ Pointer to parent node
- ▶ Doubly-linked list of children
- ▶ Its degree (number of children)
- ▶ "Mark bit" (see next slide)

Heap property: all non-roots have $\text{priority}(\text{parent}) \leq \text{priority}(\text{self})$

Whole heap = object with doubly-linked list of tree roots
and a pointer to the root with the smallest priority

No requirements limiting the shape or number of trees
(But we will see that operations cause their shapes to be limited)

Mark bits and potential function

Each node has a binary “mark bit”

- ▶ True: This node is a non-root, and after becoming a non-root it had one of its children removed
- ▶ False: It's a root, or has not had any children removed

(We do not allow the removal of more than one child from non-root nodes)

Potential function Φ : $2 \times (\text{number of marked nodes}) + 1 \times (\text{number of trees})$

Some easy operations

Create new heap: Make each element its own tree root

Actual time: $O(n)$ $\Delta\Phi = +n$ Total: $O(n)$

Insert new element: Make it the root of its own tree

Actual time: $O(1)$ $\Delta\Phi = +1$ Total: $O(1)$

Merge two heaps: Concatenate their lists of roots

Actual time: $O(1)$ $\Delta\Phi = 0$ Total: $O(1)$

Find min-priority element: Follow pointer to minimum root

Actual time: $O(1)$ $\Delta\Phi = 0$ Total: $O(1)$

Decrease-priority operation, main idea


Cut edge from decreased node to its parent, making it a new root

If this removes 2nd child from any node, make it a new root too

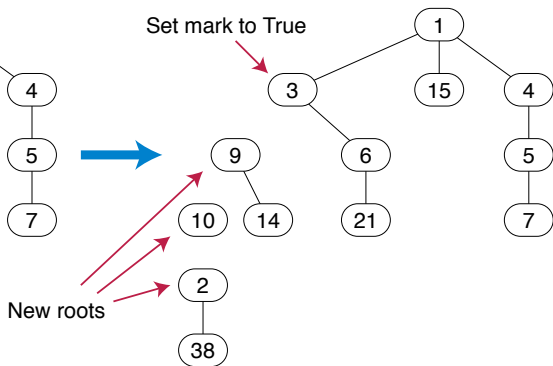
Mark = False,
ok to remove child

Mark = True, would
remove 2nd child

Decrease
priority



Set mark to True



Decrease-priority details

To decrease the priority of element at node x :

1. $y = \text{parent}(x)$
2. Cut link from x to y (making x a root)
3. Set $\text{mark}(x) = \text{False}$
4. While y is marked:
 - ▶ $z = \text{parent}(y)$
 - ▶ Cut link from y to z (making y a root)
 - ▶ Set $\text{mark}(y) = \text{False}$
 - ▶ Set $y = z$
5. If y is not a root, set $\text{mark}(y) = \text{True}$

If we cut k links, then the total actual time is $O(k)$. We unmark $k - 1$ nodes, create k new trees, and mark at most one node, so $\Delta\Phi \leq -2(k - 1) + k + 2 = 4 - k$. The $-k$ term in $\Delta\Phi$ cancels the $O(k)$ actual time leaving $O(1)$ amortized time.

Delete-min intuition

When we delete the minimum root,
we may add many new roots (its children)
we need to find the new minimum among the old and new roots

This can be slow, but we can pay for it by decreasing $\#trees$
and reducing Φ for every tree we eliminate

We need to merge the trees in a careful way
in order to control $\#children$ of deleted root

Delete-min operation

The only operation that creates bigger trees from smaller ones...

1. Make all children of deleted node into new roots
2. Merge trees with equal degree until no two are equal
3. Find the new minimum root

Actual time: $O(\# \text{ children} + \# \text{ merges} + \# \text{ remaining trees})$


$\Delta\Phi$: number of children - number of merges

Total amortized time:

$$O(\# \text{ children} + \# \text{ remaining trees}) = O(\max \# \text{ children})$$

Merging two trees

Make the larger root become the child of the smaller root!



We will only do this when both trees have equal degree
(number of children)

⇒ Degree goes up by one

Merging until all trees have different degrees

1. Make a dictionary D , where $D[x]$ will be a tree of degree x
 x will be a small integer, so D can just be an array
2. Make a list L of trees needing to be processed
Initially $L = \text{all trees}$
3. While L is non-empty:
 - ▶ Let T be any tree in L ; remove T from L
 - ▶ Let x be the degree of tree T
 - ▶ If key x is not already in D , set $D[x] = T$
 - ▶ Otherwise merge $D[x]$ and T , remove x from D , and add the merged tree to L
4. Return the collection of trees stored in D

Total times through loop = # merges + # additions to D
 $= 2 \times \# \text{ merges} + \# \text{ remaining trees}$

Degrees of children

Lemma: Order the children of any node by the time they became children (from a merge operation)


Then their degrees are $\geq 0, 0, 1, 2, 3, 4, 5, \dots$

Proof: When i th child was merged in, earlier children were already there \Rightarrow parent degree was $\geq i - 1$

At the time of the merge, it had same degree as parent, $\geq i - 1$

After taking away at most one child, its degree is $\geq i - 2$

The smallest trees obeying the degree lemma



Tree size is exponential in degree

Lemma: Subtree of degree x has $\#$ nodes $\geq \text{Fibonacci}(x)$

Proof: By induction

Base cases $x = 0$ (one node) and $x = 1$ (two nodes)

For $x > 1$:

number of nodes \geq root + numbers in each subtree

$$\geq 1 + \sum_{i=1}^x \max(1, F_{i-2})$$

$$\geq (\text{same sum for one fewer child}) + F_{x-2}$$

$$= F_{x-1} + F_{x-2} = F_x$$

Consequences of the degree lemma

Corollary 1: Maximum degree in an n -node Fibonacci heap is $\leq \log_{\varphi} n$ where $\varphi = (1 + \sqrt{5})/2$ is the golden ratio

Corollary 2: Amortized time for delete-min operation is $O(\log n)$

Ongoing research question

Can we achieve the same time bounds as Fibonacci
(constant amortized decrease-priority, logarithmic delete-min)
with a structure as practical as binary heaps?