Binomial heaps,
Fibonacci heaps,
and applications

1

Binomial trees

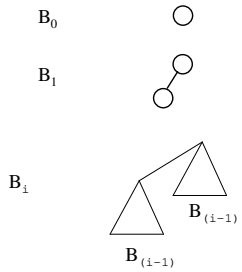

2

Binomial trees

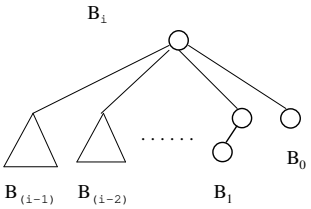

3

Properties of binomial trees

1) $|B_k| = 2^k$
2) degree(root(B_k)) = k
3) depth(B_k) = k

==> The degree and depth of a binomial tree with at most n nodes is at most log(n).

Define the rank of B_k to be k

4

Binomial heaps (def)

A collection of binomial trees at most one of every rank.
Items in the nodes, heap ordered.

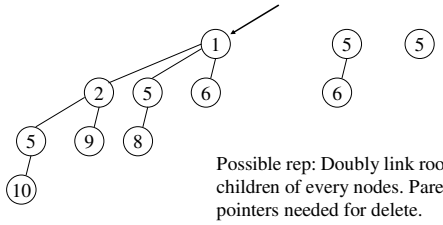

Possible rep: Doubly link roots and children of every nodes. Parent pointers needed for delete.

5

Binomial heaps (operations)

Operations are defined via a basic operation, called linking, of binomial trees:
Produce a B_k from two B_{k-1}, keep heap order.

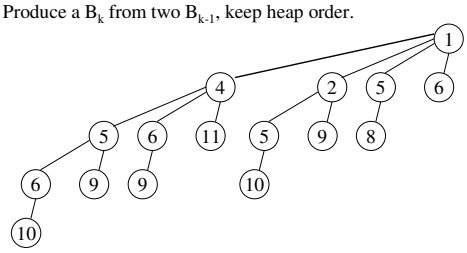

6

Binomial heaps (ops cont.)

Basic operation is meld(h1,h2):

Like addition of binary numbers.

	B_5	B_4		B_2	B_1	
h1:		B_4	B_3		B_1	B_0
h2:		B_4	B_3			B_0
	B_5	B_4		B_2		

+

Binomial heaps (ops cont.)

Findmin(h): obvious

Insert(x,h) : meld a new heap with a single B_0 containing x, with h

deletemin(h) : Chop off the minimal root. Meld the subtrees with h. Update minimum pointer if needed.

delete(x,h) : Bubble up and continue like delete-min

decrease-key(x,h,δ) : Bubble up, update min ptr if needed

All operations take O(log n) time on the worst case, except find-min(h) that takes O(1) time.

Binomial heaps - amortized ana.

Φ (collection of heaps) = #(trees)

Amortized cost of insert O(1)

Amortized cost of other operations still O(log n)

Binomial heaps + lazy meld

Allow more than one tree of each rank.

Meld (h1,h2) :

•Concatenate the lists of binomial trees.

•Update the minimum pointer to be the smaller of the minimums

O(1) worst case and amortized.

Binomial heaps + lazy meld

As long as we do not do a delete-min our heaps are just doubly linked lists:

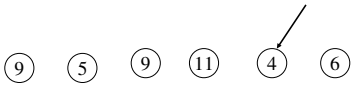

Delete-min : Chop off the minimum root, add its children to the list of trees.

Successive linking: Traverse the forest keep linking trees of the same rank, maintain a pointer to the minimum root.

Binomial heaps + lazy meld

Possible implementation of delete-min is using an array indexed by rank to keep at most one binomial tree of each rank that we already traversed.

Once we encounter a second tree of some rank we link them and keep linking until we do not have two trees of the same rank. We record the resulting tree in the array

Amortized(delete-min) = #trees - #links

$\leq$ (log(n) + #links) - #links = O(log(n))

Fibonacci heaps (Fredman & Tarjan 84)

Want to do decrease-key(x,h,δ) faster than delete+insert.

Ideally in O(1) time.

Why ?

36

Dijkstra's shortest path algorithm

Let G = (V,E) be a weighted (weights are non-negative) undirected graph, let $s \in V$. Want to find the distance (length of the shortest path), d(s,v) from s to every other vertex.

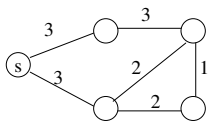

37

Dijkstra's shortest path algorithm

Dijkstra: Maintain an upper bound d(v) on d(s,v).

Every vertex is either scanned, labeled, or unlabeled.

Initially: $d(s) = 0$ and $d(v) = \infty$ for every $v \neq s$.
s is labeled and all others are unlabeled.

• Pick a labeled vertex with d(v) minimum. Make v scanned.

For every edge (v,w) if $d(v) + w(v,w) < d(w)$ then

1) $d(w) := d(v) + w(v,w)$

2) label w if it is not labeled already

38

Dijkstra's shortest path algorithm (implementation)

Maintain the labeled vertices in a heap, using d(v) as the key of v.

We perform n delete-min operations and n insert operations on the heap. O(n log(n))

For each edge we may perform a decrease-key.

With regular heaps O(m log (n)).

But if you can do decrease-key in O(1) time then you can implement Dijkstra's algorithm to run in O(n log(n) + m) time !

39

Back to Fibonacci heaps

Suggested implementation for decrease-key(x,h,δ):

If x with its new key is smaller than its parent, cut the subtree rooted at x and add it to the forest. Update the minimum pointer if necessary.

40

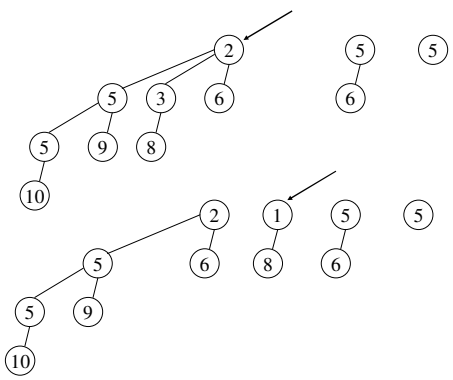

41

Decrease-key (cont.)

Does it work ?

Obs1: Trees need not be binomial trees any more..

Do we need the trees to be binomial ?
Where have we used it ?

In the analysis of delete-min we used the fact that at most log(n) new trees are added to the forest. This was obvious since trees were binomial and contained at most n nodes.

Decrease-key (cont.)

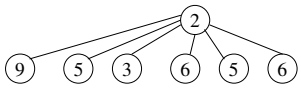

Such trees are now legitimate.

So our analysis breaks down.

Fibonacci heaps (cont.)

We shall allow non-binomial trees, but will keep the degrees logarithmic in the number of nodes.

Rank of a tree = degree of the root.

Delete-min: do successive linking of trees of the same rank and update the minimum pointer as before.

Insert and meld also work as before.

Fibonacci heaps (cont.)

Decrease-key (x,h,δ): indeed cuts the subtree rooted by x if necessary as we showed.

in addition we maintain a mark bit for every node. When we cut the subtree rooted by x we check the mark bit of p(x). If it is set then we cut p(x) too. We continue this way until either we reach an unmarked node in which case we mark it, or we reach the root.

This mechanism is called cascading cuts.

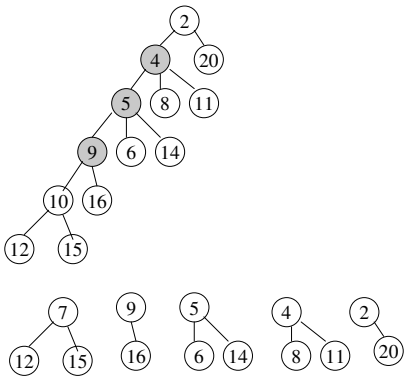

Fibonacci heaps (delete)

Delete(x,h) : Cut the subtree rooted at x and then proceed with cascading cuts as for decrease key.

Chop off x from being the root of its subtree and add the subtrees rooted by its children to the forest

If x is the minimum node do successive linking

Fibonacci heaps (analysis)

Want everything to be O(1) time except for delete and delete-min.

==> cascading cuts should pay for themselves

Φ (collection of heaps) = #(trees) + 2#(marked nodes)

Actual(decrease-key) = O(1) + #(cascading cuts)

$\Delta\Phi$(decrease-key) = O(1) - #(cascading cuts)

==> amortized(decrease-key) = O(1) !

Fibonacci heaps (analysis)

What about delete and delete-min ?

Cascading cuts and successive linking will pay for themselves. The only question is what is the maximum degree of a node ?
How many trees are being added into the forest when we chop off a root ?

Fibonacci heaps (analysis)

Lemma 1 : Let x be any node in an F-heap. Arrange the children of x in the order they were linked to x, from earliest to latest. Then the i-th child of x has rank at least i-2.

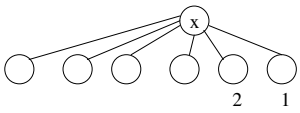

Proof:

When the i-th node was linked it must have had at least i-1 children.
Since then it could have lost at most one. ∎

Fibonacci heaps (analysis)

Corollary1 : A node x of rank k in a F-heap has at least ϕ^k descendants, where $\phi = (1 + \sqrt{5})/2$ is the golden ratio.

Proof:

Let s_k be the minimum number of descendants of a node of rank k in a F-heap.

By Lemma 1 $s_k \geq \Sigma_{i=0}^{k-2} s_i + 2$

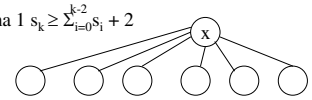

$s_0=1$, $s_1= 2$

Fibonacci heaps (analysis)

Proof (cont):

Fibonnaci numbers satisfy

$F_{k+2} = \Sigma_{i=2}^{k} F_i + 2$, for $k \geq 2$, and $F_2=1$

so by induction $s_k \geq F_{k+2}$

It is well known that $F_{k+2} \geq \phi^k$

It follows that the maximum degree k in a F-heap with n nodes is such that $\phi^k \leq n$

so $k \leq \log(n) / \log(\phi) = 1.4404 \log(n)$

Application #2 : Prim's algorithm for MST

Start with T a singleton vertex.

Grow a tree by repeating the following step:

Add the minimum cost edge connecting a vertex in T to a vertex out of T.

Application #2 : Prim's algorithm for MST

Maintain the vertices out of T but adjacent to T in a heap.

The key of a vertex v is the weight of the lightest edge (v,w) where w is in the tree.

Iteration: Do a delete-min. Let v be the minimum vertex and (v,w) the lightest edge as above. Add (v,w) to T.

For each edge (w,u) where $u \notin T$,

if key(u) = ∞ insert u into the heap with key(u) = w(w,u)
if w(w,u) < key(u) decrease the key of u to be w(w,u).

With regular heaps O(m log(n)).

With F-heaps O(n log(n) + m).

Application #3 : A better MST algorithm

Iteration i: We grow a forest, tree by tree, as follows.

Start with a singleton vertex and continue as in Prim's algorithm until either

1) The size of the heap is larger than k_i

2) Next edge picked is connected to an already grown tree

3) Heap is empty (if the graph is connected this will happen only at the very end)

Application #3 : A better MST algorithm

Contract each tree into a single vertex and start iteration i+1.

How do we contract ?

Do a DFS on the tree, marking for each vertex the # of the tree which contains it. Each edge e gets two numbers l(e), h(e) of the trees at its endpoints.

If h(e) = l(e) remove e (self loop).

(stable) Bucket sort by h(e) and by l(e), parallel edge then become consecutive so we can easily remove them.

O(m) time overall.

Complexity of iteration i:

Let n_i be the number of vertices in the i-th iteration.

O(m) inserts, O(m) decrease-key, $O(n_i)$ delete-min

total : $O(n_i \log(k_i) + m)$

Set $k_i = 2^{(2m/ni)}$ so the work per phase is O(m).

How many phases ?

Every tree in phase i is incident with at least k_i edges.

So $n_{i+1}\, k_i \le 2m_i \le 2m$

$\Rightarrow n_{i+1} \le 2m_i / k_i \le 2m / k_i$

$\Rightarrow k_{i+1} = 2^{(2m/n_{i+1})} \ge 2^{ki} \ge$

$$2^{2^{2^{2m/n}}}$$

Once $k_i \ge n$ we stop.

So the number of iterations is bounded by the minimum j such that

$$j \uparrow \quad 2^{2^{2^{2m/n}}} \ge n$$

$j = \min\{i \mid \log^{(i)}(n) \le 2m/n\} + 1 = \beta(m,n) + 1$

The overall complexity of the algorithm is $O(m\ \beta(m,n))$

Where $\beta(m,n) = \min\{i \mid \log^{i}(n) \le 2m/n\}$

for every $m \ge n$

$\beta(m,n) \le \log^{*}(n)$

For $m > n \log(n)$ the algorithm degenerates to Prim's so it actually $O(\min\{m\ \beta(m,n), n\log(n) + m\})$

Further research

Obtain similar bounds on the worst case.

Other generalizations/variations.

If the heap pointer is not given with the decrease-key and delete what then can be done ?

Better bound if the keys are integers in the RAM model.