

6.854 Problem Set 1

Nada Amin
namin@mit.edu

September 12, 2007

Collaborators

- Thomas Belulovich (thobel@mit.edu)
- Thomas Mildorf (tmildorf@mit.edu)

Problem 1

For each n , I exhibit a sequence of Fibonacci heap operations on n items that produce a heap ordered tree of depth $\geq n - 3 = \Omega(n)$.

For $n \leq 3$, the solution is trivial as the required lower-bound on the depth is ≤ 0 .

For $n \geq 4$, below is a sequence of operations on n items with keys from 1 to n (each item will be noted by its key) that produce a heap-ordered tree of depth $n - 3$, specifically either the tree $3 \rightarrow 4$ if $n = 4$ or, if $n > 4$, the tree $2 \rightarrow \dots \rightarrow n - 3 \rightarrow n - 1 \rightarrow n$, i.e. the tree with item 2 as a root and with each greater item except item $n - 2$ present as the only child of the previous smaller item present.

MAKE-TREE(n)

```
1  INSERT( $n$ )
2  INSERT( $n - 1$ )
3  INSERT( $n - 2$ )
4  DELETE-MIN()
5  for  $k \leftarrow 4$  to  $n - 1$ 
6      do INSERT( $n - 2$ )
7          INSERT( $n - k + 1$ )
8          INSERT( $n - k$ )
9          DELETE-MIN()
10         ▷ to DELETE( $n - 2$ ):
11         DECREASE-KEY( $n - 2, -\infty$ )
12         DELETE-MIN()
```

Problem 2

I show that modifying Fibonacci heaps so that a node is cut only after losing k children improves the amortized cost of DECREASE-KEY (to a better constant) at the cost of a worse cost for DELETE-MIN (by a constant factor).

DECREASE-KEY

I analyze the amortized cost of DECREASE-KEY, when a node is cut only after losing k children (i.e. having k marks). I choose the potential function:

$$\Phi = \frac{2}{k-1}(\text{number of marks}) + (\text{number of roots})$$

so that the cost of a cascading cut is 0:

$$\begin{array}{r} \\ + \quad \quad \quad 1 \quad (\text{for cutting the node}) \\ + \quad \quad \quad 1 \quad (\text{for adding a root}) \\ - \quad (k-1)2/(k-1) \quad (\text{for removing the } k-1 \text{ marks on this node}) \\ \hline \hline 0 \quad (\text{total}) \end{array}$$

Hence, the amortized cost of decrease-key is $2 + 2/(k-1)$:

$$\begin{array}{r} \\ + \quad \quad \quad 1 \quad (\text{for cutting the node}) \\ + \quad \quad \quad 1 \quad (\text{for adding a root}) \\ + \quad \quad 2/(k-1) \quad (\text{for adding a mark}) \\ + \quad \quad \quad 0 \quad (\text{for cascading cuts}) \\ \hline \hline 2 + 2/(k-1) \quad (\text{total}) \end{array}$$

As $k > 2$, the constant cost of DECREASE-KEY decreases.

DELETE-MIN

For DELETE-MIN, I show that the trees are exponential in degree, so that adding the children of the min as roots result in $O(\log(n))$ new roots.

By the union-by-rank procedure, the i^{th} added child will have degree $\geq i-k$. Indeed, by virtue of being the i^{th} child added, it must have been in the $(i-1)^{\text{th}}$ bucket, so it must have started with $i-1$ children. Since it loses at most $k-1$ children, it still has at least $i-k$ children. Now, let

$$S_k = \text{the number of descendants of a tree with } k \text{ children}$$

$$\begin{aligned}
S_0 &= 1 \\
S_1 &= 2 \\
S_n &\geq \sum_{i=0}^{n-k} S_i \\
S_n &\geq S_{n-1} + S_{n-k} \\
S_n &= \Omega(C^k)
\end{aligned}$$

Hence, DELETE-MIN will add $O(\log_C(n))$ roots. When $k = 2$, $C = \Phi$ (the golden number). Clearly, for $k > 2$, $C < \Phi$, so $\log_C(n) < \log_\Phi(n)$. Therefore, as $k > 2$, the performance of DELETE-MIN will be worse by a constant factor (precisely $\log_C(\Phi)$).

Problem 3

Part (a)

I show how to use a priority queue P that performs INSERT, DELETE-MIN and MERGE in $O(\log(n))$ time and MAKE-HEAP in $O(n)$ time to construct a priority queue Q that performs INSERT in $O(1)$ amortized time while still performing DELETE-MIN and MERGE in $O(\log(n))$ amortized time.

For my new priority queue Q , I maintain two structures, a list m and the priority queue p . On INSERT, I simply add the new element to the list m . On DELETE-MIN and MERGE, I make a priority queue out of the list m and merge it with the priority queue p , before calling the DELETE-MIN or MERGE procedure of p .

Q:MAKE-HEAP(l)

- 1 $m \leftarrow ()$
- 2 $p \leftarrow P:\text{MAKE-HEAP}(l)$

Q:INSERT(i)

- 1 $m.\text{ADD}(i)$

Q:DELETE-MIN(i)

- 1 $p2 \leftarrow P:\text{MAKE-HEAP}(m)$
- 2 $m \leftarrow ()$
- 3 $p.\text{MERGE}(p2)$
- 4 **return** $p.\text{DELETE-MIN}()$

Q:MERGE($p2$)

- 1 $p3 \leftarrow P:\text{MAKE-HEAP}(m)$
- 2 $m \leftarrow ()$
- 3 $p.\text{MERGE}(p3)$
- 4 $p.\text{MERGE}(p2)$

Amortized Analysis

Let $\Phi(Q)$ = number of elements in the list m . Then:

The amortized cost of INSERT is $O(1)$:

$$\begin{array}{r} \text{real cost : } 1 \\ \Delta\Phi : 1 \\ \hline \text{amortized cost: } 2 \end{array}$$

The amortized cost of DELETE-MIN is $O(\log(n))$:

$$\begin{array}{r} \text{real cost : } \Phi + O(\log(n)) \\ \Delta\Phi : -\Phi \\ \hline \text{amortized cost: } O(\log(n)) \end{array}$$

Similarly, the amortized cost of MERGE is $O(\log(n))$.

Part (b)

I show how even binary heaps can be modified to support INSERT in $O(1)$ amortized time while maintaining an $O(\log(n))$ time bound for DELETE-MIN.

As in part (a), on INSERT, I add the new element to a list m , which I make into a binary heap during DELETE-MIN. However, instead of merging, I add the new heap in a heap of heaps, a “super” heap.

Therefore, for my new priority queue Q , I maintain two structures: a list m and a heap of heaps s . The super-heap is keyed by the smallest element in each inner heap. On DELETE-MIN, I make a new inner heap out of the list m and add it to the super-heap. I find the min inner heap in the super-heap, perform DELETE-MIN on the min inner heap, and then MIN-HEAPIFY the super-heap to maintain the (super-)heap order (or delete the inner heap from the super-heap if it’s empty).

Q:MAKE-HEAP(l)

- 1 $m \leftarrow ()$
- 2 $s \leftarrow \text{P:MAKE-HEAP}(\text{P:MAKE-HEAP}(l))$

Q:INSERT(i)

- 1 $m.\text{ADD}(i)$

```

Q:DELETE-MIN( $i$ )
1  s.INSERT(P:MAKE-HEAP( $m$ ))
2   $m \leftarrow ()$ 
3   $h \leftarrow s.MIN()$ 
4   $x \leftarrow h.DELETE-MIN()$ 
5  if  $h.EMPTY()$ 
6      then  $s.DELETE(h)$ 
7      else  $s.MIN-HEAPIFY()$ 
8  return  $x$ 

```

Amortized Analysis

The amortized analysis is similar to part (a). In DELETE-MIN, deleting the min from the min inner heap and maintaining the heap order of the super-heap each takes $O(\log(n))$. Indeed, an inner heap has at most n elements while the super-heap has at most n inner heaps (with 1 element each). Therefore, as in part (a), the amortized cost of DELETE-MIN is $O(\log(n))$:

$$\begin{array}{r}
 \text{real cost : } \Phi + O(\log(n)) \\
 \Delta\Phi : \quad -\Phi \\
 \hline
 \text{amortized cost: } O(\log(n))
 \end{array}$$

Problem 4

I show how to use the techniques of persistent data structures to preprocess a tree in $O(n \log(n))$ time so as to allow LCA queries to be answered in $O(\log(n))$ time.

The time axis consists of the nodes of the tree in postorder. For each node, I maintain a timestamped pointer to its parent, a pointer to a binary search tree of timestamped names and an ephemeral rank for the union-by-rank heuristic.

Initially, at time $t = 0$, each node is isolated, having a null parent pointer and an empty tree of names and a rank of 0.

I support 3 operations:

UNION-WITH-NAME(a, b) increments time, UNION node a with node b and set the name of the representative to b .

FIND-AT-TIME(a, t) finds the representative of node a at time t .

NAME-AT-TIME(r, t) finds the name associated with the representative r at time t .

UNION-WITH-NAME(a, b)

```

1   $t \leftarrow t + 1$ 
2   $r_a \leftarrow \text{FIND-AT-TIME}(a, t)$ 
3   $r_b \leftarrow \text{FIND-AT-TIME}(b, t)$ 
4   $\triangleright$  apply the union-by-rank heuristic
5  if  $\text{rank}(r_a) = \text{rank}(r_b)$ 
6      then set parent pointer of  $r_b$  to  $r_a$  and timestamp it
7           $r \leftarrow r_a$ 
8           $\text{rank}(r_a) \leftarrow \text{rank}(r_a) + 1$ 
9  elseif  $\text{rank}(r_a) < \text{rank}(r_b)$ 
10     then set parent pointer of  $r_a$  to  $r_b$  and timestamp it
11          $r \leftarrow r_b$ 
12 else  $\triangleright \text{rank}(r_b) < \text{rank}(r_a)$ 
13     set parent pointer of  $r_b$  to  $r_a$  and timestamp it
14          $r \leftarrow r_a$ 
15 add the name  $b$  at time  $t$  to the binary search tree of names of  $r$ 

```

FIND-AT-TIME(a, t)

```

1  if  $\text{parent}(a)$  is null or  $\text{timestamp-parent}(a) > t$ 
2      then return  $a$ 
3      else return  $\text{FIND-AT-TIME}(\text{parent}(a), t)$ 

```

NAME-AT-TIME(r, t)

```

1  return the name at time  $t$  in the binary search tree of names of  $r$ 

```

During pre-processing, I simply call UNION-WITH-NAME(node, parent) for each node in postorder.

During a query for the pair (a, b) , I proceed as follows. Assuming, without loss of generality that node a occurs before node b in postorder, I let t be the time just before b was processed. The query call is then simply NAME-AT-TIME(FIND-AT-TIME(a, t), t).

Analysis

Thanks to the union-by-rank heuristic, I assure the union-find has only trees of depth $O(\log(n))$. Therefore, FIND-AT-TIME is $O(\log(n))$. UNION-WITH-NAME is therefore $O(\log(n) + \log(t))$ and NAME-AT-TIME is $O(\log(t))$. Since $t \leq n$, all operations are $O(\log(n))$. Therefore pre-processing is $O(n \log(n))$ and querying is $O(\log(n))$.

Problem 5

I show that the expected behavior of markless Fibonacci heaps (where, each time I do a cut, I flip an unbiased coin to decide whether to cascade the cut to

the parent) is like that of the standard ones. Indeed, the expected number of children cut before a node is cut, E_{cut} , is 2:

$$\begin{aligned}
 E_{\text{cut}} &= \frac{1}{2} && + 2 \cdot \frac{1}{4} && + 3 \cdot \frac{1}{8} && + \dots \\
 &= \left(\frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots\right) && + \left(\frac{1}{4} + \frac{1}{8} + \dots\right) && + \left(\frac{1}{8} + \dots\right) && + \dots \\
 &= 1 && + \frac{1}{2} && + \frac{1}{4} && + \dots \\
 &= \sum_{i=0}^{\infty} \frac{1}{2^i} \\
 &= 2
 \end{aligned}$$

It is possible to bias the coin so that a cascade is more than 50% likely to achieve the effect of cascading after (say) one and a half children are cut. Let p be the probability of cascading a cut, i.e. the bias of the coin. Then:

$$\begin{aligned}
 E_{\text{cut}} &= p + 2 \cdot p^2 + 3p^3 + \dots \\
 &= \sum_{i=0}^{\infty} p^i \\
 &= \frac{1}{1-p}
 \end{aligned}$$

So in order to have $E_{\text{cut}} = 1.5 = \frac{3}{2}$, we need a bias of $p = \frac{2}{3}$.

This can be used to improve the expected amortized time for DELETE-MIN at the cost of increasing the expected amortized time for DECREASE-KEY. The analysis is similar to Problem 2, with $k = 1.5 < 2$: DECREASE-KEY is more costly because the expected number of cascading cuts is higher, DELETE-MIN is less costly because the expected exponent in the exponential-in-degree trees is greater.