

External Priority Queue

Goal: Insert: $O_B(\frac{1}{B} \cdot \log_{N/B} N/B)$
DelMin: $O_B(\frac{1}{B} \cdot \log_{N/B} N/B)$


[Finkel et al, 99]
(97)

Space: $O(N/B)$

[CPU: $O(\log_2 N)$ for ops.]

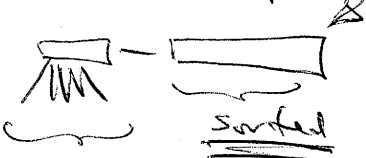
\Rightarrow Efficient external HeapSort ($= N \cdot \text{Insert} + N \cdot \text{DelMin}$)

Note for this article: $B = P$
Min = Max (w/ max heap)

RAM = Insert buffer
(an internal heap) 
 $M' = \Theta(M)$ elms. (Eg $M' = M/2$)


Heap: Tree of fanout $\Theta(M/B)$
Shape:



Up to M' elements in node buffer
Nodes: 
 $\Theta(M/B)$ children
sorted

Heap ordered: Elms in node buffer dominates any elms in subtree.

Last node



- Parent ^{is the} only internal node of fanout $< \frac{M'}{D}$
- Node buffer may contain $\leq \frac{M'}{2}$ elems.

All other nodes : Buffer contains at least $\frac{M'}{2}$ elems and at most M elems.

Operations

Insert : Insert new elm. into input buffer (heap in RAM)

If this buffer gets size $\geq M'$:


Make a new last leaf

Move input buffer contents to new nodes' buffer (first sort elems)

Swap new node with previous last leaf.

SiftUp on both nodes

Sift Up



← Some elements in conflicts (heap-order-wise) with ancestors. Else tree OK.

[Note: can be checked by comparing strongest elem. in v with weakest in $p(v)$]

Merge the two (sorted) node buffers.

Whether (wrong?) stated in paper →

Split into same sized buffers as before.
↑
sorted outcome

Think/vari-
f's

Note: Heap order restored between v and $p(v)$.

Heap order between $p(v)$ and $p(p(v))$ may be violated. ⇒ recurse SiftUp on $p(v)$.

Delete Best :

Check root of tree [best block of its buffer kept in RAM].

and insert buffer (a heap).

If best is in insert buffer, remove & return.

Else remove from root node buffer

If |root buffer| < $N/2$:

Refill (root).

Clean Up leaves

Refill (u)

Do $N/2$ merge steps on childrens buffers

If not possible (has last leaf as only child),

Add to list of leaves to be cleaned up. (removing child)

Else for each child w :

If buffer size < $N/2$

If not leaf :

Refill (w)


Else

Add to list of leaves to be cleaned up.

do as many steps as possible (\Rightarrow becomes a leaf)

Note: If there exists just one child \neq last leaf, then these steps can be performed (due to size of childrens buffers).

~~After~~ recursive
Post-ord, Before
CleanUpLeaves:



x = leaves to clean up. (= all leaves with $< \frac{M'}{2}$ elems. in buffers) \uparrow
["underfull"]

CleanUpLeaves

current

While \exists underfull leaf $v \neq$ last leaf w :

i) If $|v| + |w| > M'$: $\left\{ \begin{array}{l} \text{may change during} \\ \text{cleanUpLeaves} \end{array} \right\}$

Move $M' - |v|$ strongest elem in w 's buffer to v 's buffer (merge)

SiftUp(v)

ii) Else if $\frac{M'}{2} \leq |v| + |w| \leq M'$,

Merge w 's buffer into v
Delete(w)
SiftUp(v)


iii) Else: // $|v| + |w| < \frac{M'}{2}$

| Same as

// v is still underfull

Analysis

N elm. stored \Rightarrow max $\frac{N}{M/2} + 1$ nodes



$$\text{Height} = \lceil \log_{\text{fanout}} (\# \text{ nodes}) \rceil$$

$$= O(\log_{M/B} (N/M))$$

There are $\Omega(M)$ insert between each emptying of input buffer. Change everything to these inserts:

1) $O(1)$ SiftUps just after input buffer emptying.

Price: Height \times Scan(M)

$$= O\left(\frac{M}{B} \cdot \log_{M/B} \left(\frac{N}{M}\right)\right)$$

I.e. $O\left(\frac{1}{B} \log_{M/B} \left(\frac{N}{M}\right)\right)$ per inserted element.

[Binary merge of two buffers in nodes]

2) Lifting of elements upwards in tree during

Refills: $M/2$ merge steps (lifting $M/2$ elms one level) cost in I/O's:

(Initiate Merge = load first block of children's buffers)

$$\text{Fanout} + \text{Scan}\left(\frac{M'}{2}\right) = O\left(\frac{M}{B}\right) + O\left(\frac{M}{B}\right) = O\left(\frac{M}{B}\right)$$

So if each of the $N'/2$ lifted elements gets charged $O(\frac{1}{B})$, then the price of the lift is covered.

Therefore we define the invariant that a node buffer with m elements, residing in a node at depth d , has

$$d \cdot m \cdot \frac{1}{B}$$

So $O(\frac{N}{B} \cdot \log_{N/B}(\frac{N}{M}))$ credits for this must be supplied when the insertion buffer overflows during an insertion

credits, then the lift will release credits covering the charging [as $N'/2$ elements moves from depth d to depth $d-1$].

Note: here we use that siftups do not change sizes of buffers [p. 3, middle], so they don't change credits in buffers.

3) Each step in CleanUp Leaves does ^{zero or} one binary merge of (parts of) node buffers, and one SiftUp. This costs $O(N/B) + O(\frac{N}{B} \log_{N/B}(\frac{N}{M}))$.

Each step involves the current last leaf.

Since $|v| < N'/2$ (v is underfull), then in case i) [on page 5] we move $N' - |v| > N'/2$

⑧

elements from current last leaf's buffer. Hence, for a given current last leaf, case i) can only happen once. Next step will be case ii) or iii), which removes the node.

Therefore we when creating a node gives it $2 \times O\left(\frac{M}{B} \log_{M/B}(N/M)\right)$ credits, to cover the cost of the (up to) two steps of CleanUpLeaves where it is the current last leaf.

[Note: since a node becomes last leaf, it will stay that way until deleted, also in between CleanUpLeaves]

Summing up: if we charge the creation of a new node [during an insert overflowing the insert buffer] $O\left(\frac{M}{B} \log_{M/B}(N/M)\right)$, then all the cost/credits of 1) + 2) + 3) are covered.

↑
(for new node and its buffer)

This is $O\left(\frac{1}{B} \log_{M/B}(N/M)\right)$ per element of the overflowing insert buffer, and all of these elements were inserted since last such overflow. Thus charging each Insert operation $O\left(\frac{1}{B} \log_{M/B}(N/M)\right)$ /s covers all costs of the structure. □