

Worst-Case Efficient

External-Memory

Priority Queues

Gerth Stølting Brodal

MPI Saarbrücken

Joint work with

Jyrki Katajainen

University of Copenhagen

Priority Queues

Maintain a set of elements from a totally ordered universe under:

$\text{Insert}(x)$ - Insert element x into the set

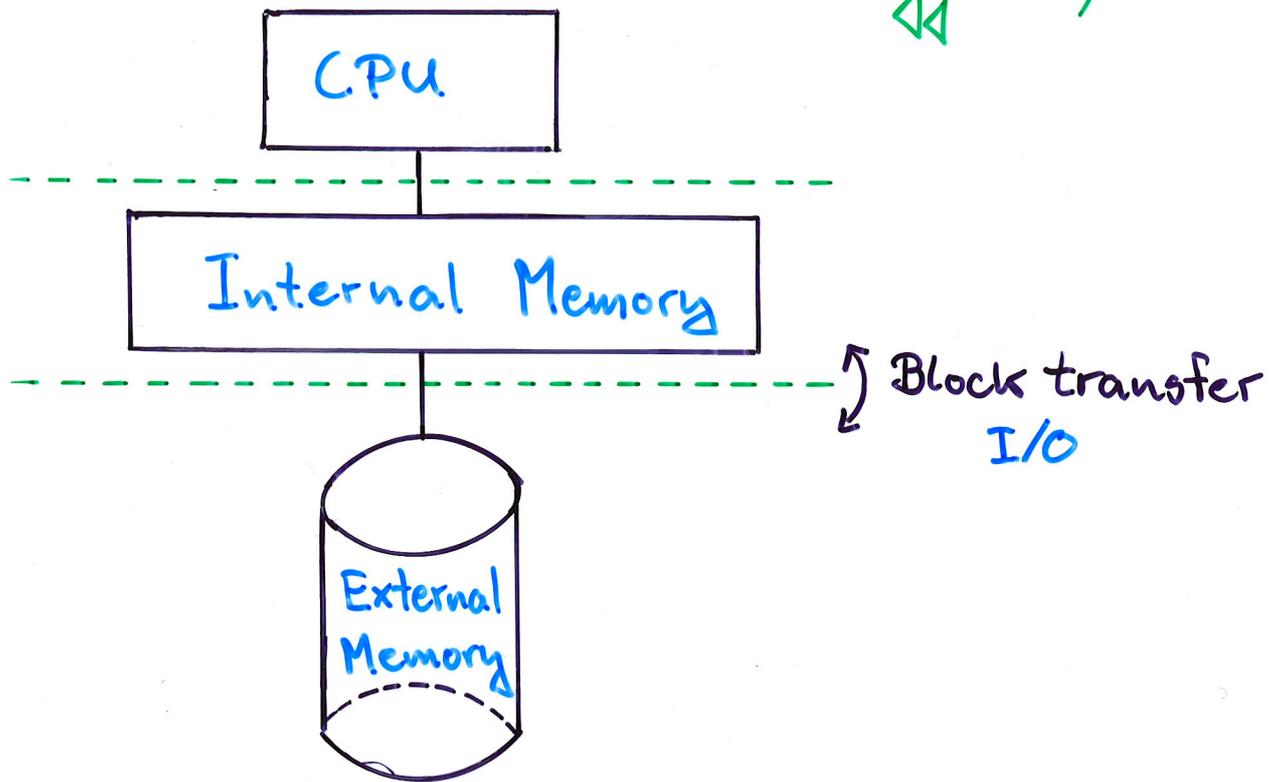
$\text{DeleteMin}()$ - Delete and return the minimum of the set

Insert and Delete can be implemented with $\Theta(\log_2 N)$ comparisons

Williams '64

External Memory Model

Aggarwal, Vitter '88



M = Size of the internal memory

B = I/O block size

N = # elements in the priority queue

Assumptions:

$$M \geq 23B, \quad B \geq \log_{\frac{M}{B}} \frac{N}{M}$$

Complexity measures:

Internal : # comparisons

External : # I/Os

Preliminaries

Merging: $\Theta(\frac{M}{B})$ sorted lists containing N elements stored in $\Theta(\frac{N}{B})$ blocks can be merged into one list with $\Theta(\frac{N}{B})$ I/Os and $\Theta(N \cdot \log_2 \frac{M}{B})$ comparisons.

Proof: $\Theta(\frac{M}{B})$ -ary merging

Sorting: N elements in $\Theta(\frac{N}{B})$ blocks can be sorted with $\Theta(\frac{N}{B} \cdot \log_{M/B} \frac{N}{M})$ I/Os and $\Theta(N \log_2 N)$ comparisons.

Proof: $\Theta(\frac{M}{B})$ -ary mergesort

(i) Sort in internal memory all sequences of m elements.

(ii) In $\log_{M/B} \frac{N}{M}$ iterations apply the above merging algorithm to obtain lists of length $M \cdot (\frac{M}{B})^i$

The sorting algorithm is optimal wrt. both comparisons and I/Os.

Aggarwal, Vitter '88

Worst - Case Efficiency

Internal: Each priority queue operation should require worst-case $\Theta(\log_2 N)$ comparisons

External: The I/Os should be divided evenly among the operations, i.e., one I/O for every $\Theta\left(\frac{B}{\log_{M/B} \frac{N}{B}}\right)$ operation.

External-Memory Priority Queues

- Heaps Williams '64

The random memory accesses implies that each CPU operation causes an I/O...

- Fishspear Fischer, Patterson '94

$\Theta\left(\frac{N}{B} \cdot \log_2 N\right)$ I/Os for N operations.
Uses $O(1)$ push down stacks

- Heap + B elements/node Wegner, Teuhola '89

$O\left(\log_2 \frac{N}{B}\right)$ I/Os for every B 'th operation.
Optimal for $M = O(B)$

→ Buffer tree Arge '95

$O\left(\frac{N}{B} \log_{M/B} \frac{N}{M}\right)$ I/Os for N operations.

→ Heap with buffers Fadel et al. '97

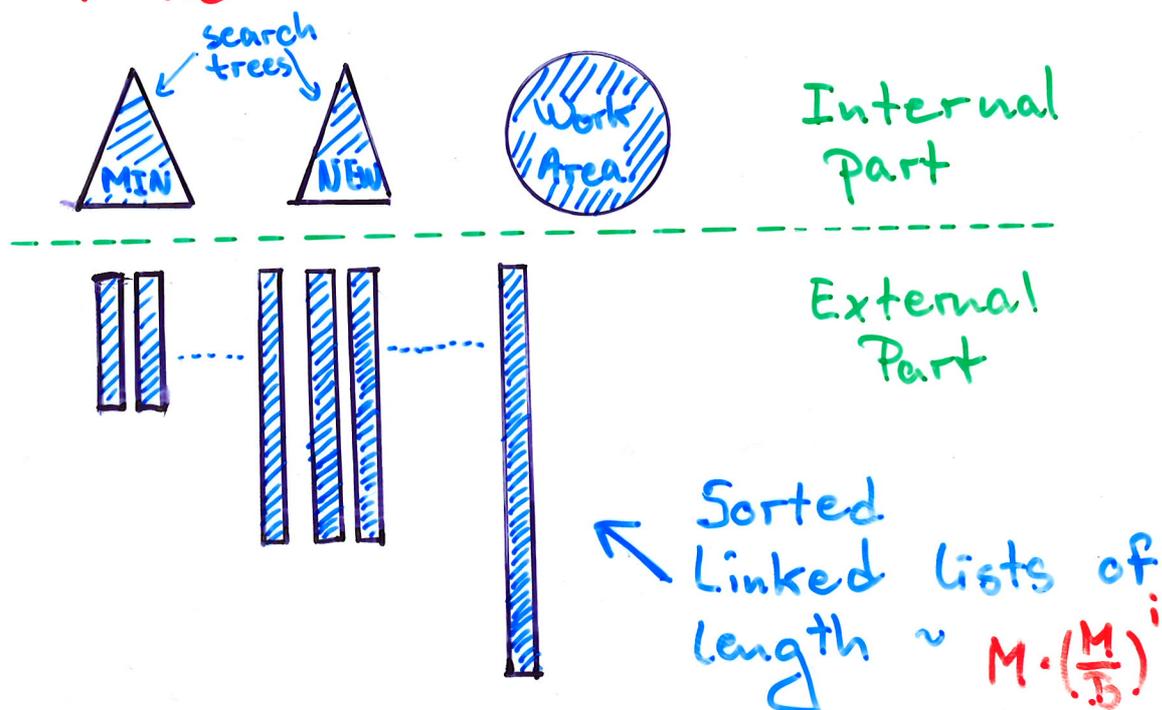
Same bounds as for buffer trees

→ No worst-case guarantee on the individual operations

New Result

DeleteMin and Insert can be done with worst-case $\Theta(\log n)$ comparisons per operation, and 1 I/O for every $\Theta(B / \log_{M/B} \frac{N}{B})$ operations, which is the best possible.

Outline of the Data Structure



- Min stores the overall $O(M)$ smallest elements \rightarrow fast DeleteMin
- New stores the most recently inserted $O(M)$ elements \rightarrow fast Insert
[Small elements are inserted directly into Min]
- The external part stores exactly B elements in each block (all larger than the elements in MIN)

Operations

$$k = O(M)$$

Insert(x)

- If x small ($\leq \max \text{MIN}$) then
Swap x and $\max \text{MIN}$
- Insert x into NEW
- If $|\text{NEW}| \geq k$ then
Perform Batch Insert_k with k
elements from NEW

Delete Min

- If $\text{MIN} = \emptyset$ then
Perform Batch Delete_k
Merge the result with NEW
Move the k smallest elements to MIN
- Delete and Return $\min \text{MIN}$

Batch Insert_k

- Create a new (ext.mem.) list of length k and assign it **rank 1**
- $i := 1$
- while #lists of rank $i = \frac{M}{B}$ do
Merge the $\frac{M}{B}$ rank i lists
and assign it rank $i+1$
 $i := i+1$

Batch Delete_k

- Delete the k least elements from the lists in external memory

$$\text{rank } L = \left\lfloor \log_{M/B} \frac{|L|}{B} \right\rfloor, \quad \# \text{lists} \leq \frac{M}{B} \cdot \log_{M/B} \frac{N}{M} \quad 5.$$

Lemma

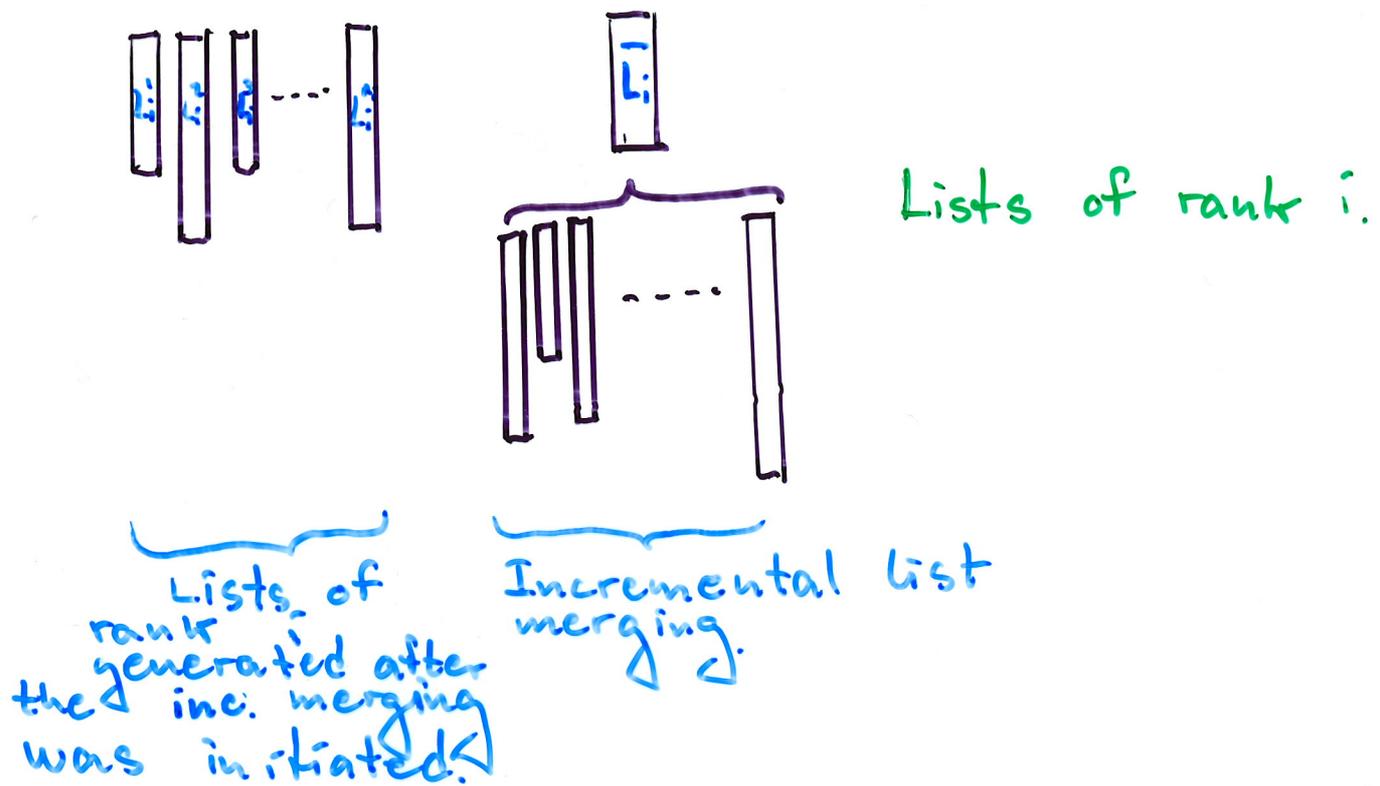
The described data structure supports Insert and DeleteMin with **amortized** $\Theta(\log u)$ comparisons and $\Theta\left(\frac{1}{B} \log_{M/B} \frac{N}{B}\right)$ I/Os.

[the same bounds follow by using buffer trees].

Worst-case bounds?

- Incremental list merging; similar as done by Thorup '96
- Incremental batch operations
- Perform BatchDeleteMin_k before MIN gets empty

Incremental list merging



Merge Step (i)

- Perform k steps of the list merging at rank i .
- If the merging is finished, then make the resulting list L_i a list of rank $i+1$ (promote the list)

Batch Insert k

- Create a new rank 1 list
- $\forall i$: Merge Step (i)

Deletions

BatchDeleteMin_k deletes and returns the k least elements of the lists.

⇓
The length of the lists decrease

⇓
The maximum rank of a list increases

Possible solution: Incremental global rebuilding

Our solution:

1) If the incremental merging of rank i finishes, but $|\bar{L}_i| < k \left(\frac{M}{B}\right)^i$, then we do not promote \bar{L}_i to a rank $i+1$ list.

2) Let R be the maximum rank of a list. If the resulting list $|\bar{L}_R| < k \left(\frac{M}{B}\right)^{R-1}$, then \bar{L}_R gets rank $R-1$.

3) If k elements are deleted from \bar{L}_i , we perform MergeStep(i)

4) Always perform MergeStep(R) after BatchDeleteMin

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

Conclusion

The first worst-case optimal external memory priority queue implementation.

Open problems

- Does "worst-case" make sense in practice?
- Applications!
- Implementation and experiments.
- Can buffer trees be deamortized, i.e., how can buffer trees support updates with worst-case $\Theta(\log_{M/B} \frac{N}{B})$ I/Os for B updates?