



Navigation Piles with Applications to Sorting, Priority Queues, and Priority Deques

Jyrki Katajainen and Fabio Vitale
Department of Computing, University of Copenhagen
Universitetsparken 1, 2100 Copenhagen East, Denmark
{jyrki|fabio}@diku.dk



Navigation Piles

n **Using of the new data structure**

- n Sorting algorithms
- n Priority queues
- n Priority dequeues
- n Implementation of other data structures



Structure Properties

Shape

n It is a complete binary tree of size $2^{\eta+1} - 1$ ($n \leq 2^\eta$)

Leaves

n The first n leaves store one element each

n The remaining are empty

Branch nodes

n Each branch node stores the index – inside the leaf sequence dominated by the node – containing the relative top element

Representation

n The sequence $A[0..n)$ stores the elements

n The sequence $B[0..2^{\eta+1})$ stores the navigation information



Structure Elements

n **Element indices:**

Indices of the elements stored in $A[0..n)$

n **Bit indices**

Indices of the bits stored in $B[0..2^{\eta+1})$

n **Branch node indices**


n **Leaf indices**

n **Levels** (no larger than η)

n **Offsets**

Indices stored in the branch nodes

Navigation Pile Example



A navigation pile of size 14 and capacity 16. The normal parent/child relationships are shown with dotted arrows and the references indicated by the offsets with solid arrows. The gray nodes are not in use.



Number of extra bits

n Number of branch nodes with depth δ ($0 \leq \delta < \eta$)

n 2^δ

n Maximum number of leaves dominated by a branch node whose height is γ

n 2^γ

n Number of bits used for the offsets

n
$$\sum_{\delta=0}^{\eta-1} 2^\delta (\eta - \delta) < 2^{\eta+1}$$



Operations

ⁿ **Construction**

- ⁿ Visiting the branch nodes in a bottom-up manner (depth-first order visit to improve the cache performances)
- ⁿ **$n-1$** element comparisons
 - ⁿ $n-1$ branch nodes with more than 1 child
- ⁿ **n** element moves
 - ⁿ n elements copied to the container sequence
- ⁿ **$O(n)$** instructions
 - ⁿ $O(1)$ instructions for every comparison and move



Operations

n ***top()***

n **$O(1)$** instructions

n Return the index stored at the root

n ***push()*** ($n < 2^n$)

n $\lceil \log_2(\lceil \log_2 n \rceil + 1) \rceil$ element comparisons

n Binary search on the path from the new last leaf to the root

n **1** element move

n The element is appended at the end of the sequence $A[0..n)$

n **$O(\log n)$** instructions



Operations

ⁿ *pop()*

ℓ : index of the last leaf ($n - 1$)

m : index of the top element

ⁿ **Case 1:** $m = \ell$

ⁿ The top element is erased and the offsets on the path from the new last leaf to the root are updated

pop() Case 1




Illustration of Case 1. The offset of the root is referring the last leaf. The leaves whose contents may change are indicated in magenta. In dark blue are indicated the branch nodes whose updating offset require one element comparison each. When updating the contents of the light blue branch nodes, no element comparisons are necessary.



Operations

i_1 : first ancestor of the last leaf having two children (in use)
($n > 1$)

i_2 : second ancestor of the last leaf having two children (in use) ($\ell \neq 2^k \forall k \in \mathbb{N}$)

k, j_1 : index of the leaves referred respectively by the offset stored at the branch node with index i_2 and at the first child of that one with index i_1 (or $j_1 = \ell - 1$ if ℓ is odd)

Case 2: $m \neq \ell$ and $k \neq \ell$ (or $m \neq \ell$ and ℓ is a power of 2)

$A[m] \beta A[n]$ and the element copied is erased

The offsets stored at the branch node on the path from $[i_1.. i_2)$ are updated to refer the leaf at the position j_1

The offsets stored on the path from the leaf at the position m to the root are updated

pop() Case 2




Illustration of Case 2. The offset of i_2 is not referring the last leaf. The offsets of the branch nodes in light blue will refer the leaf whose content is the number 4, and the the top element (44) will be overwritten by the current last leaf (22).



Operations

j_2 : index of the element referred by the offset stored at the first child of the branch node with index i_2

Case 3: $m \neq \ell$ and $k = \ell$

$A[m] \beta A[j_2]$

$A[j_2] \beta A[\ell]$

The element at the last leaf is erased

The offsets stored at the branch node on the path from $[i_1.. i_2)$ are updated to refer the leaf at the position j_1

The offsets stored on the path from i_2 (including it) upwards are updated to refer the leaf in the position j_2 if they referred earlier the last leaf

The offsets stored on the path from the leaf at the position m to the root are updated

pop() Case 3




Illustration of Case 3. The offset of i_2 is referring the last leaf. The offsets of the branch nodes in light blue from the bottom til i_2 (excluding it) will refer the leaf j_1 (4); the offset relative to the branch node i_2 will refer the leaf j_2 (15); the the top element (44) will be overwritten by the element in the position j_2 (15); finally the element of the current last leaf (22) will overwrite that one of the leaf j_2 .



Operations

- n At most $\lceil \log_2(n-1) \rceil$ element comparisons
 - n In each of the three cases only one path updating involves element comparisons
 - n The depth of the root becomes $\lceil \log_2(n-1) \rceil$ and at most one element comparisons is done at each level
- n **0, 1 or 2** moves
- n **$O(\log n)$** instructions



Sorting

- n In connection with *pop()* the top element is not overwritten, but it is saved at the earlier last leaf (*pop_and_save()* function)
- n The basic version of the sorting algorithm (Pilesort) is a sequence of $n-1$ *pop_and_save()* operations
- n The amount of extra space needed is at most **$4n$ bits** because $2^{\lceil \log_2 n \rceil} < 2n$



Sorting

n Element comparisons

n $n-1$ comparisons for the construction

n
$$\sum_{k=2}^{n-1} \lceil \log_2 k \rceil < n \log_2 n - 0.91n$$

n If we find the bottom element during the construction with at most $\lceil n/2 \rceil + 1$ comparisons (and keeping a separate copy of it) in order to reduce the number of moves, the bound becomes $n \log_2 n + 0.59n + O(1)$

n Element moves

n Not more than $4n + O(1)$ in the basic version

n Finding at the beginning the bottom element the bound becomes $2.5n + O(1)$ (for every odd value of ℓ if we need 3 moves for the execution of *pop_and_save()*, we will need at most 2 moves in the following iteration)



Priority Queues

- n **Dynamization of the container storing the elements**
 - n Dynamization techniques developed by Katajainen and Mortensen
 - n Space-efficient resizable arrays
- n **Dynamization of the bit sequence**
 - n Bit sequence divided into blocks B_* of size 2^5 and B_h of size 2^h for $h \geq 5$
 - n Bit sequence of the block B_* is kept in memory at all time
 - n The largest block can be empty (deamortization of the cost of allocations and deallocations in proximity of block boundaries)
 - n If there is an empty block and the largest nonempty block lacks more than 2^5 elements, the bit sequences for the largest block is freed

Priority Queues (Illustration)




Illustration of the Priority Queue structure. A collection of navigation piles storing 77 elements. A space-efficient resizable array is used as the container for the elements. The gray areas denote empty memory segments allocated for dynamization purposes.



Priority Queues

Space

- 2^{h+1} bits for every B_h ($h \geq 5$)
- $4n + O(1)$ extra bits for all the priority queue

Construction

- $n + \Theta((\log_2 n)(\log_2 \log_2 n))$ element comparisons
- n element moves
- $O(n)$ instructions

top()

- We assume that *push()* and *pop()* maintain the top and iterator headers
- $O(1)$ instructions (following the cursor of the iterator headers)



Priority Queues

push()

- n* The element is inserted into the largest nonempty block (or, if it is full, into a new larger block and a new navigation pile)
- n* $2\log_2 \log_2 n + O(1)$ element comparisons (the iterator and the top headers can need to be updated)
- n* 1 element move
- n* $O(\log n)$ instructions

pop()

- n* The top element is erased by overwriting with an element taken from the largest nonempty block
- n* We need to update also the top and iterator header
- n* $\log_2 n + \log_2 \log_2 n + O(1)$ element comparisons
- n* 2 element moves
- n* $O(\log n)$ instructions



Priority Deques

- n We pair the elements and we sort the resulting pairs
- n If n is odd, we keep one element in a special block B_0
- n The elements get partitioned in two distinct collections:
 top-element candidates and bottom-element candidates
- n We build the extra information for the two priority queues
- n **Construction**
 - n Not more than $1.5n + \Theta((\log_2 n)(\log_2 \log_2 n))$ comparisons and $n+1$ moves
- n ***top()*** and ***bottom()***
 - n B_0 can cause 1 comparison



Priority Deques

ⁿ *push()*

- ⁿ If B_0 is empty the element being inserted is copied there
- ⁿ If B_0 contains an element this and the new one become twins, and this pairs is added to the data structure
- ⁿ In the latter case the resource bounds are **2 times** those compared to the *push()* function



Priority Deques

- n ***pop_top()*** and ***pop_bottom()***
 - n These two functions are symmetric
 - n We consider ***pop_top()***
 - n If B_0 is empty we move the twin of the top to B_0 , take a pair of elements from the last nonempty block
 - n We carry out a partial path updating in the last nonempty block
 - n We move the pair of elements from the last nonempty block to the block containing the top element and its twin updating a full path of extra information for both the elements



Priority Deques

- n If B_0 contains an element this is used now as the new twin of the twin of the top element
- n After a single comparison the two elements are moved in their correct locations and we update at most two paths of extra bits
- n The resource bounds are about twice as high as those for the priority queues



Conclusions

- n It have been showed how it is possible to build and use a new data structure with which, in spite of low memory requirements, the worst-case bounds for the number of element comparisons, element moves and other instructions are close to the absolute minimum
- n Is it possible to get similar performance bounds with a sublinear number of extra bits (for example, a sorting algorithm using $n \log_2 n + o(n \log_2 n)$ comparisons, $O(n)$ moves, and $o(n)$ extra bits)?
- n Is it possible to improve the *push()* function requiring a constant time without making worse the worst-case bounds for the other operations?