

# Partition-based Simple Heaps

Gerth Stølting Brodal\*       John Iacono†   
Casper Moldrup Rysgaard\*       Sebastian Wild‡ 

March 17, 2026

## Abstract

We introduce a new family of priority-queue data structures: *partition-based simple heaps*. The structures consist of  $\mathcal{O}(\lg n)$  doubly-linked lists; order is enforced among data in different lists, but the individual lists are unordered. Our structures have amortized  $\mathcal{O}(\lg n)$  time extract-min and amortized  $\mathcal{O}(\lg \lg n)$  time insert and decrease-key.

The structures require nothing beyond binary search over  $\mathcal{O}(\lg n)$  elements, as well as binary partitions and concatenations of linked lists in natural ways as the linked lists get too big or small. We present three different ways that these lists can be maintained in order to obtain the stated amortized running times.

## 1. Introduction

The *priority queue* is a fundamental comparison-based abstract data type used in numerous efficient algorithms, such as discrete event simulation, Dijkstra’s shortest path algorithm and Prim’s minimum spanning tree algorithm. A priority queue (a.k.a. *heap*) stores a set of elements and has to implement the following operations, where we assume each element has an associated key from a comparison-based universe:

- $\text{INSERT}(e)$ : Add an element  $e$  to the the heap and return a pointer  $ptr$  to it.
- $\text{DELETE-MIN}()$ : Remove and return the element in heap with smallest key.
- $\text{DECREASE-KEY}(ptr, key)$ : Change the key of the element pointed to by  $ptr$  in the heap to a smaller key  $key$ .

A standard undergraduate computer-science curriculum discusses binary heaps [Wil64], which realize all operations in  $\mathcal{O}(\lg n)$  time<sup>1</sup> when the data structure currently contains  $n$  elements (as would a balanced binary search tree). However, much faster updates for priority queues are possible: The most efficient heaps such as (strict) Fibonacci heaps [FT87, BLT25] achieve *constant* (amortized) time for  $\text{INSERT}$  and  $\text{DECREASE-KEY}$  and  $\mathcal{O}(\lg n)$  time for  $\text{DELETE-MIN}$ .

---

\*Aarhus University, Denmark, {gerth,rysgaard}@cs.au.dk

†Université libre de Bruxelles, Ixelles, Belgium, john.iacono@ulb.be

‡University of Marburg, Germany, wild@informatik.uni-marburg.de; University of Liverpool, UK

<sup>1</sup>Here and throughout we let  $\lg$  denote the binary logarithm.

While optimal in theory, these *efficient priority queues* have some downsides. First, their structure and analysis are less intuitive to teach and implement. Second, the constant factors in their running time and space usage are substantially higher than for binary heaps; not least because they inherently rely on pointer-based representations of trees, with the ensuing poor locality of reference.

The desire to simplify efficient priority queues and to make them more efficient in practice has resulted in a number of alternative data structures, which we review briefly in Section 1.1. They all retain somewhat complicated structures and invariants, are nontrivial to analyze, or require  $\Omega(\lg n)$  time for updates.

In this paper, we present simple priority-queue implementations solely based on linear data structures, showing that simple quicksort-style partitioning of elements into unsorted sets suffices for amortized  $\mathcal{O}(\lg \lg n)$  time INSERT and DECREASE-KEY operations and amortized  $\mathcal{O}(\lg n)$  time DELETE-MIN operations – provided it is governed by a carefully chosen “*pivot-forgetting rule*” to keep the number of sets in  $\mathcal{O}(\lg n)$ . We present several such rules, demonstrating a rich design space. This brings us much closer to the efficient priority queues above, but with entirely elementary techniques and simple designs. As a case in point, for the arguably simplest rule (“lazy partition heaps”), we give a fully self-contained description including its amortized analysis suitable for teaching in undergraduate algorithms classes in Sections 2 and 3.

## 1.1. Previous Work

See [Bro13] for a full survey of priority queues. We begin the history of priority queues with fast DECREASE-KEY which begins with the Fibonacci heap [FT87]. They support constant-amortized time DECREASE-KEY and INSERT, and  $\mathcal{O}(\lg n)$  amortized-time DELETE-MIN. Since then a variety of other heaps with fast DECREASE-KEY have been introduced:

- The pairing heap [FSST86], which is a self-adjusting structure modeled after splay trees, and which we do not know the asymptotic amortized running time of DECREASE-KEY. Only  $\mathcal{O}(2^{\sqrt{\lg \lg n}})$  [Pet05] and  $\Omega(\lg \lg n)$  [Fre99] are known. A multipass variant has been shown to have an  $\mathcal{O}(\lg \lg n \cdot \lg \lg \lg n)$  amortized-time DECREASE-KEY [ST23a].
- The strict Fibonacci heap, which has worst-case running times [BLT25].
- Chan’s quake heaps, invented to be a simpler alternative to Fibonacci heap with  $\mathcal{O}(1)$  amortized time DECREASE-KEY [Cha13].
- Rank-pairing heaps [HST09] which are another alternative with  $\mathcal{O}(1)$  amortized time DECREASE-KEY, but where the structure is closer to the pairing heap.
- Violation heaps [Elm10], hollow heaps [HKTZ17], and Fibonacci heaps revisited [KTZ14] are other variants with with  $\mathcal{O}(1)$  amortized-time DECREASE-KEY.
- Slim and smooth heaps [ST23b] have  $\mathcal{O}(\lg \lg n)$  DECREASE-KEY, but are in the pointer-based model with constant indegree.

### 1.1.1. Lower bounds

There are two lower bounds for DECREASE-KEY in heaps, one due to Fredman [Fre99] and the other due to Iacono and Özkan [IÖ14]. Both of these bounds only apply to heap-based priority queues that follow two different particular models, each of which includes pairing heaps, and neither of which applies to our new structures. The Fredman lower bound focuses

on the tradeoff between augmented data and the decrease-key time. The Iacono and Özkan bound focuses on structures that are pointer-based. Both of these results point to the fact that  $\mathcal{O}(\lg \lg n)$  is arguably a natural time for heaps, but both only apply to (different) classes of generalized heaps which are quite specific for their adversary arguments to work.

All of the above points to  $\mathcal{O}(\lg \lg n)$  as the right running time for DECREASE-KEY in a pointer model structure with constant indegree. Slim and smooth heaps achieve this, and pairing heaps are conjectured to. All of the other heaps require either RAM-model tricks or non-constant indegree. For example, Fibonacci heaps need random access into an array of size  $\mathcal{O}(\lg n)$  in order to combine heaps with the same rank (which is an integer with logarithmic range). One could of course replace this array with a balanced binary search tree, but this would introduce an  $\mathcal{O}(\log \log n)$  overhead. In [FT87], it is shown how this array can be removed, albeit by using nodes of linear indegree.

### 1.1.2. Cache-oblivious inspiration

We point to the works of Chowdhury and Ramachandran [CR18] and Brodal, Fagerberg, Meyer, and Zeh [BFMZ04] as sources of inspiration as they are devoid of treeology. Both of these, which were designed for the cache-oblivious model, use two sequences of arrays of increasing size; superficially it looks like two copies of our data structure where items are inserted into one set, percolate down, at some point move to the other set, and then percolate up. Being designed for the cache-oblivious model, the lists are stored in arrays.


### 1.1.3. Lazy search trees

Also related are lazy search trees [SW20, SZ22, RW25], a data structure that smoothly interpolates between priority queues and binary search trees; depending on what queries are used, INSERT can be as fast as in the best priority queues, but arbitrary sorted dictionary queries are supported, as well. Our structures can be seen as a one-sided instance of the lower two tiers of the original lazy search trees with slightly modified operations.

### 1.1.4. Quickheaps

Other related work takes a more application-focused view of making priority queue implementations fast in practice. While asymptotic performance matters, constant-factor overhead and locality of reference can, for all reasonable input sizes, dwarf the distinction between logarithmic and sublogarithmic complexities. Quickheaps [PN06, NP10] and their “stronger” refinements [NPPS11] keep data in a single array and conceptually apply successive quickselect invocations for queries, where past partitioning steps are remembered and skipped.

This works well if updates come in random order; for robustness against adversarial inputs, stronger quickheaps take inspiration from schemes for balancing binary search trees to ensure good amortized performance in quickheaps: either using randomization or using a form of weight-balanced binary search tree. Being designed for the much harder problem of keeping an entire binary search tree in shape, the resulting concatenation rules for quickheaps are complicated. We point out that not only the concatenation rule itself, but also the amortized analysis in [NPPS11] is substantially different from those in this paper. They also cite favorable performance in the external-memory model as a benefit of quickheaps (in any variant). Our heaps can also have these same favorable properties. Note that the authors of quickheaps excluded pointer-based operations, in particular DECREASE-KEY, from their consideration of quickheaps in external memory.



**Figure 1:** Illustration of our partition-based heap framework. The pivots  $p_2 \dots, p_\ell$  partition the  $n$  elements in the heap into sets  $\ell$  sets  $S_1, S_2, \dots, S_\ell$ . Black dots are the  $n$  elements in the heap. Pivots can be elements in the heap (black) or elements previously removed from the heap (white).

## 2. Partition-based Heap Framework

Here we describe the generic framework of partition-based heaps common to all three of our data structures in Sections 3–5. We assume that the elements have keys from a totally ordered universe and that elements are ordered with respect to their keys. If ever duplicate keys are inserted into a partition-based heap, standard techniques are used so that all ties are broken consistently.

Partition-based heaps maintain the elements currently stored in the heap as a partition of sets  $S_1, S_2, \dots, S_\ell$ , for some  $\ell$  that is  $\mathcal{O}(\lg n)$ . The number of elements currently stored in the heap is  $n = |S_1| + \dots + |S_\ell|$ . While the sets themselves are unordered, if  $i < j$ , all elements in  $S_i$  are less than all elements of  $S_j$ .

Each set  $S_i$  is stored as a *doubly-linked list* as constant-time concatenation of sets is vital to all of our structures. For each set, the size  $|S_i|$  is explicitly maintained. As we store all elements in linked lists, we support *stable pointers* to any element inserted.

In addition to the sets and their sizes, a partition-based heap stores *pivots*  $p_2 \leq \dots \leq p_\ell$  which are elements such that  $S_i \subseteq [p_i, p_{i+1})$  for  $i = 1, \dots, \ell$ ; for notational convenience, assume  $p_1$  and  $p_{\ell+1}$  are pivots with keys  $-\infty$  and  $+\infty$ , respectively. Pivots may be either elements currently or previously in the heap. See Figure 1 for an illustration.

To implement each priority queue operation all of our partition-based heaps have the following steps in common:

- **INSERT( $e$ ):** Search for  $e$  among the  $\mathcal{O}(\lg n)$  pivots in  $\mathcal{O}(\lg \lg n)$  time to find the set  $S_i$  with  $p_i \leq e < p_{i+1}$ . Insert  $e$  into  $S_i$  and increment the size  $|S_i|$ . Return a pointer  $ptr$  to the new node containing  $e$  in the linked list for  $S_i$ .
- **DELETE-MIN():** Remove the smallest element from the first nonempty set  $S_i$ , and decrement the size  $|S_i|$ .
- **DECREASE-KEY( $ptr, key$ ):** Let  $e$  denote the element in the linked-list node  $ptr$  points to. Remove this node from its linked list. Binary search on the pivots in  $\mathcal{O}(\lg \lg n)$  to find the set  $S_i$  that the node was removed from and decrement the size  $|S_i|$ . Decrease the key of  $e$  to  $key$ , and re-insert the node as in INSERT, using a second search among the pivots to discover which set to insert it into, insert it at the end of the linked list for the set, and increment the set size.

These descriptions are not enough to maintain the crucial invariant that  $\ell = \mathcal{O}(\log n)$ . This is where the three heaps in Sections 3–5 differ; each has its own invariants towards this goal, and each augments the basic implementations above with restructuring operations which maintain said invariants. All structures use the technique of combining sets in constant time by concatenating linked-lists and splitting sets in linear time using a linear time selection algorithm


to maintain their invariants. Further, all structures use amortized analysis to argue for the running times.

At a very high level, the philosophy of each structure is as follows: In *lazy partition heaps* (LP heaps) in Section 3, the sets are allowed to become arbitrarily large, and there is no lower bound on the number of sets. A bound on the size of a set compared to the sum of all previous sets is maintained. In *Fibonacci heaps: the next generation* (FH:TNG) in Section 4 and the heaps in Section 5, sets are allowed to be empty and there are strict upper bounds on their size. In FH:TNG, there are lower bounds on non-empty set sizes, where both the upper and lower bounds are given by the Fibonacci numbers, which naturally is amiable to concatenating and splitting while maintaining the size invariants. The heaps in Section 5 do not maintain a lower bound on the set sizes which results in simpler restructuring algorithms.

### 3. Lazy Partition Heaps

To keep the number of sets  $\ell$  in  $\mathcal{O}(\lg n)$  at all times, *lazy partition heaps* (LP heaps) use the following pivot-forgetting rule, which is adapted from lazy search trees [SW20]:

*A pivot  $p$  shall be abandoned if the two sets that it separates together contain fewer elements than there are smaller elements in the heap (Figure 2).*



**Figure 2:** The pivot-forgetting rule for sets in LP heaps. Here, the lists for sets  $S_{i-1}$  and  $S_i$  would be concatenated (thus forgetting  $p_i$ ) since there are more elements smaller than  $p_{i-1}$  than between  $p_{i-1}$  and  $p_{i+1}$ .

In LP heaps, every DELETE-MIN operation partitions  $S_1$  around its median. Since no stringent bound on  $|S_1|$  is guaranteed, we need to amortize the occasional high costs for DELETE-MIN. For that, sets  $S_j$ ,  $j = 1, \dots, \ell$ , accumulate potential when they have grown “too large”, namely  $\varphi_j := \max\{0, |S_j| - \mathbf{s}(S_j)\}$ , where  $\mathbf{s}(S_j) := |S_1| + \dots + |S_{j-1}|$  denotes the number of elements smaller than all elements in  $S_j$  (their “safety buffer” from the minimum).

This in particular includes the set  $S_1$ , whose potential  $\varphi_1 = |S_1|$  can thus pay for the scanning and partitioning of  $S_1$  upon the next DELETE-MIN.

Apart from the sets, we store a (static) sorted array  $T$  of the current sets  $S_1, S_2, \dots, S_\ell$  (in this order). (Whenever the collection of sets changes because of a DELETE-MIN, we can afford to rebuild  $T$  from scratch.) For simplicity of the pseudocode, we let each set  $S_i$  store the corresponding pivot  $p_i$  as  $S_i.p$ . We point out that  $\mathbf{s}(S_j)$  is only needed during the pivot-forgetting phase and is computed on demand there.

#### 3.1. Operations

Pseudo code for the operations on a LP heap is shown in Figure 3. Throughout the lifetime of an LP heap, we maintain the following invariant.

**Invariant (S):** The number of set is at most  $\ell \leq 2 \lg n + 1$ .

If the invariant is at risk of violation—which is only possible upon DELETE-MIN—we apply the following concatenation rule (cf. Figure 2). Recall that  $\mathbf{s}(S_j) := |S_1| + \dots + |S_{j-1}|$  is the number of *smaller elements* (smaller than  $p_j$ , that is):

**Concatenation Rule (C):** If  $|S_j| + |S_{j+1}| < \mathbf{s}(S_j)$ , concatenate  $S_j$  and  $S_{j+1}$ .

Concatenation here simply means we concatenate the linked lists of elements for the two sets and “forgets” the pivot formerly separating the two sets. This operation takes constant time. FORGET-PIVOTS computes  $\mathbf{s}(S_j)$  and applies rule (C) to  $\mathcal{S}$  successively. Clearly, FORGET-PIVOTS uses  $\mathcal{O}(\ell)$  time if the (initial) number of sets is  $\ell$ . Note that FORGET-PIVOTS also removes all empty sets present, thereby ensuring that  $|S_j| \geq 1$  for all  $1 \leq j \leq \ell$ . After FORGET-PIVOTS has concatenated all pairs of sets that satisfy (C), we obtain a bound on the number of sets from the following lemma.

**Lemma 3.1 ((C) implies (S)):** *Suppose after applying rule (C), we have sets  $S_1, \dots, S_\ell$ , with  $|S_j| \geq 1$  for all  $1 \leq j \leq \ell$ . Then  $\ell \leq 2 \lg n + 1$ , for  $n$  the total size of the sets.*  $\triangleleft$

*Proof:* We first show that prefixes of sets grow exponentially: for all  $1 \leq j \leq \ell$ , we have  $\Sigma(j) := |S_1| + \dots + |S_j| \geq 2^{(j-1)/2}$  (\*).

The proof is by induction on  $j$ . The claim holds for  $1 \leq j \leq 2$ , since  $|S_i| \geq 1$  for  $1 \leq i \leq \ell$ . In the inductive step for  $j \geq 3$ , the inductive hypothesis implies  $\Sigma(j-2) \geq 2^{(j-3)/2}$ . By concatenation rule (C), we have  $|S_{j-1}| + |S_j| \geq \mathbf{s}(S_{j-1})$ , and by definition  $\mathbf{s}(S_{j-1}) = \Sigma(j-2)$ . Therefore  $\Sigma(j) = \Sigma(j-2) + |S_{j-1}| + |S_j| \geq 2 \cdot \Sigma(j-2) \geq 2 \cdot 2^{(j-3)/2} = 2^{(j-1)/2}$ , which proves (\*).

Since (\*) holds for any prefix of sets, it does so in particular for  $j = \ell$ . Then,  $\Sigma(\ell) = n$  (all elements), and we get  $\ell \leq 2 \lg n + 1$  as claimed.  $\square$

It follows from Lemma 3.1 that FORGET-PIVOTS ensures (S). This completes the description of the data structure.

### 3.2. Potential

For analyzing the amortized cost of operations, we use the following potential:

$$\Phi := \sum_{j=1}^{\ell} \varphi_j \quad \text{with} \quad \varphi_j := \beta \cdot \max\{0, |S_j| - \mathbf{s}(S_j)\}.$$

(The sum is over all sets currently in the data structure.) Here  $\beta \geq 2$  is a fixed scaling parameter (chosen depending on the constant factor in the linear cost of partitioning).

### 3.3. Amortized Cost of Insert

Upon an INSERT operation, the set  $S_j$  containing  $e$  is located using binary search in  $T$  in  $\mathcal{O}(\lg \lg n)$  time. Adding  $e$  to the linked list of  $S_j$  takes constant time.

In  $\Phi$ ,  $\mathbf{s}(S_{j'})$  increases by 1 for all  $j' > j$ ; this never increases the value of  $\Phi$ . The size of  $S_j$  grows by 1 which increases  $\Phi$  by at most  $\beta = \mathcal{O}(1)$ . INSERT thus costs  $\mathcal{O}(\lg \lg n)$  (both amortized and worst case).

LP-HEAP-INSERT( $e$ )

```

1   $ptr := \text{NEW-NODE}(e)$  // obtain pointer to element
2   $S := T.\text{FIND}(e.\text{key})$  // binary search the pivots
3   $S.\text{APPEND}(ptr)$  // insert element
4  return  $ptr$ 

```

LP-HEAP-DECREASE-KEY( $ptr, key$ )

```

1   $e := *ptr$  // obtain element from pointer
2   $S := T.\text{FIND}(e.\text{key})$ 
3   $S.\text{REMOVE}(ptr)$  // remove element from current set
4   $e.\text{key} := key$  // update key
5   $S := T.\text{FIND}(e.\text{key})$ 
6   $S.\text{APPEND}(ptr)$  // insert element in new set

```

LP-HEAP-DELETE-MIN()

```

1   $min := S_1.\text{DELETE-MIN}()$  // scans  $S_1$ 
2   $p := \text{MEDIAN}(S_1)$  // (larger) median (deterministic linear-time selection [BFP+73])
3   $\bar{S}_1, \bar{S}_2 := \text{PARTITION}(S_1, p)$ 
4   $\bar{S}_2.p := p$ 
5   $\mathcal{S} := [\bar{S}_1, \bar{S}_2, S_2, \dots, S_\ell]$  // doubly-linked list of sets
6   $\text{FORGET-PIVOTS}(\mathcal{S})$  // modifies  $\mathcal{S}$ 
7  Rebuild  $T$  from  $\mathcal{S}$ 
8  return  $min$ 

```

FORGET-PIVOTS( $\mathcal{S}$ )

```

// Assume  $\mathcal{S} = [S_1, \dots, S_\ell]$ 
1   $\mathcal{S}.\text{REMOVE-ALL-EMPTY-SETS}()$ 
2   $A := S_1$ 
3   $s(A) := 0$ 
4  while  $\mathcal{S}.\text{HAS-NEXT}(A)$ 
5       $B := \mathcal{S}.\text{NEXT}(A)$ 
6      if  $|A| + |B| < s(A)$ 
7          Concatenate  $B$  into  $A$ 
8           $\mathcal{S}.\text{REMOVE}(B)$ 
9      else
10          $s(B) := s(A) + |A|$ 
11          $A := B$ 
12     end if
13 end while

```

**Figure 3:** LP heaps implementation.

### 3.4. Amortized Cost of Decrease-Key

When  $\text{DECREASE-KEY}(ptr, key)$  is performed, the element  $e$  at pointer  $ptr$  has its key decreased to  $key$ . Re-inserting  $e$  costs  $\mathcal{O}(\lg \lg n)$  (both amortized and worst case). This changes the size of at most two sets, with one increasing and the other decreasing, which increases the potential by at most  $\beta$ . It may also change  $\mathbf{s}(S_{j'})$  for some sets  $S_{j'}$ , but since we move  $e$  from  $S_i$  to an  $S_j$  with  $1 \leq j \leq i$  when we make the key of the element smaller,  $\mathbf{s}(S_j)$  can only *increase*, which does not increase the potential. In total,  $\text{DECREASE-KEY}$  uses  $\mathcal{O}(\lg \lg n + \beta) = \mathcal{O}(\lg \lg n)$  time (both amortized and worst case).

$\text{DECREASE-KEY}$  may leave a set  $S_i$  empty, where  $2 \leq i \leq \ell$ , however, these may be removed upon a later  $\text{FORGET-PIVOTS}$ , without an increase in time or potential, as the number of sets remains bounded, maintaining invariant (S).

### 3.5. Amortized Cost of Delete-Min


A  $\text{DELETE-MIN}$  scans  $S_1$  to find the minimum element  $e$  and removes it; moreover, we split  $S_1$  by partitioning it around its median, resulting in two new sets  $\bar{S}_1$  and  $\bar{S}_2$ . For the partition, we use deterministic selection. Partitioning  $S_1$  has overall actual cost of  $\mathcal{O}(|S_1|)$ .

Finally, we establish invariant (S) by calling  $\text{FORGET-PIVOTS}$ . By the invariant (S) on the number of sets, and since partitioning set  $S_1$  creates two new sets,  $\text{FORGET-PIVOTS}$  uses amortized  $\mathcal{O}(\lg n)$  actual time.

It remains to analyze the change in potential. Note first that the concatenation rule (C) is chosen precisely so that when concatenating  $S_{j+1}$  into  $S_j$ , the contributions  $\varphi_{j+1}$  and  $\varphi_j$ , both before and after the concatenation, are all zero, so concatenating does not change the potential at all. More specifically, assume that  $|S_j| + |S_{j+1}| < \mathbf{s}(S_j)$ , and a concatenation of  $S_j$  and  $S_{j+1}$  is performed to obtain the new set  $\bar{S}_j$ , by the concatenation rule (C). Then the potential is decreasing by  $\varphi_j$  and  $\varphi_{j+1}$  and increasing by the new potential  $\bar{\varphi}_j$ . By definition these potentials are

$$\begin{aligned} \varphi_j &= \beta \cdot \max\{0, |S_j| - \mathbf{s}(S_j)\} &= 0 \\ \varphi_{j+1} &= \beta \cdot \max\{0, |S_{j+1}| - (\mathbf{s}(S_j) + |S_j|)\} &= 0 \\ \bar{\varphi}_j &= \beta \cdot \max\{0, (|S_j| + |S_{j+1}|) - \mathbf{s}(S_j)\} &= 0. \end{aligned}$$

The change in potential  $\Delta\Phi$  is therefore 0 upon a concatenation.



**Figure 4:** Sketch of the sets of an LP heap and the partition of  $S_1$  into  $\bar{S}_1$  and  $\bar{S}_2$  in  $\text{DELETE-MIN}$ .

It remains to analyze the influence on the potential of partitioning the set  $S_1$  into the new sets  $\bar{S}_1$  and  $\bar{S}_2$ . The situation after partitioning  $S_1$  is illustrated in Figure 4. For the set  $S_1$ , we lose  $\varphi_1 = \beta|S_1|$  and gain  $\bar{\varphi}_1 + \bar{\varphi}_2 = \beta|\bar{S}_1|$  from  $\bar{S}_1$  and  $\bar{S}_2$ , as we choose the larger median s. t.  $\mathbf{s}(\bar{S}_2) = |\bar{S}_1| \geq |\bar{S}_2|$ , i.e.,  $\bar{\varphi}_2 = 0$ . All other sets lose one smaller element, so we gain  $\leq \ell \cdot \beta$  in potential. By construction,  $|\bar{S}_1| \leq \frac{1}{2}|S_1|$  (recall that we delete an element from  $S_1$  before partitioning into  $\bar{S}_1$  and  $\bar{S}_2$ ), so  $\Delta\Phi \leq -\frac{\beta}{2}|S_1| + \ell\beta$ . For a sufficiently large  $\beta \geq 2$ , the total amortized cost is thus

$$\mathcal{O}(\Delta\Phi + |S_1| + \lg n) = \mathcal{O}(\lg n).$$

### 3.6. Extensions

Here we discuss a few further properties of LP heaps; they are not needed to understand (or teach) LP heaps.

#### 3.6.1. Further Operations

We can support LP-HEAP-FIND-MIN in  $\mathcal{O}(1)$  time, by maintaining the minimum upon LP-HEAP-INSERT and LP-HEAP-DECREASE-KEY in constant time, and by scanning  $\overline{S}_1$  after a LP-HEAP-DELETE-MIN, which maintains the amortized time. Note that  $|S_1|$  can be arbitrary large, i.e., we cannot afford scanning  $S_1$  for each LP-HEAP-FIND-MIN operation.

We can also support an LP-HEAP-INCREASE-KEY operation or a general LP-HEAP-DELETE(*ptr*) operation, at total amortized cost  $\mathcal{O}(\beta \lg n + \lg \lg n) = \mathcal{O}(\lg n)$ . For both operations  $S_1$  might become empty, and further for general LP-HEAP-DELETE(*ptr*),  $n$  decreases which may violate invariant (S), and therefore FORGET-PIVOTS is needed to re-establish (S). This also adds  $\mathcal{O}(\lg n)$  to the cost.

An operation LP-HEAP-BUILD from a given collection of  $n$  elements can be supported in  $\mathcal{O}(n)$  time. For that, we simply declare all elements in  $S_1$ . The actual cost (if elements are given in a linked list) can even be constant, however we have to pay for the increase in potential of  $\beta n$ , so total amortized costs are  $\mathcal{O}(n)$ .

The LP-HEAP-MELD operation of taking the union of two heaps is not easily supported efficiently. One can trivially create a new LP heap from the concatenation of the two input heaps using LP-HEAP-BUILD, at linear amortized cost in the total number of elements.

#### 3.6.2. Partitioning

We can replace the deterministic linear-time selection [BFP<sup>+</sup>73] by randomized selection [Hoa61, FR75], leading to the same expected amortized time.

More importantly, we can also simply run a *single round* of (random) partitioning: Upon the partition, the potential decreases by  $\min\{|\overline{S}_1|, |\overline{S}_2|\}$ . By choosing the pivot not as the median, but at random between the elements of  $S_1$ , the decrease is therefore *expected* linear in  $|S_1|$ , and the amortized time of DELETE-MIN remains *expected*  $\mathcal{O}(\lg n)$ .


#### 3.6.3. LP Quickheaps

The *quickheaps layout* (described in more detail in Section 1.1) offers an enticing alternative to the linked lists for each set: At the expense of more data movement during insertions and deletions, we store all sets contiguously in a single array, with the pivots separating them. The only extra space is a static array  $T$  with the  $\mathcal{O}(\lg n)$  pivot positions.

Concatenations of adjacent sets can still be done in constant time (by simply demoting a pivot to an ordinary element), and partitioning can now be done in place, which improves practical overhead in running time. While we can also still binary search the pivots to *find* a set to insert an element into, insertions or deletions must now be implemented with up to one swap per set, leading to overall  $\Theta(\lg n)$  running time for all operations.

#### 3.6.4. Pointer-Machine Version

LP heaps already stay within the pointer-machine model of computation with bounded indegree nodes except for the index  $T$ . By replacing the sorted array and binary search by any balanced binary search tree, then LP heaps is entirely in the model.



**Figure 5:** Illustration of a next generation Fibonacci heap. While each set is stored as a linked list, this illustration uses colors to show the relationship between each list’s size, the Fibonacci numbers, and the size potential function.

## 4. Fibonacci Heaps: The Next Generation

The next generation Fibonacci heap (FH:TNG)<sup>2</sup> follows our framework for partition-based simple heaps with but a few distinctive features: Sets can either be present or absent. If a set is present, its size is required to be between the  $i$ th Fibonacci number, denoted as  $F_i$  ( $F_1 := F_2 := 1$ ;  $F_i := F_{i-1} + F_{i-2}$ ), and the  $i + 3$ rd,  $F_{i+3}$ , with the one exception that the first set,  $S_3$ , has no minimum size. Finally we require that there are at most eight empty sets in a row (*the consecutive empty invariant*) and at most two non-empty sets in a row (*the consecutive nonempty invariant*) to allow room for splits and concatenations.

The implementation of the operations follows our basic framework. The only difference is what to do when the invariants are violated. We restore them as follows using *invariant-restoring operations*: If a set  $S_i$  has  $F_{i+3}$  elements we call it *full* and there are two options: If the set below,  $S_{i+1}$  is empty,  $S_i$  is simply moved to  $S_{i+1}$ . We call this a *overflow-down operation*. If  $S_{i+1}$  is not empty,  $S_{i+2}$  must be empty by the consecutive nonempty invariant. We concatenate all  $F_{i+3}$  elements of  $S_i$  to the set  $S_{i+1}$  and then remove the  $F_{i+3}$  largest elements of the resultant set using linear-time selection and place them in  $S_{i+2}$ ; this leaves the size of  $S_{i+1}$  unchanged but the contents will certainly change. We call this an *overflow-thru operation*. If a set  $S_i$  has reached its minimum size  $F_i$ , it is *underfull*, and we execute either an *underflow-up operation* or an *underflow-thru operation* depending on whether  $S_{i-1}$  is empty. These are symmetric to the two overflow operations. If the consecutive-nonempty invariant is violated, there are three nonempty sets in a row with an empty list below. We concatenate the top two of them and place the resultant set in the empty slot below, we call this a *merge-down operation*.

If the consecutive-empty invariant is violated, there are nine empty sets in a row with a non-empty set below. We take the nonempty set below and split using linear-time selection into two set slots immediately above. We call this a *split-up*, and this takes linear time in the sets involved. Merge-down takes constant time as it is combining two linked lists, whereas the split-up takes linear time in the size of the sets to perform a selection and traverse the set to split it in two.

The asymmetry of the linear-time split-up with the constant time merge-down is crucial and is deeply connected to the fact that we allow keys to be decreased but not increased! The intuition is that repeated DECREASE-KEY operations can force inexpensive merge-downs to be performed as everyone is slowly moved to bigger sets. For example, this can happen by repeatedly DECREASE-KEY-ing the maximum item to make it the minimum. However, split-up operations are only triggered by the more expensive DELETE-MIN operations and thus their more expensive cost can be paid for.

<sup>2</sup>The name is meant to reflect that fact that in the original Fibonacci heaps and the the structure here, Fibonacci numbers play crucial, and entirely different, roles.

Here we present the potential function used in the analysis. The potential of the data structure is the sum of three components, each of which is defined for each set in the structure:

**Nonempty potential.** Define  $\varphi_i^{\text{nonempty}}$  to be 1 if  $S_i$  is nonempty and 0 if  $S_i$  is empty. This is used to pay for the merge-down operation where the number of nonempty sets decreases by one at unit cost.

**Size potential.** Gives potential as  $S_i$  nears being full ( $F_{i+3}$  elements) or underfull ( $F_i$  elements) and 0 potential when its size is in the middle of its allowable range (from  $F_{i+1}$  to  $F_{i+2}$  elements). This is the classic potential function used in array doubling and is used to pay for the underflow and overflow operations. Figure 5 is color-coded to highlight these three ranges.

$$\varphi_i^{\text{size}} := \begin{cases} 0 & \text{if } S_i \text{ is empty} \\ F_{i+1} - |S_i| & \text{if } F_i \leq |S_i| < F_{i+1} \text{ (green in the figure)} \\ 0 & \text{if } F_{i+1} \leq |S_i| \leq F_{i+2} \text{ (blue in the figure)} \\ |S_i| - F_{i+2} & \text{if } F_{i+2} < |S_i| \leq F_{i+3} \text{ (purple in the figure)} \end{cases}$$

**Up potential.** This gives potential to nonempty sets that have few items above. This is used to pay for the split-up operation. The intuition is that this is a mechanism that allows the DELETE-MIN operation to pay unit cost to each of  $\Theta(\lg n)$  sets, which the sets save up in the case most items above have been removed and a split-up needs to happen. Let  $\uparrow_i$  be defined to be  $\sum_{j=1}^{i-1} |S_j|$ . Define the up potential of set  $S_i$ ,  $\varphi_i^\uparrow$ , as

$$\varphi_i^\uparrow := \begin{cases} 0 & \text{if } i \leq 2 \text{ or } S_i \text{ is empty} \\ \max\{0, F_{i-3} - \uparrow_i\} & \text{otherwise} \end{cases}$$

The potential of the data structure,  $\Phi$ , is the sum of the up, size, and set potentials of every set; we refer to this sum as  $\varphi_i$ :  $\Phi := \sum_{i=0}^k \overbrace{(\varphi_i^\uparrow + \varphi_i^{\text{size}} + \varphi_i^{\text{nonempty}})}^{\varphi_i}$ .

**Facts.** We now present a series of facts which relating to these potentials.

**Fact 4.1:** If  $S_i$  is nonempty and at least of one of  $S_{i-1}$ ,  $S_{i-2}$   $S_{i-3}$  is nonempty then  $\varphi_i^\uparrow = 0$ .  $\triangleleft$

*Proof:* If any of  $S_{i-1}$ ,  $S_{i-2}$  and  $S_{i-3}$  are nonempty they will have least  $F_{i-3}$  items. Thus  $\uparrow_i \geq F_{i-3}$  and thus  $\max\{0, F_{i-3} - \uparrow_i\} = 0$ .  $\square$

**Fact 4.2:**  $\uparrow_i \leq F_{i+4} - 1$   $\triangleleft$

*Proof:*  $\uparrow_i = \sum_{j=1}^{i-1} |S_j| \leq \sum_{j=1}^{i-1} F_{i+3} = \sum_{j=4}^{i+2} F_i \leq F_{i+4} - 1$  via the classic identity that  $\sum_{i=1}^n F_i = F_{i+2} - 1$   $\square$

**Fact 4.3:** If an operation redistributes items among  $S_i, S_{i+1}, \dots, S_j$  without insertion or deletion,  $\varphi_\ell^\uparrow$  does not change outside of  $\ell \in [i, j]$ .  $\triangleleft$

*Proof:* Recall that  $\varphi_\ell^\uparrow$  is a function of whether  $S_\ell$  is empty as well as  $\uparrow_\ell$ , which is defined to be  $\sum_{m=1}^{\ell-1} |S_m|$  and does not change if the sets being redistributed are either entirely inside or entirely outside of this sum.  $\square$

**Fact 4.4:** For  $i = O(1)$ ,  $\varphi_i = O(1)$ . ◁

The next fact, follows from the intuition that set  $S_i$  can be split into  $S_{i-1}$  and  $S_{i-2}$  or an inverse concatenation performed without gaining size potential. This follows from the properties of Fibonacci numbers, where a set of a certain fullness can be split into two with the same fullness and thus the same potential.

**Fact 4.5:** The elements of a nonempty set  $S_i$  can be distributed to previously empty sets  $S_{i-1}$  and  $S_{i-2}$  so that the set potential remains unchanged. Inversely, the elements of nonempty sets  $S_{i-1}$  and  $S_{i-2}$  can be concatenated into the previously empty set  $S_i$  without increasing potential. ◁

*Proof:* In general, if  $S_i$  has between  $F_{i+j}$  and  $F_{i+j+1}$  elements, for  $j \in 0, 1, 2$  we can proportionally distribute them so that  $S_{i-1}$  has between  $F_{i+j-1}$  and  $F_{i+j}$  elements and  $S_{i-2}$  has between  $F_{i+j-2}$  and  $F_{i+j-1}$  elements; this follows naturally from the definition of Fibonacci numbers. In any such distribution, for the case where  $j = 0$  the size potential was and remains,  $F_{i-1} - |S_i|$ , for  $j = 1, 0$ , and for  $j = 2$ ,  $|S_i| - F_{i+2}$ .

For the inverse operation of merging, set  $j$  so that  $S_i$  will have between  $F_{i+j}$  and  $F_{i+j+1}$  elements. If  $j = 1$  then this holds trivially as the potential will become zero. For the cases where  $j = 0$  and  $j = 2$  a little math confirms the intuition:

$$\underbrace{\varphi_i^{\text{size}}}_{\text{after}} = |S_i| - F_{i+1} = \overbrace{|S_{i-1}| - F_i + |S_{i-2}| - F_{i-1}}^{\substack{|S_{i-1}| - F_i \leq \varphi_{i-1}^{\text{size}} \\ |S_{i-2}| - F_{i-1} \leq \varphi_{i-2}^{\text{size}}}} \leq \underbrace{\varphi_{i-1}^{\text{size}} + \varphi_{i-2}^{\text{size}}}_{\text{before}}$$

$$\underbrace{\varphi_i^{\text{size}}}_{\text{after}} = F_{i+3} - |S_i| = \overbrace{F_{i+2} - |S_{i-1}| + F_{i+1} - |S_{i-2}|}^{\substack{F_{i+2} - |S_{i-1}| \leq \varphi_{i-1}^{\text{size}} \\ F_{i+1} - |S_{i-2}| \leq \varphi_{i-2}^{\text{size}}}} \leq \underbrace{\varphi_{i-1}^{\text{size}} + \varphi_{i-2}^{\text{size}}}_{\text{before}} \quad \square$$

## 4.1. Split-up


See Figure 6. Split-up is called when the invariant that there are at most 8 consecutive empty sets in a row is violated, that is, there are 9 consecutive empty sets. Specifically, it is called on set  $i$  when  $S_i$  and  $S_{i-10}$  are nonempty and  $S_{i-1} \dots S_{i-9}$  are empty. This requires  $i \geq 11$ .

The moving of items from  $S_i$  to  $S_{i-1}$  and  $S_{i-2}$  should be done in a way proportional their sizes, that is in respect of the golden ratio. Very specifically, if there are between  $F_l$  and  $F_{l-1}$  elements in  $S_i$ , these should be moved so there are between  $F_{l-1}$  and  $F_{l-2}$  elements in  $S_{i-1}$  and  $F_{l-2}$  and  $F_{l-3}$  elements in  $S_{i-3}$ ; this is of course possible due to the definition of the Fibonacci numbers.

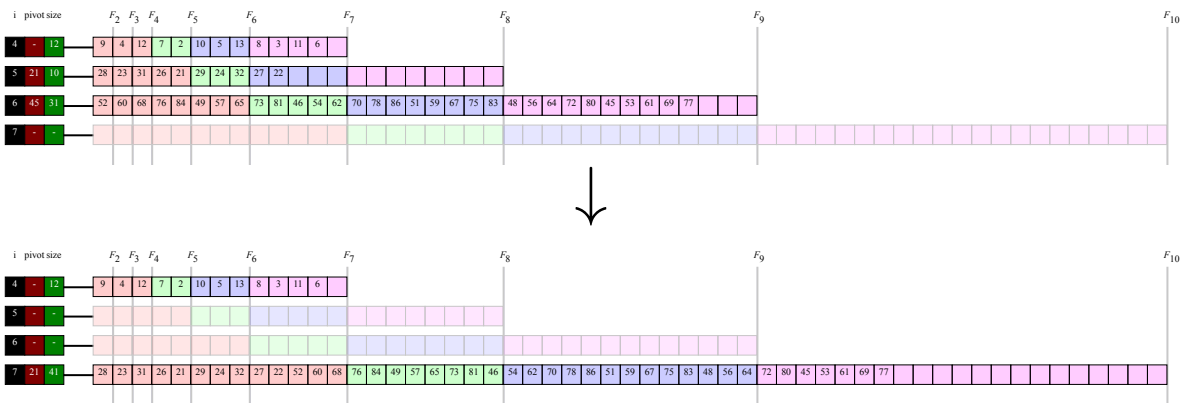
The new pivots are set to the minimum in each set via linear scan.

**Lemma 4.6:** *The amortized cost of the split-up operation is at most 0 when called on  $S_i$ ,  $i \geq 11$ .* ◁

*Proof:* The nominal cost of the split-up operation is defined to be  $F_{i-6}$  as it takes linear time in the sets involved. To compute the change in potential we look at each of the three components of the potential function.



**Figure 6:** Split-up operation. There are 9 consecutive empty sets in a row,  $S_2 \dots S_{10}$ . The set  $S_{11}$  is then proportionally split into sets  $S_9$  and  $S_{10}$ . In all cases the nonempty sets remain in the same color range, in this case green, which represents one of the three cases of the size potential. Empty cells continue off the page.



**Figure 7:** A merge-down operation performed when there are three non-empty sets in a row, here  $S_4, S_5, S_6$ , and the set below  $S_7$  is empty. sets  $S_5$  and  $S_6$  are concatenated into the empty slot of set  $S_7$ .

**Change in  $\varphi_i^\uparrow$ :** This is  $\sum_{i=0}^k \varphi_i^\uparrow - \sum_{i=0}^k \bar{\varphi}_i^\uparrow$ . However, this can be simplified using Fact 4.3 and recalling that empty sets have zero up potential to  $\bar{\varphi}_{i-1}^\uparrow + \bar{\varphi}_{i-2}^\uparrow - \varphi_i^\uparrow$ . We can now bound each of these quantities separately. First,  $\bar{\varphi}_{i-1}^\uparrow = 0$  by Fact 4.1. Next, for  $\bar{\varphi}_{i-2}^\uparrow$ , this is  $\max\{0, F_{i-5} - \bar{\uparrow}_{i-2}\} \leq F_{i-5}$ .

Finally, for  $-\varphi_i^\uparrow$ : we use the fact that  $S_{i-1} \dots S_{i-9}$  are empty to bound  $\uparrow_i$ :

$$\begin{aligned}
-\varphi_i^\uparrow &= -\max\{0, F_{i-3} - \uparrow_i\} && \text{By definition} \\
&= -\max\{0, F_{i-3} - \uparrow_{i-9}\} && \uparrow_i = \uparrow_{i-9} \\
&< -\max\{0, F_{i-3} - F_{i-5} - 1\} \uparrow_{i-9} && < F_{i-5} - 1 \text{ by Fact 4.2} \\
&= -F_{i-4} + 1 && i \geq 11
\end{aligned}$$

Thus, bringing it together gives:

$$\sum_{i=0}^k \varphi_i^\uparrow - \sum_{i=0}^k \bar{\varphi}_i^\uparrow = \bar{\varphi}_{i-1}^\uparrow + \bar{\varphi}_{i-2}^\uparrow - \varphi_i^\uparrow \leq 0 + F_{i-5} - F_{i-4} + 1 = -F_{i-6} + 1$$

**Change in  $\varphi_i^{\text{size}}$ :** 0.


**Change in  $\varphi_i^{\text{nonempty}}$ :** +1, as there is now one more nonempty set.

In summary, by definition, the amortized cost is the nominal cost plus the change in potential which is at most, when  $i \geq 11$ :  $F_{i-6} - F_{i-6} - 1 + 0 + 1 = 0$   $\square$

## 4.2. Merge-down

See Figure 7. The merge-down operation is called when the invariant that there are most two non-empty consecutive sets is violated. That is, there are three non-empty consecutive sets  $S_i, S_{i-1}$  and  $S_{i-2}$  with  $S_{i+1}$  being empty (or non-existent). The sets  $S_i$  and  $S_{i-1}$  are concatenated to form the new set  $\bar{S}_{i+1}$ ; this is where we use fact that the sets are stored as doubly-linked sets and not as arrays so that this operation takes constant time. The pivot  $\bar{p}_{i-2}$  is set to  $p_{i-1}$ .

**Lemma 4.7:** *The amortized cost of the merge-down operation is at most 0.*  $\triangleleft$



**Figure 8:** An overflow-down operation performed on  $S_5$  which has  $F_8 = 21$  elements when  $S_6$  is empty.

*Proof:* The nominal cost of the operation is 1, To compute the change in potential we look at each of the three components of the potential function:

**Change in  $\varphi_i^\uparrow$ :** 0. By Fact 4.3 this is  $\bar{\varphi}_{i+1}^\uparrow - \varphi_i^\uparrow - \varphi_{i-1}^\uparrow$ . By Fact 4.1, since  $S_{i-2} = \bar{S}_{i-2}$  is nonempty, all three of these are zero.

**Change of  $\varphi_i^{\text{size}}$ :**  $\leq 0$  by Fact 4.5.

**Change of  $\varphi_i^{\text{nonempty}}$ :**  $-1$  as the number of non-empty sets decreases by one.

In summary, by definition, the amortized cost is the nominal cost plus the change in potential:  
 $1 + 0 + 0 - 1 = 0$  □

### 4.3. Overflow-down

See Figure 8. The overflow-down operation is called when a set  $S_i$  is full, and thus has  $F_{i+3}$  items, and the set below,  $S_{i+1}$ , is empty. The set is moved from  $S_i$  to  $S_{i+1}$  in constant time.

**Lemma 4.8:** *The amortized cost of the overflow-down on set  $S_i$  is at most 0 when  $i \geq 3$*  ◁


*Proof:* The nominal cost of overflow up is 1, since it is a constant-time operation. To compute the change in potential we look at each of the three components of the potential function:

**Change in  $\varphi_i^\uparrow$ :** By Fact 4.3 this is  $\bar{\varphi}_{i+1}^\uparrow - \varphi_i^\uparrow$  which is at most  $\bar{\varphi}_{i+1}^\uparrow$ . Using the definition of the up potential,  $\bar{\varphi}_{i+1}^\uparrow = \max\{0, F_{i-2} - \uparrow_{i+1}\} \leq F_{i-2}$ .

**Change of  $\varphi^{\text{size}}$ :** As  $|\bar{S}_{i+1}| = F_{i+3}$ ,  $\bar{\varphi}_{i+1}^{\text{size}} = 0$ . Thus we need only consider  $-\varphi_i^{\text{size}}$ :  $-\left(\overbrace{|S_i|}^{F_{i+3}} - F_{i+2}\right) = F_{i+2} - F_{i+3} = -F_{i+1}$

**Change of  $\varphi_i^{\text{nonempty}}$ :** 0 as the number of nonempty sets does not change.

In summary, by definition, the amortized cost is the nominal cost plus the change in potential  
 is  $1 + \underbrace{F_{i-2} - F_{i+1}}_{\leq -F_i} + 0 \leq 0$ . □



**Figure 9:** An overflow-thru operation performed on  $S_5$  which has  $F_8 = 21$  elements when  $S_6$  is nonempty.

#### 4.4. Overflow-thru

See Figure 9. The overflow-thru operation is called when a set  $S_i$  is full, and thus has  $F_{i+3}$  items, and the set below,  $S_{i+1}$  is nonempty. As there can not be more than two non-empty consecutive sets by the consecutive nonempty invariant, this implies  $S_{i+2}$  is empty.

Recall that in this operation,  $S_i$  and  $S_{i+1}$  are combined, and the largest items are moved from the combined set to  $\bar{S}_{i+2}$  so that the size of  $\bar{S}_{i+1}$  is unchanged and thus  $|\bar{S}_{i+2}| = |S_i| = F_{i+3}$ . This movement requires a linear time median/order statistic algorithm and thus has nominal cost  $F_{i-4}$  since  $S_{i+2}$  is the largest set involved. The pivot  $\bar{p}_{i+1}$  and  $\bar{p}_{i+2}$  are set to the minimum in each set.

**Lemma 4.9:** *The amortized cost of overflow-thru is  $\leq 0$ , when  $i \geq 5$ .*  $\triangleleft$

*Proof:* The nominal cost of the operation is  $F_{i-4}$  as it takes linear time in the size of the sets involved. To compute the change in potential we look at each of the three components of the potential function:

**Change in  $\varphi^\uparrow$ :** By Fact 4.3 this is  $\bar{\varphi}_{i+2}^\uparrow + \bar{\varphi}_{i+1}^\uparrow - \varphi_{i+1}^\uparrow - \varphi_i^\uparrow$ . By Fact 4.1,  $\varphi_{i+1}^\uparrow$  and  $\bar{\varphi}_{i+2}^\uparrow$  are both zero. Combined with the fact that  $\varphi_{i+1}^\uparrow$  is nonnegative, the change in  $\varphi^\uparrow$  is at most  $\bar{\varphi}_{i+1}^\uparrow = \max\{0, F_{i-2} - \uparrow_{i+1}\} \leq F_{i-2}$ .

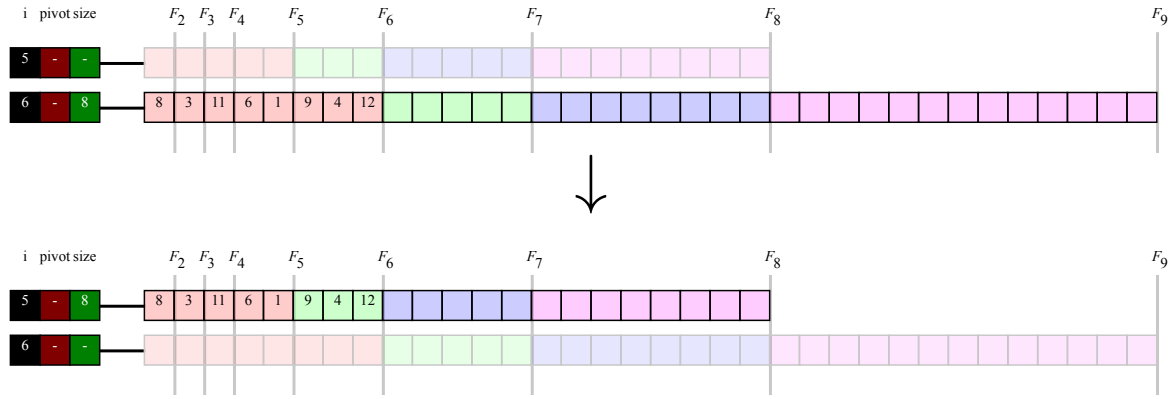
**Change in  $\varphi^{\text{size}}$ :** This is  $\bar{\varphi}_{i+2}^{\text{size}} + \bar{\varphi}_{i+1}^{\text{size}} - \varphi_{i+1}^{\text{size}} - \varphi_i^{\text{size}}$ . As  $|S_{i+2}| = F_{i+3}$ ,  $\bar{\varphi}_{i+2}^\uparrow = 0$  by definition. Since  $\varphi_{i+1}^{\text{size}} = \bar{\varphi}_{i+1}^{\text{size}}$ , this means the change in  $\varphi^{\text{size}}$  is  $-\varphi_i^{\text{size}}$ . By definition  $-\varphi_i^{\text{size}} = -(|S_i| - F_{i+2})$ , which since  $|S_i| = F_{i+3}$  is  $-F_{i+1}$ .

**Change in  $\varphi^{\text{nonempty}}$ :** 0, since the total number of nonempty sets does not change.

In summary, by definition, the amortized cost is the nominal cost plus the change in potential:  $F_{i-4} + F_{i-2} - F_{i+1} + 0 \leq 0$   $\square$

#### 4.5. Underflow-up

See Figure 10. The underflow-up operation is called when a set  $S_i$  is underfull, and thus has  $F_i$  items, and the set above,  $S_{i-1}$ , is empty. In this case the set and the pivot is moved from  $S_i$  to  $S_{i-1}$  in constant time.



**Figure 10:** An underflow-up operation performed on  $S_6$  which has  $F_6 = 8$  elements when  $S_5$  is empty.

**Lemma 4.10:** *The amortized cost of the underflow-up operation on  $S_i$  is at most 0 when  $i \geq 6$ .*  $\triangleleft$

*Proof:* The nominal cost of the operation is 1, as it is a constant-time operation. To compute the change in potential we look at each of the three components of the potential function:

**Change in  $\varphi^\uparrow$ :** By Fact 4.3 this is  $\bar{\varphi}_{i-1}^\uparrow - \varphi_i^\uparrow$ . We have  $\bar{\varphi}_{i-1}^\uparrow - \varphi_i^\uparrow \leq \bar{\varphi}_{i-1}^\uparrow \leq F_{i-4}$  since  $\bar{\varphi}_{i-1}^\uparrow = \max\{0, F_{i-4} - \bar{\tau}_{i-1}\}$  by definition.

**Change in  $\varphi^{\text{size}}$ :** This is  $\bar{\varphi}_{i-1}^{\text{size}} - \varphi_i^{\text{size}}$ . As  $|\bar{S}_{i-1}| = F_i$ ,  $\bar{\varphi}_{i-1}^{\text{size}} = 0$  by definition. We can now

$$\text{bound } -\varphi_i^{\text{size}}: -(F_{i+1} - \overbrace{|S_i|}^{F_i}) \leq -F_{i+1} + F_i = -F_{i-1}.$$

**Change in  $\varphi^{\text{nonempty}}$ :** 0. The number of nonempty sets does not change.  $\square$

In summary, by definition, the amortized cost is the nominal cost plus the change in potential:  $1 + F_{i-4} - F_{i-1} + 0 < 0$  as  $i \geq 6$ .

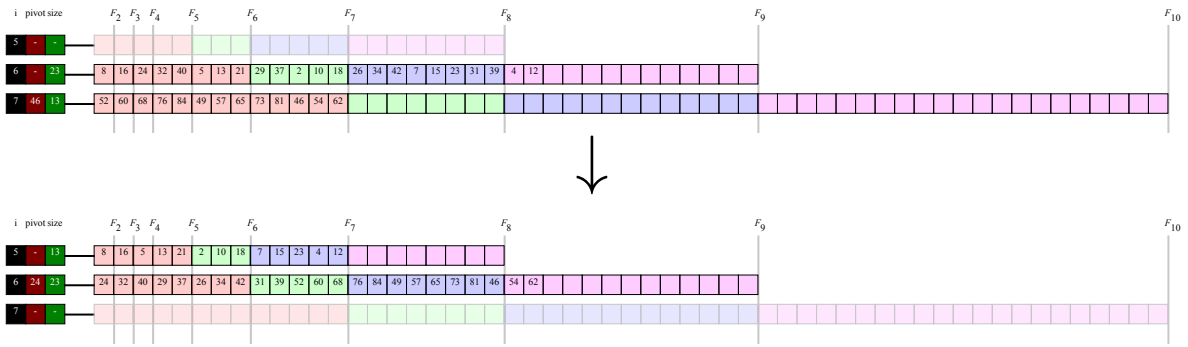
#### 4.6. Underflow-thru

See Figure 11. The underflow-thru operation is called when a set  $S_i$  is underfull, and thus has  $F_i$  items, and the set above,  $S_{i-1}$  is nonempty. As there can not be more than two non-empty consecutive sets by the consecutive nonempty invariant, this implies  $S_{i-2}$  is nonempty.

Recall that in this operation,  $S_i$  and  $S_{i-1}$  are combined, and the smallest items are moved from the combined set to  $\bar{S}_{i-2}$  so that the size of  $S_{i-1}$  is unchanged, thus  $|\bar{S}_{i-1}| = |S_{i-1}|$  and  $|\bar{S}_{i-2}| = |S_i| = F_i$ . This requires a linear time median/order statistic algorithm and thus has nominal cost  $F_{i-6}$  since  $S_i$  is the largest set involved. The pivots  $\bar{p}_{i-1}$  and  $\bar{p}_{i-2}$  are set to the minimum in each set.

**Lemma 4.11:** *The amortized cost of the underflow-thru operation is at most 0 when  $i \geq 3$ .*  $\triangleleft$

*Proof:* The nominal cost of the operation is  $F_{i-6}$ . To compute the change in potential we look at each of the three components of the potential function:



**Figure 11:** An underflow-thru operation performed on  $S_7$  which has  $F_7 = 13$  elements when  $S_5$  is nonempty.

**Change in  $\varphi^\uparrow$ :** By Fact 4.3 this is  $\overline{\varphi}_{i-1}^\uparrow + \overline{\varphi}_{i-2}^\uparrow - \varphi_i^\uparrow - \varphi_{i-1}^\uparrow$ . By Fact 4.1,  $\varphi_i^\uparrow$  and  $\overline{\varphi}_{i-1}^\uparrow$  are both zero. Also,  $-\varphi_{i-1}^\uparrow$  is at most 0. Thus the change in  $\varphi^\uparrow$  is at most  $\overline{\varphi}_{i-2}^\uparrow$  which is  $\max\{0, F_{i-5} - \overline{\tau}_{i-2}\} \leq F_{i-5}$ .

**Change in  $\varphi^{\text{size}}$ :** This is  $\overline{\varphi}_{i-2}^{\text{size}} + \overline{\varphi}_{i-1}^{\text{size}} - \varphi_{i-1}^{\text{size}} - \varphi_i^{\text{size}}$ . Note that while the contents of  $S_{i-1}$  will change, the size will not and thus  $\varphi_{i-1}^{\text{size}} = \overline{\varphi}_{i-1}^{\text{size}}$ . Also note that by definition, the size of  $\overline{S}_{i-2}$  is  $F_i$  and thus by definition  $\overline{\varphi}_{i-2}^{\text{size}} = 0$ . So, the change in  $\varphi^{\text{size}}$  is  $-\varphi_i^{\text{size}}$ :

$$-\varphi_i^{\text{size}} - (F_{i+1} - \overbrace{|S_i|}^{F_i}) \leq -(F_{i+1} - F_i) = -F_{i-1}$$

**Change in  $\varphi^{\text{nonempty}}$ :** **0** The number of nonempty sets does not change.

In summary, by definition, the amortized cost is the nominal cost plus the change in potential:  $F_{i-6} + F_{i-5} - F_{i-1} + 0 < 0$ .  $\square$

## 4.7. Insert

In an INSERT operation, a binary search is performed on the pivots to determine which set the newly inserted item is added to. If the item is smaller than the smallest pivot, it is added to the first set,  $S_3$ . Denote this set as  $j$ . The size of  $j$  is updated and the pivot of  $j$  is not updated as it is still valid. If this becomes overfull, a single overflow-down or overflow-thru is performed which eliminates the invariant violation on the set size but may cause too many empty sets in a row below  $j$  and too many nonempty sets in a row above  $j$ . The former is fixed via split-up, and the latter via merge-down, which can cascade downwards.

**Lemma 4.12:** *The amortized cost of the INSERT operation is  $O(\lg \lg n)$ .*  $\triangleleft$

*Proof:* We need not analyze any invariant-restoring operations as these have been shown to have amortized cost at most zero (except for those called on constant sized sets which can happen at most a constant number of times with a constant amortized cost).

The nominal cost of this operation is  $O(\lg \lg n)$  due to the binary search on the pivots of the  $\Theta(\lg n)$  sets.

To compute the change in potential we look at each of the three components of the potential function.

**Change in  $\varphi_i^\uparrow$ :** As this is defined as  $\max\{0, F_{i-3} - \uparrow_i\}$ , and the  $\uparrow_i$  are only incremented if they change, this potential increases by at most 0.

**Change in  $\varphi_i^{\text{size}}$ :**  $\leq 1$ . The set  $j$  now has one more element and thus could gain one unit of size potential.

**Change in  $\varphi_i^{\text{nonempty}}$ :** 0, as the number of sets do not change.

In summary, by definition, the amortized cost excluding any invariant-restoring operations is the nominal cost plus the change in potential which is at most:  $O(\lg \lg n) + 0 + 1 + 0 = O(\lg \lg n)$ . Including any calls to invariant-restoring operations, which have a total of  $O(1)$  amortized cost, gives the total amortized cost of  $O(\lg \lg n)$ .  $\square$

#### 4.8. Delete-Min

In the DELETE-MIN operation, we find and remove the smallest element from the smallest non-empty set, which we denote as  $S_j$ . Due to the consecutive empty invariant,  $j \leq 11$  and thus this takes constant time as  $|S_j| \leq F_{14} = 377$ . We update the size of the set  $j$ . If the set becomes underfull, we perform a single underflow-up or underflow-thru, which can cause a cascading chain of merge-downs to be performed.

**Lemma 4.13:** *The amortized cost of the DELETE-MIN operation is  $O(\lg n)$ .*  $\triangleleft$

*Proof:* We need not analyze any invariant-restoring suboperations as these have been shown to have amortized cost at most zero (except for those called on constant sized sets which can happen at most a constant number of times with a constant amortized cost).

The actual cost of this operation excluding the invariant-restoring sub-operations is  $O(1)$ .

To compute the change in potential we look at each of the three components of the potential function.

**Change in  $\varphi_i^\uparrow$ :** As this is defined as  $\max\{0, F_{i-3} - \uparrow_i\}$  for the nonempty sets (and 0 for the empty sets), and each of the  $O(\lg n)$   $\uparrow_i$  will decrease by at most 1, the up potential increases by at most  $O(\lg n)$ .

**Change in  $\varphi_i^{\text{size}}$ :**  $\leq 1$ . The set  $j$  now has one less element and thus could gain one unit of size potential.

**Change in  $\varphi_i^{\text{nonempty}}$ :** At most 0.

In summary, by definition, the amortized cost is the actual cost plus the change in potential which is:  $O(1) + O(\lg n) + 1 + 0 = O(\lg n)$ . Including any calls to invariant-restoring operations, which have a total of  $O(1)$  amortized cost, gives the total amortized cost of  $O(\lg n)$ .  $\square$

#### 4.9. Decrease-Key

This operation is carried out by first removing the item from its current set,  $S_j$ , decreasing it, and re-inserting it in the appropriate set  $S_k$ ; crucially  $k \leq j$  as the key is only decreased. Two binary searches are needed on the previous value (to update the size of the set it was in) and the new set (to insert and also update the size). The pivots of both sets remain valid with no change needed. After the removal, if the set is underfull a single underflow-up or underflow-thru is performed, and then split-up and merge-downs are performed as needed until all invariants are restored; after the insertion if the set is overfull a single overflow-down or overflow-thru is performed, and then split-up and merge-down are performed as needed until all invariants are restored.

**Lemma 4.14:** *The amortized cost of DECREASE-KEY is  $\mathcal{O}(\lg \lg n)$ .* ◁

*Proof:* We need not analyze any invariant-restoring suboperations as these have been shown to have amortized cost at most zero (except for those called on constant sized sets which can happen at most a constant number of times with a constant amortized cost).

The nominal cost of this operation is  $\mathcal{O}(\lg \lg n)$  due to the two binary search operations.

To compute the change in potential we look at each of the three components of the potential function.

**Change in  $\varphi_i^\uparrow$ :** This is defined as  $\max\{0, F_{i-3} - \uparrow_i\}$  for the nonempty sets, and 0 for the empty sets. For  $i < k$  and  $i > j$  this will not change, but for  $i$  between  $k$  and  $j$ , that is for sets between the removal and re-insertion,  $\uparrow_i$  will increase by 1, and thus  $\varphi_i^\uparrow$  will decrease by 1. This is where we crucially use the fact that keys can only decrease, as if we allowed an increase the potential here could grow by  $\Theta(\lg n)$ , but instead it may only shrink.

**Change in  $\varphi_i^{\text{size}}$ :**  $\leq 2$ . The insertion and deletion could each cause a unit change in the size potential of the affected sets.

**Change in  $\varphi_i^{\text{nonempty}}$ :** 0, as the number of sets do not change.

In summary, by definition, the amortized cost is the actual cost plus the change in potential which is at most:  $\mathcal{O}(\lg \lg n) + 0 + 2 + 0 = \mathcal{O}(\lg \lg n)$ . Including any calls to invariant-restoring operations, which have a total of  $\mathcal{O}(1)$  amortized cost, gives the total amortized cost of  $\mathcal{O}(\lg \lg n)$ . □

## 5. Heaps with Exponential Upper Bounded Sets

In this section we consider another variant of a partition-based simple heap consisting of sets  $S_1, \dots, S_\ell$  and pivots  $p_2 \leq \dots \leq p_\ell$ . Again, the time obtained for DELETE-MIN is amortized  $\mathcal{O}(\lg n)$  and for INSERT and DECREASE-KEY amortized  $\mathcal{O}(\lg \lg n)$ . In this variant we have exponential increasing upper bounds on the sizes of the sets, without any lower bounds, and allowing an arbitrary number of the sets to be empty. This is in contrast to the stronger structural invariants of the FH:TNG in Section 4.

**Invariant:**  $|S_i| < 3 \cdot 2^i$  for  $1 \leq i \leq \ell$ .

The heap operations are implemented as follows:

- INSERT( $e$ ):  
Binary search over the pivots to find  $i$  where  $p_i \leq e < p_{i+1}$ , and append  $e$  to  $S_i$ . If  $|S_i| = 3 \cdot 2^i$ , perform a recursive *push* of  $S_i$  to  $S_{i+1}$  (see below).
- DELETE-MIN():  
If  $S_1$  is empty, first recursively *pull* elements into  $S_1$  from  $S_2$  (see below). Delete and return the minimum from  $S_1$ . If  $\ell > 1 + \lg n$ , let  $S_{\ell-1}$  be the concatenation of  $S_{\ell-1}$  and  $S_\ell$ , and discard  $S_\ell$  and  $p_\ell$  before returning, i.e.,  $\ell$  decreases by one (this ensures  $\ell \leq 1 + \lg n$ ).
- DECREASE-KEY( $ptr$ ,  $key$ ):  
Delete the element  $e$  pointed to by  $ptr$  from its current set  $S_j$  (to decrement the size  $|S_j|$  we find  $j$  by a binary search over the pivots). Change the key of  $e$  to  $key$ , and perform a binary search over the pivots to find the set  $S_i$  where to add  $e$ , and append  $e$  to  $S_i$ . If  $|S_i| = 3 \cdot 2^i$ , perform a recursive *push* of  $S_i$  to  $S_{i+1}$ .

**Push.** For INSERT and DECREASE-KEY, the set  $S_i$  receiving one more element might get size  $|S_i| = 3 \cdot 2^i$ . We *push* all  $3 \cdot 2^i$  elements of  $S_i$  to  $S_{i+1}$ , and set  $S_i = \emptyset$ . In general, a recursive push to  $S_j$  of a set of elements  $X$  always satisfies  $2^{j-1} \leq |X| \leq 3 \cdot 2^{j-1}$ . If  $|S_j| < 2^j$  before the push (i.e.,  $S_j$  is too small to be pushed to  $S_{j+1}$ ),  $X$  is concatenated to  $S_j$  in constant time, such that  $|S_j| < 5 \cdot 2^{j-1}$ , and the push terminates. Otherwise, the set  $X$  becomes  $S_j$  and the old  $S_j$  is recursively pushed to  $S_{j+1}$ . The pivots need to be updated accordingly, and when  $j = \ell + 1$  a new set  $S_{\ell+1} = X$  is created, i.e.,  $\ell$  increases by one.


**Pull.** The recursive *pull* of elements into  $S_1$ , or more generally into  $S_i$  when  $S_i$  is empty, is performed as follows: If  $S_{i+1}$  is empty, we first recursively pull elements into  $S_{i+1}$ . If  $|S_{i+1}| \leq 2^{i-1}$ , we swap the roles of  $S_{i+1}$  and the empty  $S_i$ , i.e., in constant time move all elements from  $S_{i+1}$  to  $S_i$ . Otherwise, we select and delete the  $2^{i-1}$  smallest elements from  $S_{i+1}$  in time  $\mathcal{O}(2^i)$  and let  $S_i$  be these (and update  $p_{i+1}$ ; here it is crucial that elements are distinct to ensure  $\max S_i < p_{i+1} = \min S_{i+1}$ ).

### 5.1. Analysis

To prove the amortized performance of this structure we apply the following potential function  $\Phi = \varphi^{\text{insert}} + \varphi^{\text{push}} + \varphi^{\text{pull}}$ , where

$$\begin{aligned} \varphi^{\text{insert}} &:= \sum_{i=1}^{\ell} \varphi_i^{\text{insert}} & \varphi^{\text{push}} &:= \sum_{i=1}^{\ell} \varphi_i^{\text{push}} & \varphi^{\text{pull}} &:= \sum_{i=1}^{\ell} \varphi_i^{\text{pull}} \\ \varphi_i^{\text{insert}} &:= \max\{0, |S_i| - 5 \cdot 2^{i-1}\} \in [0, 2^{i-1}] & \varphi_i^{\text{push}} &:= \lfloor |S_i|/2^i \rfloor \in [0, 3] \\ \varphi_i^{\text{pull}} &:= \max\left\{0, 2^{i-1} - \sum_{j=1}^i |S_j|\right\} \in [0, 2^{i-1}] \end{aligned}$$

The basic idea is to let each set  $S_i$  have a potential. A big set will have a push potential  $\varphi_i^{\text{push}}$  that releases one unit of potential when the set is moved to  $S_{i+1}$ . An insertion into  $S_i$  can trigger a push of  $S_i$ . The saved up insertion potential  $\varphi_i^{\text{insert}}$  will be released by the push of  $S_i$  to cover the increase in pull potential  $\varphi_i^{\text{pull}}$ . Finally, pull potential is intuitively associated with each “empty slot” among the first  $2^{i-1}$  “slots” in  $S_i$ , with a potential to cover the cost to be recursively filled from each of the subsequent sets  $S_{i+1}, \dots, S_\ell$ . Figure 12 illustrates the relation between set sizes and the different potentials.



**Figure 12:** Illustration of set sizes and potentials for the  $S_i$  sets.

**Lemma 5.1:** For a heap with  $n$  elements  $\ell \leq 1 + \lg n$ . ◁

*Proof:*  $S_\ell$  is created when at least  $2^{\ell-1}$  elements are pushed from  $S_{\ell-1}$  to  $S_\ell$  during INSERT, i.e.,  $n \geq 2^{\ell-1}$  implying  $\ell \leq 1 + \lg n$ . During DELETE-MIN  $n$  decreases by one. The concatenation of  $S_{\ell-1}$  and  $S_\ell$  when  $\ell > 1 + \lg n$ , ensures  $\ell$  decreases by one and  $\ell \leq 1 + \lg n$  after DELETE-MIN.  $\square$

**Lemma 5.2:** An INSERT operation, excluding recursive pushing, takes worst-case  $\mathcal{O}(\lg \ell)$  time and the change in potential is  $\Delta\Phi \leq 2$ .  $\triangleleft$

*Proof:* The binary search for  $i$  and appending to  $S_i$  takes worst-case  $\mathcal{O}(\lg \ell)$  time. Increasing  $|S_i|$  by one can at most increase  $\varphi_i^{\text{insert}}$  and  $\varphi_i^{\text{push}}$  by one, and decrease  $\varphi_j^{\text{pull}}$  for each  $j \geq i$ . It follows that  $\Delta\Phi \leq 2$ .  $\square$

**Lemma 5.3:** A DECREASE-KEY operation, excluding recursive pushing, takes worst-case  $\mathcal{O}(\lg \ell)$  time and the change in potential is  $\Delta\Phi \leq 2$ .  $\triangleleft$

*Proof:* Performing DECREASE-KEY involves two binary searches on the  $\ell - 1$  pivots, requiring worst-case  $\mathcal{O}(\lg \ell)$  time, and worst-case  $\mathcal{O}(1)$  time for moving an element from  $S_i$  to  $S_j$ , where  $1 \leq j \leq i \leq \ell$ . The potentials  $\varphi_m^{\text{pull}}$  can only change for  $j \leq m < i$ , in which case they decrease by one,  $\varphi_i^{\text{insert}}$  and  $\varphi_i^{\text{push}}$  can only decrease, and  $\varphi_j^{\text{insert}}$  and  $\varphi_j^{\text{push}}$  at most increase by one. It follows that  $\Delta\Phi \leq 2$ .  $\square$

**Lemma 5.4:** The recursive pushing of  $S_i$  into  $S_{i+1}$ , when  $|S_i| = 3 \cdot 2^i$ , takes worst-case  $\mathcal{O}(m - i)$  time and the change in potential is  $\Delta\Phi \leq -m + i + 3$ , for some  $m$  where  $i < m \leq \ell + 1$ .  $\triangleleft$

*Proof:* Assuming the recursive pushing continues until pushing elements into  $S_m$  for  $i < m \leq \ell + 1$ , i.e.,  $|S_j| \geq 2^j$  for  $i < j < m$  and  $|S_m| < 2^m$  before the push. Each of the  $m - i$  pushes takes worst-case constant time to either rename sets or to concatenate two doubly-linked lists.

In the following, recall a bar over a value denotes the value after an operation. We first consider the change to the potential  $\varphi^{\text{insert}}$ . Before the push, we have  $|S_i| = 6 \cdot 2^{i-1}$ , i.e.,  $\varphi_i^{\text{insert}} = 2^{i-1}$ . Since after the push  $|\bar{S}_i| = 0$ , we have  $\bar{\varphi}_i^{\text{insert}} = 0$ . For  $i + 1 \leq j \leq m$ , between  $2^{j-1}$  and  $3 \cdot 2^{j-1}$  elements are pushed to  $S_j$ . It follows that after the recursive push  $|S_j| \leq 5 \cdot 2^{j-1}$ , for  $i \leq j \leq m$ , i.e.,  $\bar{\varphi}_i^{\text{insert}} = \dots = \bar{\varphi}_m^{\text{insert}} = 0$ . It follows that  $\Delta\varphi^{\text{insert}} \leq -2^{i-1}$ .

For  $\varphi^{\text{pull}}$ , note that  $|S_j|$  only changes for  $i \leq j \leq m$  and only these  $\bar{\varphi}_j^{\text{pull}}$  can change. For  $i < j < m$ , we have  $|S_j| \geq 2^{j-1}$  before and after the push, i.e.,  $\varphi_j^{\text{pull}} = \bar{\varphi}_j^{\text{pull}} = 0$ ,  $|S_m|$  increases, i.e.,  $\bar{\varphi}_m^{\text{pull}}$  can only decrease, and the maximal increase in  $\varphi_i^{\text{pull}}$  is from 0 to  $2^{i-1}$ , i.e.,  $\varphi^{\text{pull}} \leq 2^{i-1}$ .

For  $\varphi^{\text{push}}$ , we have  $|\bar{S}_i| = 0$ , i.e.,  $\bar{\varphi}_i^{\text{push}} = 0$ . For  $i < j < m$ , since we have  $\bar{S}_j = S_{j-1}$  and  $|S_{j-1}| \geq 2^{j-1}$  before the recursive push, we have  $\varphi_{j-1}^{\text{push}} \geq 1$  and  $\bar{\varphi}_j^{\text{push}} \leq \varphi_{j-1}^{\text{push}} - 1$ . Since  $\bar{\varphi}_i^{\text{push}} = 0$  and  $\bar{\varphi}_m^{\text{push}} \leq 2$ , it follows that  $\Delta\varphi^{\text{push}} \leq 2 - (m - i - 1) = -m + i + 3$ .  $\square$

**Lemma 5.5:** A DELETE-MIN operation, excluding recursive pulling, takes worst-case  $\mathcal{O}(1)$  time and the change in potential is  $\Delta\Phi \leq \ell$ .  $\triangleleft$

*Proof:* Removing the minimum from  $S_1$  takes worst-case constant time, since  $|S_1| \leq 6$ . Decreasing  $|S_1|$  can only decrease  $\varphi_1^{\text{insert}}$  and  $\varphi_1^{\text{push}}$ , or leave them unchanged. All other  $\varphi_j^{\text{insert}}$  and  $\varphi_j^{\text{push}}$  remain unchanged. Decreasing  $|S_1|$  by one can at most increase  $\varphi_j^{\text{pull}}$  by one for all  $1 \leq j \leq \ell$ . The total change in potential becomes  $\Delta\Phi \leq \ell$ .

Since the deletion of an element causes  $n$  to decrease by one, we might have  $\ell > 1 + \lg n$ . In this case DELETE-MIN in worst-case constant time concatenates  $S_{\ell-1}$  and  $S_\ell$ , and  $\ell$  decreases by one, ensuring  $\ell \leq 1 + \lg n$ . Since  $\ell > 1 + \lg n$  before the concatenation, we have  $|\bar{S}_{\ell-1}| \leq n < 2^{\ell-1}$ , and  $\varphi_{\ell-1}^{\text{push}} = \varphi_{\ell-1}^{\text{insert}} = 0$  and  $\varphi_{\ell-1}^{\text{pull}}$  can only have decreased. It follows that  $\Phi$  can only decrease by performing this additional check during DELETE-MIN.  $\square$

**Lemma 5.6:** The recursive pulling of elements into  $S_1$ , when  $|S_1| = 0$ , takes worst-case  $\mathcal{O}(m + 2^{i-1})$  time and the change in potential is  $\Delta\Phi \leq -m - 2^{i-1} + 1$ , for some  $i$  and  $m$ , where  $1 \leq i < m \leq \ell$ .  $\triangleleft$

*Proof:* Assume when pulling into  $S_1$  we have  $|S_1| = \dots = |S_{m-1}| = 0$  and  $|S_m| > 0$  for some  $1 < m \leq \ell$ . Furthermore, let  $i$  be the largest  $i < m$  where  $2^{i-1} < |S_m|$ . I.e.,  $S_m$  will be iteratively swapped with  $S_{m-1}, \dots, S_{i+1}$ , and for  $1 \leq j \leq i$  we fill  $S_j$  by selecting the  $2^{j-1}$  smallest elements from  $S_{j+1}$  in  $\mathcal{O}(2^j)$  time. It follows that the worst-case time is  $\mathcal{O}(m + \sum_{j=1}^i 2^j) = \mathcal{O}(m + 2^i)$ .

After the pull,  $|S_j| \leq 2^{j-1}$  for  $1 \leq j < m$ , i.e.,  $\bar{\varphi}_j^{\text{insert}} = 0$  and  $\bar{\varphi}_j^{\text{push}} = 0$ . Since  $|S_m|$  decreases, we have that  $\Delta\varphi_m^{\text{insert}} \leq 0$  and  $\Delta\varphi_m^{\text{push}} \leq 0$ . In total  $\varphi^{\text{insert}}$  and  $\varphi^{\text{push}}$  can only decrease.

For  $\varphi^{\text{pull}}$ , we have that  $\bar{\varphi}_m^{\text{pull}}, \dots, \bar{\varphi}_\ell^{\text{pull}}$  are unchanged, since we only move elements from  $S_m$  to  $S_0, \dots, S_{m-1}$ . Since at least one element is moved to  $S_0$ , we have  $\varphi_j^{\text{pull}}$  decreases by at least one for all  $1 \leq j < m$ . For  $1 \leq j \leq i$ , exactly  $2^j$  elements are moved to  $S_j$ , and subsequently subsets hereof to  $S_{j-1}, \dots, S_1$ , i.e., we have  $\Delta\varphi_j^{\text{pull}} = -2^{j-1}$ . It follows that  $\Delta\varphi^{\text{pull}} \leq -\sum_{j=1}^i 2^{j-1} - (m-1-i) = -2^i - m + i + 2 \leq -2^{i-1} - m + 1$ , which is also a bound on  $\Delta\Phi$ .  $\square$

By Lemma 5.1 we have  $\ell = \mathcal{O}(\lg n)$ . Lemmas 5.1, 5.2 and 5.4 imply INSERT takes amortized  $\mathcal{O}(\lg \lg n)$  time. Lemmas 5.1, 5.3 and 5.4 imply DECREASE-KEY takes amortized  $\mathcal{O}(\lg \lg n)$  time. Finally, Lemmas 5.1, 5.5 and 5.6 imply DELETE-MIN takes amortized  $\mathcal{O}(\ell) = \mathcal{O}(\lg n)$  time.

**Corollary 5.7:** *The amortized time for INSERT and DECREASE-KEY is  $\mathcal{O}(\lg \lg n)$  and the amortized time for DELETE-MIN is  $\mathcal{O}(\lg n)$ .*  $\triangleleft$

## 6. Conclusion

We presented three priority queue implementation, following a simple partition-based heap framework that uses elementary data structures and algorithms. In particular for the simplest design, the lazy partition heap, we give a fully self-contained and complete description and analysis.

Replacing the binary search on an array by a binary search tree, the implementation works in the pointer-machine model and requires only constant indegree on nodes, and can hence be made persistent efficiently.

Due to its simplicity and linear structure, our implementation appears amenable to adaptations for external memory, but supporting pointer-based decrease-key operations becomes more complicated (as noted by [NPPS11]). We reserve this for future work.

Using our rule as the concatenation rule in quickheaps promises to yield a highly practical and robust priority queue implementation, combining the best of previously known quickheaps variants. A detailed empirical evaluation is planned for future work.

The astute reader may have noticed that INSERT and DECREASE-KEY take amortized constant time, except for the binary searches over the  $\mathcal{O}(\lg n)$  pivots. By using a non-pointer-machine model data structure such as a fusion tree [FW93] to store the pivots, the running time can be reduced to  $\mathcal{O}(\lg_w \lg n)$ , which under the standard assumption that  $w = \Omega(\log n)$  results in  $\mathcal{O}(1)$  amortized running times for these operations.

**Acknowledgements.** We thank our reviewers for their helpful comments. The second author's work is supported by the Fonds de la Recherche Scientifique — FNRS. The last author is partly supported by EPSRC grant EP/X039447/1.

## References

- [BFMZ04] Gerth Stølting Brodal, Rolf Fagerberg, Ulrich Meyer, and Norbert Zeh. Cache-oblivious data structures and algorithms for undirected breadth-first search and shortest paths. In Torben Hagerup and Jyrki Katajainen, editors, *Algorithm Theory - SWAT 2004, 9th Scandinavian Workshop on Algorithm Theory, Humlebaek, Denmark, July 8-10, 2004, Proceedings*, volume 3111 of *Lecture Notes in Computer Science*, pages 480–492. Springer, 2004. doi:10.1007/978-3-540-27810-8\_41.
- [BFP<sup>+</sup>73] Manuel Blum, Robert W. Floyd, Vaughan R. Pratt, Ronald L. Rivest, and Robert Endre Tarjan. Time bounds for selection. *J. Comput. Syst. Sci.*, 7(4):448–461, 1973. doi:10.1016/S0022-0000(73)80033-9.
- [BLT25] Gerth Stølting Brodal, George Lagogiannis, and Robert E. Tarjan. Strict Fibonacci heaps. *ACM Trans. Algorithms*, 21(2):15:1–15:18, 2025. doi:10.1145/3707692.
- [Bro13] Gerth Stølting Brodal. A survey on priority queues. In Andrej Brodnik, Alejandro López-Ortiz, Venkatesh Raman, and Alfredo Viola, editors, *Space-Efficient Data Structures, Streams, and Algorithms - Papers in Honor of J. Ian Munro on the Occasion of His 66th Birthday*, volume 8066 of *Lecture Notes in Computer Science*, pages 150–163. Springer, 2013. doi:10.1007/978-3-642-40273-9\_11.
- [Cha13] Timothy M. Chan. Quake heaps: A simple alternative to Fibonacci heaps. In Andrej Brodnik, Alejandro López-Ortiz, Venkatesh Raman, and Alfredo Viola, editors, *Space-Efficient Data Structures, Streams, and Algorithms - Papers in Honor of J. Ian Munro on the Occasion of His 66th Birthday*, volume 8066 of *Lecture Notes in Computer Science*, pages 27–32. Springer, 2013. doi:10.1007/978-3-642-40273-9\_3.
- [CR18] Rezaul Alam Chowdhury and Vijaya Ramachandran. Cache-oblivious buffer heap and cache-efficient computation of shortest paths in graphs. *ACM Trans. Algorithms*, 14(1):1:1–1:33, 2018. doi:10.1145/3147172.
- [Elm10] Amr Elmasry. The violation heap: a relaxed Fibonacci-like heap. *Discret. Math. Algorithms Appl.*, 2(4):493–504, 2010. doi:10.1142/S1793830910000838.
- [FR75] Robert W. Floyd and Ronald L. Rivest. Expected time bounds for selection. *Commun. ACM*, 18(3):165–172, 1975. doi:10.1145/360680.360691.
- [Fre99] Michael L. Fredman. On the efficiency of pairing heaps and related data structures. *J. ACM*, 46(4):473–501, 1999. doi:10.1145/320211.320214.
- [FSST86] Michael L. Fredman, Robert Sedgewick, Daniel Dominic Sleator, and Robert Endre Tarjan. The pairing heap: A new form of self-adjusting heap. *Algorithmica*, 1(1):111–129, 1986. doi:10.1007/BF01840439.
- [FT87] Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM*, 34(3):596–615, July 1987. doi:10.1145/28869.28874.
- [FW93] Michael L. Fredman and Dan E. Willard. Surpassing the information theoretic bound with fusion trees. *J. Comput. Syst. Sci.*, 47(3):424–436, 1993. doi:10.1016/0022-0000(93)90040-4.

- [HKTZ17] Thomas Dueholm Hansen, Haim Kaplan, Robert E. Tarjan, and Uri Zwick. Hollow heaps. *ACM Trans. Algorithms*, 13(3):42:1–42:27, 2017. doi:10.1145/3093240.
- [Hoa61] C. A. R. Hoare. Algorithm 65: find. *Commun. ACM*, 4(7):321–322, 1961. doi:10.1145/366622.366647.
- [HST09] Bernhard Haeupler, Siddhartha Sen, and Robert Endre Tarjan. Rank-pairing heaps. In Amos Fiat and Peter Sanders, editors, *Algorithms - ESA 2009, 17th Annual European Symposium, Copenhagen, Denmark, September 7-9, 2009. Proceedings*, volume 5757 of *Lecture Notes in Computer Science*, pages 659–670. Springer, 2009. doi:10.1007/978-3-642-04128-0\_59.
- [IÖ14] John Iacono and Özgür Özkan. Why some heaps support constant-amortized-time decrease-key operations, and others do not. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 637–649. Springer, 2014. doi:10.1007/978-3-662-43948-7\_53.
- [KTZ14] Haim Kaplan, Robert Endre Tarjan, and Uri Zwick. Fibonacci heaps revisited. *CoRR*, abs/1407.5750, 2014. URL: <http://arxiv.org/abs/1407.5750>, arXiv:1407.5750.
- [NP10] Gonzalo Navarro and Rodrigo Paredes. On sorting, heaps, and minimum spanning trees. *Algorithmica*, 57:585–620, 2010. doi:10.1007/s00453-010-9400-6.
- [NPPS11] Gonzalo Navarro, Rodrigo Paredes, Patricio V. Poblete, and Peter Sanders. Stronger quickheaps. *Int. J. Found. Comput. Sci.*, 22:945–969, 2011. doi:10.1142/S0129054111008507.
- [Pet05] Seth Pettie. Towards a final analysis of pairing heaps. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 174–183. IEEE Computer Society, 2005. doi:10.1109/SFCS.2005.75.
- [PN06] Rodrigo Paredes and Gonzalo Navarro. Optimal incremental sorting. In *2006 Proceedings of the Eighth Workshop on Algorithm Engineering and Experiments (ALENEX)*, page 171–182. Society for Industrial and Applied Mathematics, January 2006. doi:10.1137/1.9781611972863.16.
- [RW25] Casper Moldrup Rysgaard and Sebastian Wild. Lazy B-trees. In *Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 345 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 87:1–87:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICS.MFCS.2025.87.
- [ST23a] Corwin Sinnamon and Robert E. Tarjan. A nearly-tight analysis of multipass pairing heaps. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 535–548. SIAM, 2023. URL: <https://doi.org/10.1137/1.9781611977554.ch23>, doi:10.1137/1.9781611977554.CH23.
- [ST23b] Corwin Sinnamon and Robert E. Tarjan. A tight analysis of slim heaps and smooth heaps. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 549–567. SIAM, 2023. URL: <https://doi.org/10.1137/1.9781611977554.ch24>, doi:10.1137/1.9781611977554.CH24.

- [SW20] Bryce Sandlund and Sebastian Wild. Lazy search trees. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 704–715. IEEE, 2020. doi:10.1109/FOCS46700.2020.00071.
- [SZ22] Bryce Sandlund and Lingyi Zhang. Selectable heaps and optimal lazy search trees. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1962–1975, 2022. doi:10.1137/1.9781611977073.78.
- [Wil64] John William Joseph Williams. Algorithm 232: Heapsort. *Communications of the ACM*, 7(6):347–348, 1964. doi:10.1145/512274.512284.

# Appendix

## A. LP Heap Proof-of-Concept Implementation

To illustrate the simplicity of LP Heaps, we provide a fully self-contained Python implementation below. The code is intended to serve as detailed specification of LP Heaps, tuned for readability. It is not intended to be efficient.

We first show the main code for LP Heaps, the class `LazyHeap`; it uses a doubly-linked list and a selection method which we provide below.

```

1 class LazyHeap:
2     def __init__(self):
3         self.counter = 0
4         self.pivots = []
5         self.sets = [LinkedList()]
6
7     def insert(self, value):
8         node = Node((value, self.counter))
9         self.counter += 1
10        self._insert(node)
11        return node
12
13    def delete(self, node):
14        self._remove(node)
15        self._merge_sets()
16
17    def extract_min(self):
18        S = self.sets[0]
19        if len(S) == 0:
20            raise ValueError("Heap is empty")
21        smallest_node = min(S.iter_nodes(),
22                           key=lambda node: node.value)
23        S.remove(smallest_node)
24        if len(S) >= 2:
25            median, first, second = self._partition_set(S)
26            self.pivots = [median] + self.pivots
27            self.sets = [first, second] + self.sets[1:]
28        self._forget_pivots()
29        return smallest_node.value[0]
30
31    def decrease_key(self, node, new_value):
32        assert new_value <= node.value[0]
33        self._remove(node)
34        node.value = (new_value, node.value[0])
35        self._insert(node)
36
37    def _insert(self, node):
38        S = self._find_set(node.value)
39        S.append(node)
40
41    def _remove(self, node):
42        S = self._find_set(node.value)
43        S.remove(node)
44
```

```

47     def _find_set(self, value):
48         if len(self.sets) == 1 or value < self.pivots[0]:
49             return self.sets[0]
50             low = 0
51         high = len(self.pivots)
52         while low + 1 < high:
53             mid = (low + high) // 2
54             if self.pivots[mid] <= value:
55                 low = mid
56             else:
57                 high = mid
58         assert value >= self.pivots[low]
59         return self.sets[low + 1]

61     def _partition_set(self, S):
62         median = select(S, -(-len(S) // 2))
63         smaller = []
64         larger = []
65         for node in S.iter_nodes():
66             if node.value < median:
67                 smaller.append(node)
68             else:
69                 larger.append(node)
70         first = LinkedList()
71         for node in smaller:
72             first.append(node)
73         second = LinkedList()
74         for node in larger:
75             second.append(node)
76         return median, first, second

78     def _forget_pivots(self):
79         new_pivots = []
80         new_sets = [self.sets[0]]
81         outside = 0
82         for pivot, S in zip(self.pivots, self.sets[1:]):
83             if len(new_sets[-1]) == 0 or
84                outside > len(new_sets[-1]) + len(S):
85                 new_sets[-1].concat(S) // merge
86             else:
87                 outside += len(new_sets[-1])
88                 new_pivots.append(pivot)
89                 new_sets.append(S)
90         if len(new_sets[-1]) == 0 and len(new_sets) >= 2:
91             new_pivots.pop()
92             new_sets.pop()
93         self.pivots = new_pivots
94         self.sets = new_sets

```

The code above uses a circularly closed doubly-linked list; since the node objects are used as pointers/references to inserted elements, we provide the simple implementation here:

```

1 class Node:
2     def __init__(self, value):
3         self.value = value

```

```

4         self.prev = None
5         self.next = None

7 class LinkedList:
8     def __init__(self):
9         self.head = Node(None)
10        self.tail = Node(None)
11        self.head.next = self.tail
12        self.tail.prev = self.head
13        self.size = 0

15    def __len__(self):
16        return self.size

18    def append(self, node):
19        self.tail.prev.next = node
20        node.prev = self.tail.prev
21        node.next = self.tail
22        self.tail.prev = node
23        self.size += 1

25    def remove(self, node):
26        node.prev.next = node.next
27        node.next.prev = node.prev
28        node.prev = None
29        node.next = None
30        self.size -= 1

32    def concat(self, other):
33        t = self.tail.prev
34        h = other.head.next
35        t.next = h
36        h.prev = t
37        self.tail = other.tail
38        self.size += other.size

40    def iter_nodes(self):
41        at = self.head.next
42        while at is not self.tail:
43            yield at
44            at = at.next

46    def __iter__(self):
47        for node in self.iter_nodes():
48            yield node.value

```

To be self-contained, we also give a simple implementation of the median-of-medians linear-time deterministic selection method [BFP<sup>+</sup>73]. Note that the method is not in-place and should be considered a “pedagogical” implementation. (A practical implementation would certainly want to use a randomized quickselect variant instead.)

```

1 def select(iterable, k):
2     def inner(arr, k):
3         buckets = [sorted(arr[i:i+5]) for i in range(0, len(arr), 5)]

```

```
5     if len(buckets) == 1:
6         return buckets[0][k]

8     medians = [buck[len(buck) // 2] for buck in buckets]
9     pivot = inner(medians, len(medians) // 2)

11    smaller = []
12    larger = []

14    for value in arr:
15        if value < pivot:
16            smaller.append(value)
17        else:
18            larger.append(value)

20    if k < len(smaller):
21        return inner(smaller, k)
22    else:
23        return inner(larger, k - len(smaller))

25    return inner(list(iterable), k)
```