

Worst-Case Efficient External-Memory Priority Queues

Gerth Stølting Brodal*

Max-Planck-Institut für Informatik
Im Stadtwald
D-66123 Saarbrücken, Germany
E-mail: brodal@mpi-sb.mpg.de

Jyrki Katajainen†

University of Copenhagen
Universitetsparken 1
DK-2100 Copenhagen East, Denmark
E-mail: jyrki@diku.dk

October 2, 1997

Abstract

A priority queue Q is a data structure that maintains a collection of elements, each element having an associated priority drawn from a totally ordered universe, under the operations INSERT, which inserts an element into Q , and DELETEMIN, which deletes an element with the minimum priority from Q . In this paper a priority-queue implementation is given which is efficient with respect to the number of block transfers or I/Os performed between the internal and external memories of a computer. Let B and M denote the respective capacity of a block and the internal memory measured in elements. The developed data structure handles any intermixed sequence of INSERT and DELETEMIN operations such that in every disjoint interval of B consecutive priority-queue operations at most $c \log_{M/B} \frac{N}{M}$ I/Os are performed, for some positive constant c . These I/Os are divided evenly among the operations: if $B \geq c \log_{M/B} \frac{N}{M}$, one I/O is necessary for every $B/(c \log_{M/B} \frac{N}{M})$ th operation and if $B < c \log_{M/B} \frac{N}{M}$, $\frac{c}{B} \log_{M/B} \frac{N}{M}$ I/Os are performed per every operation. Moreover, every operation requires $O(\log_2 N)$ comparisons in the worst case. The best earlier solutions can only handle a sequence of S operations with $O(\sum_{i=1}^S \frac{1}{B} \log_{M/B} \frac{N_i}{M})$ I/Os, where N_i denotes the number of elements stored in the data structure prior to the i th operation, without giving any guarantee for the performance of the individual operations.

1 Introduction

A *priority queue* is a data structure that stores a set of elements, each element consisting of some *information* and a *priority* drawn from some totally ordered universe. A priority queue supports the operations:

INSERT(x): Insert a new element x with an arbitrary priority into the data structure.

DELETEMIN(): Delete and return an element with the minimum priority from the data structure.

In the case of ties, these are broken arbitrarily. The precondition is that the priority queue is not empty.

*Supported by the Carlsberg foundation under grant No. 96-0302/20. Partially supported by the ESPRIT Long Term Research Program of the EU under contract No. 20244 (project ALCOM-IT).

†Supported partially by the Danish Natural Science Research Council under contract No. 9400952 (project Computational Algorithmics).

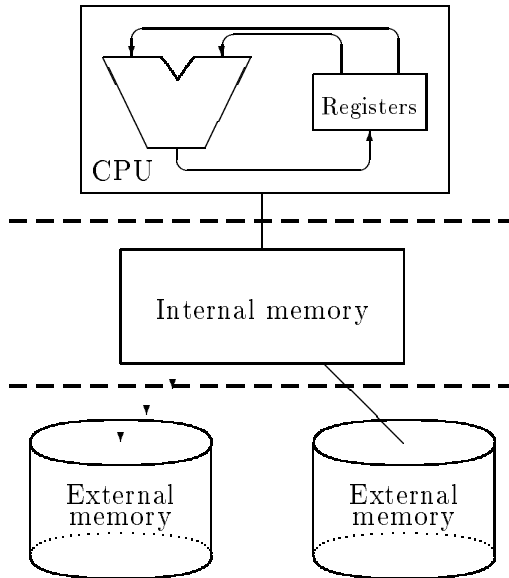


Figure 1: The model of computation

Priority queues have numerous applications, a few listed by Sedgewick [28] are: sorting algorithms, network optimization algorithms, discrete event simulations and job scheduling in computer systems. For the sake of simplicity, we will not hereafter make any distinction between the elements and their priority.

In this paper we study the problem of maintaining a priority queue on a computer with a two-level memory: a fast *internal memory* and a slow *external memory* (see Fig. 1). We assume that the computer has a processing unit, the *processor* or CPU, and a collection of hardware, the *I/O subsystem*, which is responsible for transferring data between internal and external memory. The processor together with the internal memory can be seen as a traditional random access machine (RAM) (see, *e.g.*, [3]). In particular, note that the processor can only access data stored in internal memory. The capacity of the internal memory is assumed to be bounded so it might be necessary to store part of the data in external memory. The I/O subsystem takes care of transferring the data between the two memory levels, and this is done in blocks of a fixed size.

The behavior of algorithms on such a computer system can be characterized by two quantities: *processor performance* and *I/O performance*. By the processor performance we mean the number of primitive operations performed by the processor. Our measure of processor performance is the number of element comparisons carried out. It is straightforward to verify that the total number of other (logical, arithmetical, etc.) operations required by our algorithms is proportional to that of comparisons. Assuming that the elements occupy only a constant number of computer words, the total number of primitive operations is asymptotically the same as that of comparisons. Our measure of I/O performance is the number of block transfers or *I/Os* performed, *i. e.*, the number of blocks read from the external memory plus the number of blocks written to the external memory by the I/O subsystem. Our main goal is to analyze the total work carried out by the processor and the I/O subsystem during the execution of the algorithms.

The *system performance*, *i. e.*, the total elapsed execution time when the algorithms are run on a real computer, depends heavily on the realization of the computer. A real computer may have

multiple processors (see, *e. g.*, [18]) and/or the I/O subsystem can transfer data between several disks at the same time (cf. [2, 25, 30]), the processor operations (see, *e. g.*, [27]) and/or the I/Os (cf. [19]) might be pipelined, but the effect of these factors is not considered here. It has been observed that in many large-scale computations the increasing bottleneck of the computation is the performance of the I/O subsystem (see, *e. g.*, [15, 26]), increasing the importance of I/O efficient algorithms.

When expressing the performance of the priority-queue operations, we use the following parameters:

B : the number of elements per block,

M : the number of elements fitting in internal memory, and

N : the number of elements currently stored in the priority queue; more specifically, the number of elements stored in the structure just prior to the execution of INSERT or DELETEMIN.

We assume that each block and the internal memory also fit some pointers in addition to the elements, and $B \geq 1$ and $M \geq 23B$. Furthermore, we use $\log_a n$ as a shorthand notation for $\max(1, \ln n / \ln a)$, where \ln denotes the natural logarithm.

Several priority-queue schemes, such as implicit heaps [33], leftist heaps [12, 20], and binomial queues [9, 31] have been shown to permit both INSERT and DELETEMIN with worst-case $O(\log_2 N)$ comparisons. Some schemes, such as implicit binomial queues [10] guarantee worst-case $O(1)$ comparisons for INSERT and $O(\log_2 N)$ comparisons for DELETEMIN. Also any kind of balanced search trees, such as AVL trees [1] or red-black trees [16] could be used as priority queues. However, due to the usage of explicit or implicit pointers the performance of these structures deteriorates on a two-level memory system. It has been observed by several researchers that a d -ary heap performs better than the normal binary heap on multi-level memory systems (see, *e. g.*, [22, 24]). For instance, a B -ary heap reduces the number of I/Os from $O(\log_2 \frac{N}{B})$ (cf. [4]) to $O(\log_B \frac{N}{B})$ per operation [24]. Of course, a B -tree [8, 11] could also be used as a priority queue, with which a similar I/O performance is achieved. However, in a virtual-memory environment a B -ary heap seems to be better than a B -tree [24].

When a priority queue is maintained in a two-level memory, it is advantageous to keep the small elements in internal memory and the large elements in external memory. Hence, due to insertions large elements are to be moved from internal memory to external memory and due to deletions small elements are to be moved from external memory to internal memory. Assuming that we maintain two buffers of B elements in internal memory, one for INSERTs and one for DELETEMINs, at most every B th INSERT and DELETEMIN will cause a buffer overflow or underflow. Several data structures take advantage of this kind of buffering. Fishspear, developed by Fischer and Paterson [14], can be implemented by a constant number of push-down stacks, implying that any sequence of S INSERT and DELETEMIN operations requires at most $O(\sum_{i=1}^S \frac{1}{B} \log_2 N_i)$ I/Os, where N_i denotes the size of the data structure prior to the i th operation. Wegner and Teuhola [32] realized that a B -ary heap, in which each node stores B elements, guarantees worst-case $O(\log_2 \frac{N}{B})$ I/Os for every B th INSERT and every B th DELETEMIN operation.

The above structures assume that the internal memory can only fit $O(B)$ elements, *i. e.*, a constant number of blocks. Even faster solutions are possible if the whole capacity of the internal memory is utilized. Arge [5, 6] introduced an (a, b) -tree structure that can be used to carry out any sequence of S INSERT and DELETEMIN operations with $O(\frac{S}{B} \log_{M/B} \frac{S}{M})$ I/Os. Fadel *et al.* [13] gave a heap structure with a similar I/O performance but their bound depends on the size profile, not on S . Their heap structure can handle any sequence of S operations with $O(\sum_{i=1}^S \frac{1}{B} \log_{M/B} \frac{N_i}{M})$

I/Os, where N_i denotes the size of the data structure prior to the i th operation. The number of comparisons required when handling the sequence is $O(\sum_{i=1}^S \log_2 N_i)$. When this data structure is used for sorting N elements, both the processor and I/O performance match the well-known lower bounds $\Omega(\frac{N}{B} \log_{M/B} \frac{N}{M})$ I/Os [2] and $\Omega(N \log_2 N)$ comparisons (see, *e.g.*, [20]), which are valid for all comparison-based algorithms.

To achieve the above bounds—as well as our bounds—the following facilities must be provided:

1. we should know the capacity of a block and the internal memory beforehand,
2. we must be able to align elements into blocks, and
3. we must have a full control over the replacement of the blocks in internal memory.

There are operating systems that provide support for these facilities (see, *e.g.*, [17, 21, 23]).

The tree structure of Arge and the heap structure of Fadel *et al.* do not give any guarantees for the performance of individual operations. In fact, one INSERT or DELETMIN can be extremely expensive, the cost of handling the whole sequence being an upper bound. Therefore, it is risky to use these structures in on-line applications. For large-scale real-time discrete event simulations and job scheduling in computer systems it is often important to have a guaranteed worst-case performance.

We describe a new data structure that gives worst-case guarantees for the cost of individual operations. Basically, our data structure is a collection of sorted lists that are incrementally merged. This idea is borrowed from a RAM priority-queue structure of Thorup [29]. Thorup used two-way merging in his internal data structure but we use multi-way merging since it behaves better in an external-memory environment. As to the processor and I/O performance, our data structure handles any intermixed sequence of operations as efficiently as the heap structure by Fadel *et al.* [13]. In every disjoint interval of B consecutive priority-queue operations our data structure requires at most $c \log_{M/B} \frac{N}{M}$ I/Os, for some positive constant c . These I/Os are divided evenly among the operations. If $B \geq c \log_{M/B} \frac{N}{M}$, one I/O is necessary for every $B/(c \log_{M/B} \frac{N}{M})$ th priority-queue operation, and if $B < c \log_{M/B} \frac{N}{M}$, $\frac{c}{B} \log_{M/B} \frac{N}{M}$ I/Os are performed per every priority-queue operation. Moreover, every operation requires $O(\log_2 N)$ comparisons in the worst case.

The outline of the remaining of the paper is as follows. In Sect. 2 the basic data structure is described. In Sect. 3 algorithms are developed for performing batch insertions and deletions on the external-memory part of the data structure. In Sect. 4 we combine internal-memory buffers with incrementally performed batch operations to achieve the claimed performance bounds. Finally, in Sect. 5 some concluding remarks are given.

2 Basic data structure

The basic component of our priority-queue data structure is a collection of sorted lists. When new elements arrive, these are added to a list which is kept in internal memory and sorted incrementally. If the capacity of internal memory is exceeded due to insertions, a fraction of the list containing the recently inserted elements is transferred to external memory. To bound the number of lists in external memory we merge the existing lists. This merging is related to the merging done by the external mergesort algorithm [2]. One particular list that is kept in internal memory contains the smallest elements. If this list is exhausted due to deletions, new smallest elements are extracted from the lists in external memory. Because we are interested in worst-case bounds the merging is accomplished incrementally. A similar idea has been applied by Thorup [29] to construct RAM priority queues but instead of two-way merging we rely on multi-way merging.

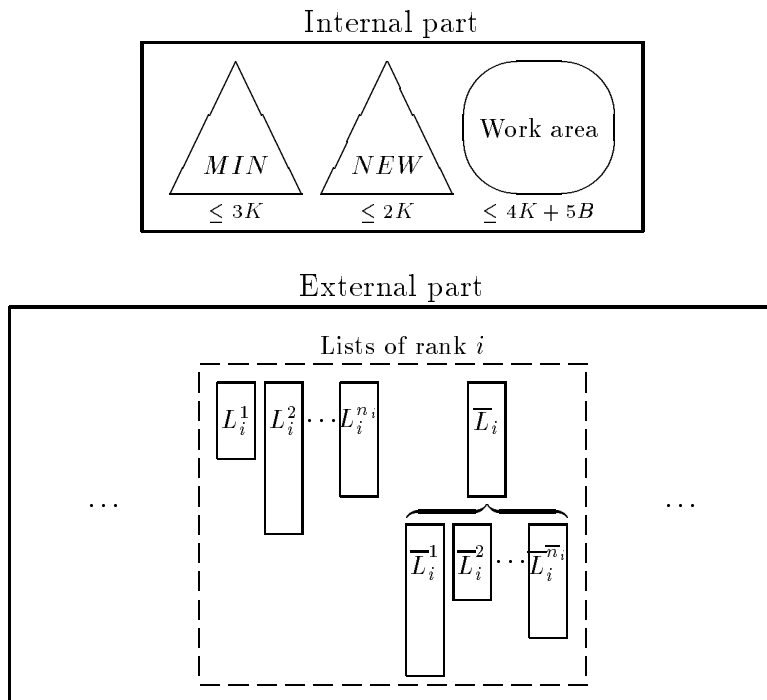


Figure 2: The data structure

Before giving the details of the data structure, let us recall the basic idea of external mergesort which sorts N elements with $O(\frac{N}{B} \log_{M/B} \frac{N}{M})$ I/Os [2]. First, the given N elements are partitioned into $\Theta(N/M)$ lists each of length $\Theta(M)$. Second, each of the lists are read into internal memory and sorted, requiring $O(N/B)$ I/Os in total. Third, $\Theta(M/B)$ sorted lists of shortest length are repeatedly merged until only one sorted list remains containing all the elements. Since each element takes part in $O(\log_{M/B} \frac{N}{M})$ merges, the total number of I/Os is $O(\frac{N}{B} \log_{M/B} \frac{N}{M})$.

Our data structure consists of two parts: an *internal part* and an *external part* (see Fig. 2). The data structure takes two parameters K and m , where K is a multiple of B , $9K + 5B \leq M$, and $m = K/B$. The internal part of the data structure stores $O(K)$ elements and is kept all the time in internal memory. The external part is a priority queue which permits the operations:

BATCHINSERT $_K(X)$: Insert a sorted list X of K elements into the external-memory data structure.

BATCHDELETEMIN $_K()$: Delete the K smallest elements from the external-memory data structure.

Both of these operations require at most $O(\frac{K}{B} \log_m \frac{N}{K})$ I/Os and $O(K \log_2 \frac{N}{K})$ comparisons in the worst case. For every K th operation on the internal part we do at most one batch operation involving K elements on the external part of the data structure.

The internal part of the data structure consists of two sorted lists MIN and NEW of length at most $3K$ and $2K$, respectively. We represent both MIN and NEW as a balanced search tree that permits insertions and deletions of elements with $O(\log_2 K)$ comparisons. The rôle of MIN is to store the current at most $3K$ smallest elements in the priority queue whereas the intuitive rôle of NEW is to store the at most $2K$ recently inserted elements. All elements in MIN are smaller than the elements in NEW and the elements in the external part of the data structure, *i. e.*, the overall minimum element is the minimum of MIN .

The external part of the data structure consists of sorted lists of elements. Each of these lists has a *rank*, which is a positive integer, and we let R denote the *maximum rank*. In Sect. 3.5 we show how to guarantee that $R \leq \log_m \frac{N}{K} + 2$. The lists with rank i , $i \in \{1, \dots, R\}$, are

$$L_i^1, L_i^2, \dots, L_i^{n_i}, \quad \bar{L}_i^1, \bar{L}_i^2, \dots, \bar{L}_i^{\bar{n}_i}, \quad \text{and} \quad \bar{L}_i.$$

For each rank i , the lists $\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i}$ are being merged incrementally and the result of this merge is to be appended to the list \bar{L}_i . The list \bar{L}_i contains the already merged part of $\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i}$. All the elements in \bar{L}_i are smaller than those in $\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i}$. When the incremental merge of lists with rank i finishes, the list \bar{L}_i is promoted to a list with rank $i+1$, provided that \bar{L}_i is sufficiently long, and a new incremental merge of lists with rank i is initiated by making $L_i^1, \dots, L_i^{n_i}$ the new \bar{L}_i^j lists.

We guarantee that the length of each of the external lists is a multiple of B . An external list L containing $|L|$ elements is represented by a single linked list of $|L|/B$ blocks, each block storing B elements plus a pointer to the next block, except for the last block which stores a null pointer. There is one exception to this representation. The last block of \bar{L}_i does not store a null pointer, but a pointer to the first block of \bar{L}_i^1 (if $\bar{n}_i = 0$, the last block of \bar{L}_i stores a null pointer). This allows us to avoid updating the last block of \bar{L}_i when merging the lists $\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i}$.

In the following, we assume that pointers to all the external lists are kept in internal memory together with their sizes and ranks. If this is not possible, it is sufficient to store this information in a single linked list in external memory. This increases the number of I/Os required by our algorithms only by a small constant.

In Sect. 3 we describe how `BATCHINSERTK` and `BATCHDELETEMINK` operations are accomplished on the external part of the data structure, and in Sect. 4 we describe how the external part can be combined with the internal part of the data structure to achieve a worst-case efficient implementation of `INSERT` and `DELETEMIN` operations.

3 Maintenance of the external part

3.1 Invariants

We start by presenting the *invariants* maintained by the `BATCHINSERTK` and `BATCHDELETEMINK` operations on the external part of the data structure. Recall that the lists with rank i , $i \in \{1, \dots, R\}$, are $L_i^1, \dots, L_i^{n_i}, \bar{L}_i$, and $\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i}$. As earlier, we let $|L|$ denote the number of elements in list L .

The aim of invariants (1) and (2) is to guarantee that the utilization of the external memory is maximized, *i. e.*, all the blocks in external memory are full of elements.

$$\sum_{j=1}^{n_i} |L_i^j|, \quad \sum_{j=1}^{\bar{n}_i} |\bar{L}_i^j|, \quad \text{and} \quad |\bar{L}_i| \quad \text{are multiples of } K \text{ (and } B), \text{ for all } i \in \{1, \dots, R\}. \quad (1)$$

$$|L_i^j| \quad \text{and} \quad |\bar{L}_i^j| \quad \text{are multiples of } B, \text{ for all } i \text{ and } j. \quad (2)$$

Invariants (3) and (4) are to avoid considering empty lists.

$$|L_i^j| > 0 \quad \text{and} \quad |\bar{L}_i^j| > 0, \quad \text{for all } i \text{ and } j. \quad (3)$$

$$\bar{n}_i > 0 \Rightarrow |\bar{L}_i| \geq K, \quad \text{for all } i \in \{1, \dots, R\}. \quad (4)$$

The remaining invariants capture the progress of the incremental merge of lists with rank i . We have two invariants especially for rank 1. The total length of the lists L_1^j is by invariant (5) bounded by the length of the list merged already. Invariant (6) bounds the total length of lists with rank 1 that are still to be merged.

$$\sum_{j=1}^{n_1} |L_1^j| - |\bar{L}_1| \leq 0 . \quad (5)$$

$$\sum_{j=1}^{n_1} |L_1^j| + \sum_{j=1}^{\bar{n}_1} |\bar{L}_1^j| \leq 2Km - K . \quad (6)$$

For rank $i \in \{2, \dots, R\}$, we have similar invariants. Because \bar{L}_{i-1} can be promoted to a new L_i^j list before the next incremental merge of lists with rank i is initiated, we “include” $|\bar{L}_{i-1}|$ in the sum $\sum_{j=1}^{n_i} |L_i^j|$.

$$(|\bar{L}_{i-1}| + \sum_{j=1}^{n_i} |L_i^j|) - |\bar{L}_i| \leq 2K \frac{m^i - m}{m - 1}, \quad \text{for all } i \in \{2, \dots, R\} . \quad (7)$$

$$(|\bar{L}_{i-1}| + \sum_{j=1}^{n_i} |L_i^j|) + \sum_{j=1}^{\bar{n}_i} |\bar{L}_i^j| \leq 2K \frac{m^{i+1} - m}{m - 1}, \quad \text{for all } i \in \{2, \dots, R\} . \quad (8)$$

The last invariant bounds the length of the result of each incremental merge.

$$|\bar{L}_i| + \sum_{j=1}^{\bar{n}_i} |\bar{L}_i^j| \leq 2K \frac{m^{i+1} - m}{m - 1}, \quad \text{for all } i \in \{1, \dots, R\} . \quad (9)$$

3.2 The MERGE_K procedure

The heart of our construction is the incremental merge procedure MERGE_K(X_1, X_2, \dots, X_ℓ), which merges and removes the K smallest elements from the sorted lists X_1, \dots, X_ℓ . All list lengths are assumed to be multiples of B . After the merging of the K smallest elements we rearrange the remaining elements in the X_i lists such that the lists still have lengths which are multiples of B . We allow MERGE_K to make the X_i lists shorter or longer. We just require that the resulting X_i lists remain sorted. For the time being, we assume that the result of MERGE_K is stored in internal memory.

The procedure MERGE_K is implemented as follows. For each list X_i we keep the block containing the current minimum of X_i in internal memory. In internal memory we maintain a heap [33] over the current minima of all the lists. We use the heap to find the next element to be output in the merging process. Whenever an element is output, it is the current minimum of some list X_i . We remove the element from the heap and the list X_i , and insert the new minimum of X_i into the heap, provided that X_i has not become empty. If necessary, we read the next block of X_i into internal memory.

After the merging phase, we have from each list X_i a partially filled block B_i in internal memory. Let $|B_i|$ denote the number of elements left in block B_i . Because we have merged K elements from the blocks read and K is a multiple of B , $\sum_{i=1}^{\ell} |B_i|$ is also a multiple of B . Now we merge the remaining elements in the B_i blocks in internal memory. This merge requires $O(\ell B \log_2 \ell)$ comparisons. Let \hat{X} denote the resulting list and let B_j be the block that contained the maximum

element of \hat{X} . Finally, we write \hat{X} to external memory such that X_j becomes the list consisting of \hat{X} concatenated with the part of X_j that already was stored in external memory. Note that X_j remains sorted.

In total, MERGE_K performs at most $K/B + \ell$ I/Os for reading the prefixes of X_1, \dots, X_ℓ (for each list X_i , we read at most one block of elements that do not take part in the merging) and at most ℓ I/Os for writing \hat{X} to external memory. The number of comparisons required for MERGE_K for each of the $K + \ell B$ elements read into internal memory is $O(\log_2 \ell)$. Hence, we have proved

Lemma 1 $\text{MERGE}_K(X_1, \dots, X_\ell)$ performs at most $2\ell + K/B$ I/Os and $O((K + \ell B) \log_2 \ell)$ comparisons. The number of elements to be kept in internal memory by MERGE_K is at most $K + \ell B$. If the resulting list is written to external memory incrementally, only $(\ell + 1)B$ elements have to be kept in internal memory simultaneously.

3.3 Batch insertions

To insert a sorted list of K elements into the external part of the data structure we increment n_1 and let $L_1^{n_1}$ contain the K new elements. To reestablish the invariants we apply the procedure $\text{MERGESTEP}(i)$, for each $i \in \{1, \dots, R\}$.

The procedure $\text{MERGESTEP}(i)$ does the following. If $\bar{n}_i = 0$, the incremental merge of lists with rank i is finished, and we make \bar{L}_i the list $L_{i+1}^{n_{i+1}+1}$, provided that $|\bar{L}_i| \geq Km^i$. Otherwise, we let \bar{L}_i be the list $L_i^{n_i+1}$ because the list is too short to be promoted. Finally, we initiate a new incremental merge by making the lists $L_i^1, \dots, L_i^{n_i}$ the new \bar{L}_i^j lists. If $\bar{n}_i > 0$, we concatenate \bar{L}_i with the result of $\text{MERGE}_K(\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i})$, *i. e.*, we perform K steps of the incremental merge of $\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i}$. Note that, by writing the first block of the merged list to the block of external memory that stored the first block of \bar{L}_i^1 earlier, we do not need to update the pointer in the previous last block of \bar{L}_i . Pseudo code for the BATCHINSERT_K and the MERGESTEP procedures is given in Fig. 3.

We now argue that this implementation of batch insertions reestablishes the invariants. That invariants (1)–(4) remain satisfied is straightforward to verify. For invariants (5)–(9) we give an inductive argument in i that the invariants become reestablished.

For the lists of rank one there are two cases to consider. If $\sum_{j=1}^{\bar{n}_1} |\bar{L}_1^j| \geq K$, $\text{MERGESTEP}(1)$ increases $|\bar{L}_1|$ by K and decreases $\sum_{j=1}^{\bar{n}_1} |\bar{L}_1^j|$ by K . Because the list containing the K new elements increases $\sum_{j=1}^{n_1} |L_1^j|$ by K , it follows that $\text{MERGESTEP}(1)$ in this case reestablishes the invariants (5) and (6). Otherwise $\sum_{j=1}^{\bar{n}_1} |\bar{L}_1^j| = 0$ and after performing $\text{MERGESTEP}(1)$ invariant (5) is satisfied because $\sum_{j=1}^{n_1} |L_1^j| = 0$. There are now two subcases depending on if \bar{L}_1 is made a list with rank 1 or 2. If \bar{L}_1 is promoted, the merging part of $\text{MERGESTEP}(1)$ decreases the left-hand side of (6) by K and reestablishes (6). On the other hand, if \bar{L}_1 is not promoted, then $|\bar{L}_1| \leq Km - K$ and by (5) we have that before inserting the K new elements $\sum_{j=1}^{n_1} |L_1^j| \leq Km - K$. We conclude that (6) is satisfied after the execution of $\text{MERGESTEP}(1)$. Because the left-hand side of (9) only changes when a new incremental merge is initialized and (6) is satisfied prior to the insertion, it follows that invariant (9) is satisfied for rank 1 after the execution of $\text{MERGESTEP}(1)$.

For rank $i \in \{2, \dots, R\}$, $\text{MERGESTEP}(i-1)$ increases the left-hand side of (7) and (8) by at most K . If $\sum_{j=1}^{\bar{n}_i} |\bar{L}_i^j| \geq K$, $\text{MERGESTEP}(i)$ decreases $\sum_{j=1}^{\bar{n}_i} |\bar{L}_i^j|$ by K and (7) and (8) are reestablished. The left-hand side of (9) does not change and (9) remains satisfied. Otherwise $\sum_{j=1}^{\bar{n}_i} |\bar{L}_i^j| = 0$. Because $\sum_{j=1}^{n_i} |L_i^j| = 0$ after the execution of $\text{MERGESTEP}(i)$ and (9) is satisfied for rank $i-1$, it follows that invariant (7) is satisfied for rank i . To see that (8) is satisfied we consider two subcases depending on whether \bar{L}_i is promoted or not. If \bar{L}_i is promoted, the merging part of


```

procedure MERGESTEP( $i$ )
  if  $\bar{n}_i = 0$  and  $|\bar{L}_i| > 0$  then
    if  $|\bar{L}_i| \geq Km^i$  then
       $n_{i+1} \leftarrow n_{i+1} + 1, L_{i+1}^{n_{i+1}} \leftarrow \bar{L}_i, \bar{L}_i \leftarrow \emptyset$ 
    else
       $n_i \leftarrow n_i + 1, L_i^{n_i} \leftarrow \bar{L}_i, \bar{L}_i \leftarrow \emptyset$ 
    fi
  fi
  if  $\bar{n}_i = 0$  and  $n_i > 0$  then
     $\bar{n}_i \leftarrow n_i, (\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i}) \leftarrow (L_i^1, \dots, L_i^{n_i}), n_i \leftarrow 0$ 
  fi
  if  $\bar{n}_i > 0$  then
     $\bar{L}_i \leftarrow \bar{L}_i \cdot \text{MERGE}_K(\bar{L}_i^1, \dots, \bar{L}_i^{\bar{n}_i})$ 
    remove empty  $\bar{L}_i^j$  lists
  fi
end

procedure BATCHINSERT $_K(X)$ 
   $n_1 \leftarrow n_1 + 1, L_1^{n_1} \leftarrow X$ 
  for  $i \leftarrow 1$  to  $R$  do MERGESTEP( $i$ ) od
end

```

Figure 3: The MERGESTEP procedure and the insertion of K elements into the external part.

MERGESTEP(i) decreases the left-hand side of (8) by K and (8) is reestablished, provided that the left-hand side of (8) is different from $|\bar{L}_{i-1}|$. Otherwise (8) is satisfied due to (9) for rank $i - 1$. If \bar{L}_i is not promoted, then $|\bar{L}_i| \leq Km^i - K$ and by (7) the left-hand side of (8) is bounded by

$$K + 2K \frac{m^i - m}{m - 1} + 2(Km^i - K) < 2K \frac{m^{i+1} - m}{m - 1}. \quad (10)$$

To see that (9) is reestablished note that the left-hand side of (9) for rank i only changes when MERGESTEP(i) has finished its merge and makes \bar{L}_i a list with rank i or $i + 1$. If \bar{L}_i keeps its rank, then the left-hand side of (9) is similarly to (8) bounded by (10). If \bar{L}_i is promoted, then (9) follows from (8), because if MERGESTEP($i - 1$) temporarily makes (8) violated for rank i by increasing the left-hand side by K , it is the case that $|\bar{L}_{i-1}| \geq K$. We conclude (9) is reestablished.

The total number of I/Os performed in a batched insertion of K elements is K/B for writing the K new elements to external memory and by Lemma 1 at most $2(\bar{n}_i + K/B)$ for incrementally merging the lists with rank i . The number of comparisons for rank i is $O((\bar{n}_i B + K) \log_2 \bar{n}_i)$. The maximum number of elements to be stored in internal memory for batched insertions is $(\bar{n}_{\max} + 1)B$, where $\bar{n}_{\max} = \max\{\bar{n}_1, \dots, \bar{n}_R\}$. To summarize, we have

Lemma 2 *A sorted list of K elements can be inserted into the external part of the data structure by performing $(1 + 2R)K/B + 2 \sum_{i=1}^R \bar{n}_i$ I/Os and $O(\sum_{i=1}^R (\bar{n}_i B + K) \log_2 \bar{n}_i)$ comparisons. At most $(\bar{n}_{\max} + 1)B$ elements need to be stored in internal memory.*

3.4 Batch deletions

The removal of the K smallest elements from the external part of the data structure is carried out in two steps. In the first step the K smallest elements are located without affecting the invariants. In the second step the actual deletion is accomplished followed by some processing to reestablish the invariants.

Let \mathcal{L} be one of the lists L_i^1 or \bar{L}_i , for some i , or an empty list. We will guarantee that the list \mathcal{L} contains the K smallest elements of the lists considered so far. Initially \mathcal{L} is empty. Note that by invariant (4) we do not have to consider lists \bar{L}_i^j when finding the minimum elements. By performing $L_i^1 \leftarrow \text{MERGE}_K(L_i^1, \dots, L_i^{n_i}) \cdot L_i^1$ no invariant changes its truth value and L_i^1 now contains the K smallest elements of $L_i^1, \dots, L_i^{n_i}$. The procedure SPLITMERGE_K takes two sorted lists as its arguments and returns (the name of) one of the lists. If the first argument is an empty list, then the second list is returned. Otherwise, we require that the length of both lists to be at least K and we rearrange the K smallest elements of both lists as follows. The two prefixes of length K are merged and split among the two lists such that the lists remain sorted and the length of the lists remain unchanged. One of the lists will now have a prefix containing K elements which are smaller than all the elements in the other list. The list with this prefix is returned. For each rank $i \in \{1, \dots, R\}$, we now carry out the assignments $L_i^1 \leftarrow \text{MERGE}_K(L_i^1, \dots, L_i^{n_i}) \cdot L_i^1$, $\mathcal{L} \leftarrow \text{SPLITMERGE}_K(\mathcal{L}, L_i^1)$, and $\mathcal{L} \leftarrow \text{SPLITMERGE}_K(\mathcal{L}, \bar{L}_i)$.

It is straightforward to verify that after performing the above, the prefix of the list \mathcal{L} contains the K smallest elements in the external part of the data structure and that the invariants are satisfied. We now delete the K smallest elements from list \mathcal{L} . If \mathcal{L} is L_i^1 , it follows that $\sum_{j=1}^{n_i} |L_i^j|$ has been decreased by K , and the invariants remain satisfied. Otherwise, \mathcal{L} is \bar{L}_i and $|\bar{L}_i|$ has been decreased by K . The only invariant that can become violated is (7) if $i \geq 2$ or (5) if $i = 1$. By performing $\text{MERGESTEP}(i)$ the invariants will become reestablished. There are two cases to consider depending on $\sum_{j=1}^{\bar{n}_i} |\bar{L}_i^j|$. If the sum is nonzero, the MERGESTEP operation will increase $|\bar{L}_i|$ by $|K|$ and the invariants again become reestablished. If the sum is zero, the argument follows as for insertions. Fig. 4 gives pseudo code for the deletion of the K smallest elements. The procedure POP_K removes and returns the first K elements of a list.

```

procedure BATCHDELETEMIN $_K$ ()
   $\mathcal{L} \leftarrow \emptyset$ 
  for  $i \leftarrow 1$  to  $R$  do
    if  $n_i > 0$  then
       $L_i^1 \leftarrow \text{MERGE}_K(L_i^1, \dots, L_i^{n_i}) \cdot L_i^1$ 
      remove empty  $L_i^j$  lists
       $\mathcal{L} \leftarrow \text{SPLITMERGE}_K(\mathcal{L}, L_i^1)$ 
    fi
    if  $|\bar{L}_i| > 0$  then  $\mathcal{L} \leftarrow \text{SPLITMERGE}_K(\mathcal{L}, \bar{L}_i)$  fi
  od
   $X \leftarrow \text{POP}_K(\mathcal{L})$ 
  if  $\mathcal{L} = \bar{L}_i$  then  $\text{MERGESTEP}(i)$  fi
  return  $X$ 
end

```

Figure 4: The deletion of K smallest elements from the external part.

By always keeping the prefix of \mathcal{L} in internal memory the total number of I/Os for the deletion of the K smallest elements (without the call to `MERGESTEP`) is $(4R-1)(K/B) + 2 \sum_{i=1}^R n_i$ because, for each rank i , $n_i + 2(K/B)$ blocks are to be read into internal memory and all blocks except the K/B blocks holding the smallest elements should be written back to external memory. The number of comparisons for rank i is $O((K + Bn_i) \log_2 n_i)$. The additional call to `MERGESTEP` requires at most $K/B + \bar{n}_i$ additional block reads and block writes, and $O((K + B\bar{n}_i) \log_2 \bar{n}_i)$ comparisons. The maximum number of elements to be stored in internal memory for the batched minimum deletions is $2K + B \max\{n_{\max}, \bar{n}_{\max}\}$, where $n_{\max} = \max\{n_1, \dots, n_R\}$ and $\bar{n}_{\max} = \max\{\bar{n}_1, \dots, \bar{n}_R\}$.

Lemma 3 *The K smallest elements can be deleted from the external part of the data structure by performing at most $4R(K/B) + 2 \sum_{i=1}^R n_i + \bar{n}_{\max}$ I/Os and $O(\sum_{i=1}^R (K + n_i B) \log_2 n_i + (K + B\bar{n}_{\max}) \log_2 \bar{n}_{\max})$ comparisons. At most $2K + B \max\{n_{\max}, \bar{n}_{\max}\}$ elements need to be stored in internal memory.*

3.5 Bounding the maximum rank

We now describe a simple approach to guarantee that the maximum rank R of the external data structure is bounded by $\log_m N/K + 2$. Whenever insertions cause the maximum rank to increase, this is because of `MERGESTEP`($R-1$) has finished an incremental merge resulting a list of length Km^{R-1} , which implies that $R \leq \log_m \frac{N}{K} + 1$. The problem we have to consider is how to decrement R when deletions are performed.

Our solution is the following. Whenever `MERGESTEP`(R) finishes the incremental merge of lists with rank R , we check if the resulting list \bar{L}_R is very small. If \bar{L}_R is very small, *i. e.*, $|\bar{L}_R| < Km^{R-1}$, and there are no other list of rank R , we make \bar{L}_R a list with rank $R-1$ and decrease R . The modifications to the `MERGESTEP` procedure are shown in Fig. 5.

```

procedure MERGESTEP( $i$ )
  if  $i \geq 2$  and  $i = R$  and  $n_R = 0$  and  $\bar{n}_R = 0$  and  $|\bar{L}_R| < Km^{R-1}$  then
     $n_{R-1} \leftarrow n_{R-1} + 1, L_i^{n_{R-1}} \leftarrow \bar{L}_R, R \leftarrow R - 1, i \leftarrow R$ 
  fi
  the code for MERGESTEP from Fig. 3
end

```

Figure 5: The modifications to the `MERGESTEP` procedure to handle deletions.

To guarantee that the same is done also in the connection with batched minimum deletions, we always call after each `BATCHDELETETEMINK` operation, described in Sect. 3.4, `MERGESTEP`(R) k times (for $m \geq 4$ it turns out that $k = 1$ is sufficient).

Assume that at some point the number of elements is $N' = Km^{R-1}$ and that the following sequence of `INSERT` and `DELETETEMIN` operations keeps $N < N'$, *i. e.*, it is guaranteed that no call to `MERGESTEP`($R-1$) creates a new list with rank R . In the following we give a bound on the number of deletions that can be performed before the maximum rank is guaranteed to decrease.

Note that for every deletion of K elements from the external part the length of \bar{L}_R is increased by kK , implying that after at most $N'/(k+1)$ batched deletions $\bar{n}_R = 0$. If $n_R > 0$, another incremental merge is required for rank R . Because we assume $N < Km^{R-1}$, it follows that $n_R = 0$

when the second incremental merge finishes. Therefore, we conclude that at most

$$N' \left(\frac{1}{k+1} + \frac{k}{k+1} \cdot \frac{1}{k+1} \right) = N' \frac{2k+1}{(k+1)^2} \quad (11)$$

deletions can be performed before R is guaranteed to decrease.

If $m = 2$, then at least $N'/2$ (in general N'/m) elements have to be present when R is decreased. From equation (11) it follows that this is guaranteed if $k \geq 3$. For $m = 3$ it is sufficient if $k = 2$, and for $m \geq 4$ it is only required that $k \geq 1$. We conclude that when R decreases, $N \geq N'/m = Km^{R-2}$, and we get $R \leq \log_m \frac{N}{K} + 2$.

3.6 The merging degree

In the previous discussion we assumed that n_i and \bar{n}_i were sufficiently small, such that we could apply MERGESTEP to the L_i^j and \bar{L}_i^j lists. Let m' denote a maximum bound on the merging degree. The L_i^j lists can be created in three different ways:

1. when MERGESTEP($i-1$) makes \bar{L}_{i-1} an L_i^j list,
2. when the resulting \bar{L}_i list is made an L_i^j list, and
3. when the resulting \bar{L}_{i+1} list is made an L_i^j list because the maximum rank decreases from $i+1$ to i .

At most one L_i^j can be created because of 2, and at most one because of 3.

From invariant (6) or (8) it follows that at most $2K \frac{m^{i+1}-m}{m-1}$ insertions can be done before a new incremental merge of lists with rank i is initiated. Actually, if an incremental merge of lists of rank $i \geq 2$ is initiated, then from (8) it follows that at most $2K \frac{m^{i+1}-m}{m-1} - |\bar{L}_{i-1}|$ insertions can follow before a new incremental merge is initiated. Because $|\bar{L}_{i-1}|$ only increases when performing insertions, and because \bar{L}_{i-1} is promoted to an L_i^j list only when its length is at least Km^{i-1} , it follows that the number of lists created due to 1 is at most $\left\lfloor 2K \frac{m^{i+1}-m}{m-1} / (Km^{i-1}) \right\rfloor$, and it follows

$$m' \leq 2 + \left\lfloor 2K \frac{m^{i+1}-m}{(m-1)(Km^{i-1})} \right\rfloor \leq 5 + 2m, \quad \text{for } m \in \{2, 3, 4, \dots\}.$$

3.7 Resource bounds for the external part

Because $m = K/B$ it follows that the maximum rank is at most $\log_{K/B} \frac{N}{K} + 2$ and that the maximum merging degree is $5 + 2K/B$. From Lemmas 2 and 3 it follows that the number of I/Os required for inserting K elements or deleting the K smallest elements is at most $O(\frac{K}{B} \log_{K/B} \frac{N}{K})$ and the number of comparisons required is $O(K \log_2 \frac{N}{K})$. The maximal number of elements to be stored in internal memory is $4K + 5B$.

4 Internal buffers and incremental batch operations

We now describe how to combine the buffers *NEW* and *MIN* represented by binary search trees with the external part of the priority-queue data structure. We maintain the invariant that $|MIN| \geq 1$, provided that the priority queue is nonempty. Recall that we also required that $|MIN| \leq 3K$ and $|NEW| \leq 2K$.

We first consider $\text{INSERT}(x)$. If x is less than or equal to the maximum of MIN or all elements of the priority queue are stored in MIN , we insert x into MIN with $O(\log_2 K)$ comparisons. If MIN exceeds its maximum allowed size, $|\text{MIN}| = 3K + 1$, we move the maximum of MIN to NEW . Otherwise, x is larger than the maximum of MIN and we insert x into NEW with $O(\log_2 K)$ comparisons. The implementation of DELETETEMIN deletes and returns the minimum of MIN . Both operations require at most $O(\log_2 K)$ comparisons.

There are two problems with the above implementation of INSERT and DELETETEMIN . Insertions can cause NEW to become too big and deletions can make MIN empty. We therefore for every K th priority-queue operation perform one batch insertion or deletion. If $|\text{NEW}| \geq K$, we remove K elements from NEW one by one and perform BATCHINSERT_K on the removed elements. If $|\text{NEW}| < K$ and $|\text{MIN}| \leq 2K$, we instead increase the size of MIN by moving K small elements to MIN as follows. First, we perform a $\text{BATCHDELETETEMIN}_K$ operation to extract the K least elements from the external part of the data structure. The K extracted elements are inserted into NEW one by one, using $O(K \log_2 K)$ comparisons. Second, we move the K smallest elements of NEW to MIN one by one. If $|\text{NEW}| < K$ and $|\text{MIN}| > 2K$, we do nothing but delay the batch operation until $|\text{MIN}| = 2K$ or $|\text{NEW}| = K$. Each batch operation requires at most $O(\frac{K}{B} \log_{K/B} \frac{N}{K})$ I/Os and at most $O(K(\log_2 \frac{N}{K} + \log_2 K)) = O(K \log_2 N)$ comparisons.

By doing one of the above described batch operations for every K th priority-queue operation it is straightforward to verify that

$$|\text{NEW}| + (3K - |\text{MIN}|) \leq 2K,$$

provided that the priority queue contains at least K elements, implying $|\text{NEW}| \leq 2K$ and $|\text{MIN}| \geq K$, because each batch operation decreases the left-hand side of the equation by K .

The idea is now to perform a batch operation incrementally over the next K priority-queue operations. Let N denote the number of elements in the priority queue, when the corresponding batch operation is initiated. Notice that N can at most be halved while performing a batch operation, because $N \geq 2K$ prior to the batch operation. Because $|\text{MIN}| \geq K$ when a batch operation is initiated, it is guaranteed that MIN is nonempty while incrementally performing the batch operation over the next K priority-queue operations.

Because a batch operation requires at most $O(\frac{K}{B} \log_{K/B} \frac{N}{K})$ I/Os and at most $O(K \log_2 N)$ comparisons, it is sufficient to perform at most $O(\log_2 N)$ comparisons of the incremental batch operation per priority-queue operation and if $B \geq c \log_{M/B} \frac{N}{M}$, one I/O for every $B/(c \log_{M/B} \frac{N}{M})$ th priority-queue operation and if $B < c \log_{M/B} \frac{N}{M}$, $\frac{c}{B} \log_{M/B} \frac{N}{M}$ I/Os for every priority-queue operation, for some positive constant c , to guarantee that the incremental batch operation is finished after K priority-queue operations.

Because $|\text{MIN}| \leq 3K$, $|\text{NEW}| \leq 2K$, and a batched operation at most requires $4K + 5B$ elements to be stored in internal memory, we have the constraint that $9K + 5B \leq M$. Let now $K = \lfloor (M - 5B)/9 \rfloor$. Recall that we assumed that $M \geq 23B$. Therefore, $K \geq 2B$. Since $M > K$, $O(\log_{M/B} \frac{N}{M}) = O(\log_{K/B} \frac{N}{M})$. Hence, we have proved the main result of this paper.

Main theorem *There exists an external-memory priority-queue implementation that supports INSERT and DELETETEMIN operations with worst-case $O(\log_2 N)$ comparisons per operation. If $B \geq c \log_{M/B} \frac{N}{M}$, one I/O is necessary for every $B/(c \log_{M/B} \frac{N}{M})$ th operation and if $B < c \log_{M/B} \frac{N}{M}$, $\frac{c}{B} \log_{M/B} \frac{N}{M}$ I/Os are performed per every operation, for some positive constant c .*

5 Concluding remarks

We have presented an efficient priority-queue implementation which guarantees a worst-case bound on the number of comparisons and I/Os required for the individual priority-queue operations. Our bounds are comparison based.

If the performance bounds are allowed to be amortized the data structure can be simplified considerably, because no list merging and batch operation is required to be incrementally performed. Then no \overline{L}_i and \overline{L}_i^j lists are required, and we can satisfy $1 \leq |MIN| \leq K$, $|NEW| \leq K$, and $n_i < m$ by always (completely) merging exactly m lists of equal rank, the rank of a list L being $\lfloor \log_m \frac{|L|}{K} \rfloor$.

What if the size of the elements or priorities is not assumed to be constant? That is, express the bounds as a function of N and the length of the priorities. How about the priorities having variable lengths? Initial research in this direction has been carried out by Arge *et al.* [7], who consider sorting strings in external memory.

Acknowledgment

We would like to thank Jukka Teuhola for helpful discussions.

References

- [1] G.M. Adel'son-Vel'skiĭ and E.M. Landis. An algorithm for the organization of information, *Soviet Mathematics*, Volume 3, pages 1259–1263, 1962.
- [2] A. Aggarwal and J.S. Vitter. The input/output complexity of sorting and related problems. *Communications of the ACM*, Volume 31, pages 1116–1127, 1988.
- [3] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley Publishing Company, Reading, 1974.
- [4] T.O. Alanko, H.H.A. Erkiö, and I.J. Haikala. Virtual memory behavior of some sorting algorithms. *IEEE Transactions on Software Engineering*, Volume SE-10, pages 422–431, 1984.
- [5] L. Arge. The buffer tree: A new technique for optimal I/O-algorithms. In *Proceedings of the 4th Workshop on Algorithms and Data Structures*, Lecture Notes in Computer Science 955, pages 334–345, Springer, Berlin/Heidelberg, 1995.
- [6] L. Arge. Efficient external-memory data structures and applications. BRICS Dissertation DS-96-3, Department of Computer Science, University of Aarhus, Århus, 1996.
- [7] L. Arge, P. Ferragina, R. Grossi, and J.S. Vitter. On sorting strings in external memory. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, ACM Press, New York, pages 540–548, 1997.
- [8] R. Bayer and E.M. McCreight. Organization and maintenance of large ordered indexes. *Acta Informatica*, Volume 1, pages 173–189, 1972.
- [9] M.R. Brown. Implementation and analysis of binomial queue algorithms. *SIAM Journal on Computing*, Volume 7, pages 298–319, 1978.
- [10] S. Carlsson, J.I. Munro, and P.V. Poblete. An implicit binomial queue with constant insertion time. In *Proceedings of the 1st Scandinavian Workshop on Algorithm Theory*, Lecture Notes in Computer Science 318, pages 1–13, Springer-Verlag, Berlin/Heidelberg, 1988.
- [11] D. Comer. The ubiquitous B -tree. *ACM Computing Surveys*, Volume 11, pages 121–137, 1979.
- [12] C.A. Crane. Linear lists and priority queues as balanced trees. Technical Report STAN-CS-72-259, Computer Science Department, Stanford University, Stanford, 1972.

- [13] R. Fadel, K. V. Jakobsen, J. Katajainen, and J. Teuhola. Heaps and heapsort on secondary storage. Submitted to *Theoretical Computer Science*. A preliminary version appeared as “External heaps combined with effective buffering” in *Proceedings of the Computing: The Australasian Theory Symposium, Australian Computer Science Communications*, Volume 19, Number 2, pages 72–78, 1997.
- [14] M. J. Fischer and M. S. Paterson. Fishspear: A priority queue algorithm. *Journal of the ACM*, Volume 41, pages 3–30, 1994.
- [15] G.A. Gibson, J.S. Vitter, J. Wilkes *et al.* Strategic directions in storage I/O issues in large-scale computing. *ACM Computing Surveys*, Volume 28, pages 779–793, 1996.
- [16] L. J. Guibas and R. Sedgewick. A dichromatic framework for balanced trees. In *Proceedings of the 19th Annual Symposium on Foundations of Computer Science*, IEEE, New York, pages 8–21, 1978.
- [17] K. Harty and D.R. Cheriton. Application-controlled physical memory using external page-cache management. In *Proceedings of the 5th International Conference on Architectural Support for Programming Languages and Operating Systems, ACM SIGPLAN Notices*, Volume 27, Number 9, pages 187–197, 1992.
- [18] J. JáJá. *An Introduction to Parallel Algorithms*. Addison-Wesley Publishing Company, Reading, 1992.
- [19] B. H. H. Juurlink and H. A. G. Wijshoff. The parallel hierarchical memory model. In *Proceedings of the 4th Scandinavian Workshop on Algorithm Theory*, Lecture Notes in Computer Science 824, pages 240–251, Springer-Verlag, Berlin/Heidelberg, 1994.
- [20] D.E. Knuth. *The Art of Computer Programming*, Volume 3/ *Sorting and Searching*. Addison-Wesley Publishing Company, Reading, 1973.
- [21] K. Krueger, D. Loftesness, A. Vahdat, and T. Anderson. Tools for the development of application-specific virtual memory management. In *Proceedings of the 8th Annual Conference on Object-Oriented Programming Systems, Languages, and Applications, ACM SIGPLAN Notices*, Volume 28, Number 10, pages 48–64, 1993.
- [22] A. LaMarca and R.E. Ladner. The influence of caches on the performance of sorting. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 370–379, ACM, New York and SIAM, Philadelphia, 1997.
- [23] D. McNamee and K. Amstrong. Extending the Mach external pager interface to accommodate user-level block replacement policies. Technical Report 90-09-05, Department of Computer Science and Engineering, University of Washington, Seattle, 1990.
- [24] D. Naor, C.U. Martel, and N.S. Matloff. Performance of priority queue structures in a virtual memory environment. *The Computer Journal*, Volume 34, pages 428–437, 1991.
- [25] M.H. Nodine and J.S. Vitter. Large-scale sorting in parallel memories, In *Proceedings of the 3rd ACM Symposium on Parallel Algorithms and Architectures*, ACM Press, New York, pages 29–39, 1991.
- [26] Y.N. Patt. Guest editor’s introduction: The I/O subsystem — A candidate for improvement. *IEEE Computer*, Volume 27, Number 3, pages 15–16, 1994.
- [27] D.A. Patterson and J.L. Hennessy, *Computer Organization & Design: The Hardware/Software Interface*. Morgan Kaufmann Publishers, San Francisco, 1994.
- [28] R. Sedgewick. *Algorithms*. Addison-Wesley Publishing Company, Reading, 1983.
- [29] M. Thorup. On RAM priority queues. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 59–67, ACM, New York and SIAM, Philadelphia, 1996.
- [30] J.S. Vitter and E. A.M. Shriver. Algorithm for parallel memory, I: Two-level memories. *Algorithmica*, Volume 12, pages 110–147, 1994.

- [31] J. Vuillemin. A data structure for manipulating priority queues. *Communications of the ACM*, Volume 21, pages 309–315, 1978.
- [32] L. M. Wegner and J. I. Teuhola. The external heapsort. *IEEE Transactions on Software Engineering*, Volume 15, pages 917–925, 1989.
- [33] J. W. J. Williams. Algorithm 232, Heapsort. *Communications of the ACM*, Volume 7, pages 347–348, 1964.