

# Pairing Heaps: Experiments and Analysis

JOHN T. STASKO and JEFFREY SCOTT VITTER

**ABSTRACT:** *The pairing heap has recently been introduced as a new data structure for priority queues. Pairing heaps are extremely simple to implement and seem to be very efficient in practice, but they are difficult to analyze theoretically, and open problems remain. It has been conjectured that they achieve the same amortized time bounds as Fibonacci heaps, namely,  $O(\log n)$  time for delete and delete\_min and  $O(1)$  for all other operations, where  $n$  is the size of the priority queue at the time of the operation. We provide empirical evidence that supports this conjecture. The most promising algorithm in our simulations is a new variant of the twopass method, called auxiliary twopass. We prove that, assuming no decrease\_key operations are performed, it achieves the same amortized time bounds as Fibonacci heaps.*

## 1. INTRODUCTION

A priority queue is an abstract data type for maintaining and manipulating a set of items based on priority [1]. Priority queues derive great theoretical

Support was provided in part by NSF research grant DCR-84-03613, an NSF Presidential Young Investigator Award, an IBM Faculty Development Award, and a Guggenheim Fellowship.

Part of this research was performed at Mathematical Sciences Research Institute, Berkeley, Calif., and the Institut National de Recherche en Informatique et en Automatique, Rocquencourt, France.

© 1987 ACM 0001-0782/87/0300-0234 75¢

and practical importance from their use in solving a wide range of combinatorial problems, including job scheduling, minimal spanning tree, shortest path, and graph traversal.

Priority queues support the operations *insert*, *find\_min*, and *delete\_min*; additional operations often include *decrease\_key* and *delete*. The *insert*( $t, v$ ) operation adds item  $t$  with key value  $v$  to the priority queue. The *find\_min* operation returns the item with minimum key value. The *delete\_min* operation returns the item with minimum key value and removes it from the priority queue. The *decrease\_key*( $t, d$ ) operation reduces item  $t$ 's key value by  $d$ . The *delete*( $t$ ) operation removes item  $t$  from the priority queue. The *decrease\_key* and *delete* operations require that a pointer to the location in the priority queue of item  $t$  be supplied explicitly, since priority queues do not support searching for arbitrary items by value. Some priority queues also support the *merge* operation, which combines two item-disjoint priority queues.

We will concentrate on the *insert*, *delete\_min*, and *decrease\_key* operations because they are the operations that primarily distinguish priority queues from other set manipulation algorithms and because they are the critical operations as far as the time bounds are concerned.

Several implementations of priority queues, such as implicit heaps [10], leftist heaps [3, 7], and binomial heaps [2, 9] have been shown to exhibit an  $O(\log n)$  worst-case time bound for all operations, where  $n$  is the size of the priority queue at the time of the operation. Fibonacci heaps [4] provide a dramatic improvement on the general logarithmic bound by achieving amortized time bounds of  $O(1)$  for *insert*, *decrease\_key*, and *find\_min* and  $O(\log n)$  for *delete\_min* and *delete*. This greatly improves the best known theoretical bounds for the time required to solve several combinatorial problems.<sup>1</sup> Following the approach of [8], a sequence of operations  $op_1, op_2, \dots, op_k$  is said to have *amortized time bounds*  $b_1, b_2, \dots, b_k$  if

$$\sum_{1 \leq i \leq j} t_i \leq \sum_{1 \leq i \leq j} b_i, \quad \text{for all } 1 \leq j \leq k,$$

where  $t_i$  is the actual time used by  $op_i$ . Intuitively, if operation  $op_i$  uses less time than its allotted  $b_i$  units, then the leftover time may be held in reserve to be used by later operations.

Fibonacci heaps achieve their time bounds by complicated invariants with significant overhead, so they are not the method of choice in practice. Recently, a self-adjusting data structure called the *pairing heap* was proposed [5]. Pairing heaps are much simpler than Fibonacci heaps, both conceptually and in implementation; and they have less overhead per operation. The best amortized bound proved so far for pairing heaps is  $O(\log n)$  time per operation. It is conjectured that pairing heaps achieve the same amortized time bounds as Fibonacci heaps, namely,  $O(1)$  per operation except  $O(\log n)$  for *delete\_min* and *delete*.

To test whether the conjecture is true, we performed several simulations of the pairing heap algorithms. These simulations differed significantly from the ones independently done in [6], since the latter ones did not address the conjecture. In our simulations, we tested several different pairing heaps and used “greedy” heuristics and the appropriate se-

quences of commands to make the pairing heaps perform as poorly as we could. The results were positive in that the pairing heaps always performed extremely well. This does not prove that the desired time bounds do hold, but it is reassuring and makes us optimistic that the conjecture is true.

In this article we study the “twopass” and “multipass” versions of pairing heaps; the names arise from the method used to do the *delete\_min* in each version [5]. We also introduce new variants called “auxiliary twopass” and “auxiliary multipass.” All versions are described in the next section. In Section 3, we discuss our simulations and the empirical data. Auxiliary twopass performed best in the simulations, based on our measure of performance. In Section 4, we provide a partial theoretical analysis of pairing heaps by introducing the concept of “batched potential.” We show, for example, that auxiliary twopass uses  $O(1)$  amortized time per *insert* and *find\_min* and  $O(\log n)$  amortized time for the other operations. Conjectures and open problems follow in Section 5.


## 2. PAIRING HEAP ALGORITHMS

A comprehensive description of pairing heaps appears in [5]. A summary is given below. Our studies involve the twopass algorithm, which was the subject of most of the analysis in [5], and the multipass algorithm.


Pairing heaps are represented by heap-ordered trees and forests. The key value of each node in the heap is less than or equal to those of its children. Consequently, the node with minimum value (for simplicity, we will stop referring to a key value, and just associate the value directly with the heap node) is the root of its tree. Groups of siblings, such as tree roots in a forest, have no intrinsic ordering.

In the general sense, pairing heaps are represented by multiway trees with no restriction on the number of children that a node may have. Because this multiple child representation is difficult to implement directly, the child-sibling binary tree representation of a multiway tree is used, as illustrated in Figure 1 (p. 236). In this representation, the left pointer of a node accesses its first child, and the right pointer of a node accesses its next sibling. In terms of the binary tree representation, it then follows that the value of a node is less than or equal to all the values of nodes in its left subtree. A third pointer, to the previous sibling, is also included in each node in order to facilitate the *decrease\_key* and *delete* operations. The number of pointers can be reduced from three to two, as explained in [5] at the expense of a constant factor increase in running time. Unless stated otherwise, the terms “child,” “parent,” and “subtree” will be used in the multiway tree sense;

<sup>1</sup> For example, a standard implementation of Dijkstra’s algorithm (which finds the shortest path from a specified vertex  $x$  to all other vertices in a graph with nonnegative edge lengths) uses a priority queue as follows: Let us denote the number of vertices in the graph by  $V$  and the number of edges by  $E$ . The key value of each item  $y$  in the priority queue represents the length of the shortest path from vertex  $x$  to vertex  $y$  using only the edges in the graph already processed. Initially, no edges are processed, and the priority queue contains  $V$  items; the key value of item  $x$  is 0 and all other items have key value  $\infty$ . The algorithm successively performs *delete\_mins* until the priority queue is empty. Each time a *delete\_min* is performed (say, the vertex  $y$  is deleted), the algorithm outputs the shortest path between  $x$  and  $y$ , and each unprocessed edge  $(y, z)$  incident to  $y$  in the graph is processed; this may require that a *decrease\_key* be performed in order to lower the key value of  $z$  in the priority queue. Thus, there are at most  $V$  *inserts*,  $V$  *delete\_mins*, and  $E$  *decrease\_keys* during the course of the algorithm. If a Fibonacci heap is used to implement the priority queue, the resulting running time is  $O(E + V \log V)$ , which is a significant improvement over  $O((E + V) \log V)$  using the other heap representations. Other examples of how Fibonacci heaps can improve worst-case running times are given in [4].



(a)



(b)

FIGURE 1. An Example of a Heap Ordered Tree (a) Multiway tree heap representation. (b) Corresponding binary tree representation.

their corresponding meaning in the binary tree representation should be clear.

The primary action performed in pairing heap operations is a *comparison-link*, in which the values of two nodes are compared. The node with larger value is demoted in the sense that it becomes the first child of the smaller-valued node. The previous first child of the smaller node becomes the second child, the previous second child becomes the third child, and so on. Ties can be broken arbitrarily. The binary tree representation of the comparison-link is given in Figure 2. This comparison-link action is performed repeatedly during the *delete\_min* operation of a priority queue. It is the primary action that we seek to minimize to reduce execution times.

The *twopass algorithm* that we examined was the variant that yielded the  $O(\log n)$ -time amortized bounds for *insert*, *decrease\_key*, and *delete\_min* in [5]. Only one tree is maintained. Hence, in the binary

tree representation, the root node always has a null right pointer. The *insert*( $t, v$ ) operation performs a comparison-link between  $t$  and the tree root; the node with smaller value becomes the root of the resulting tree. The *decrease\_key*( $t, d$ ) operation begins by reducing  $t$ 's value by  $d$ . This means that  $t$  may now have a value smaller than its parent. Consequently, it must be removed from the tree (with its own subtree intact) and comparison-linked with the tree root. Again, the node with smaller value becomes the root of the resulting tree.

The *delete\_min* operation gives the twopass algorithm its name. First, the tree root node is deleted and its value returned. This leaves a forest of former children and their subtrees. Next, two comparison-linking passes are made over the roots of this forest. Pass 1 is a left-to-right pass, in which a comparison-link is performed on successive pairs of root nodes. Pass 2 then proceeds from right-to-left. In each step, the two rightmost trees are replaced by the tree resulting from a comparison-link; the "cumulative" rightmost tree is continually updated in this manner until it is the only remaining tree. The root of this final tree is the minimum of all the nodes in the tree. Figure 3 illustrates the *delete\_min* procedure in terms of the binary tree representation.

The *delete*( $t$ ) operation works as follows: If the node  $t$  to be deleted is the root of the main tree, then a *delete\_min* operation is performed. Otherwise,  $t$  is deleted from the tree. The former subtrees of  $t$  are recombined into a single tree via the twopass linking procedure. This tree is then comparison-linked to the root of the main tree.

The *multipass algorithm* that we studied was also presented in [5]. Both the *insert*( $t, d$ ) and *decrease\_key*( $t, d$ ) operations function exactly as those in the twopass algorithm. The *delete\_min* operation, however, distinguishes multipass from twopass. The first operation in the multipass *delete\_min* is the deletion of the root node; its value is returned. This leaves the heap with some number of trees, say  $r$ . Next, we repeatedly perform pairwise linking passes on the roots of these trees until the heap is left with only one tree. Each comparison-link reduces the number of trees by one, and each pass cuts the number of trees roughly in half. For  $r$  trees, a total of  $\lceil \log r \rceil^2$  passes are made. A simpler heuristic is to comparison-link the first two tree roots and place the "winning" root (smaller value) at the tail of the forest list. Alternatively, a circular list could be used to store the siblings. Both ways, a round robin effect emerges, and we see that for  $r$  tree roots, exactly  $r - 1$  link operations are performed. Following the linking phase, the heap is again left with a single tree; its root is the

<sup>2</sup> All logarithms in this article are base 2.




FIGURE 2. Two Binary Tree Heap Configurations and the Resulting Structures From a Comparison-Link Between Nodes X and Y




FIGURE 3. Twopass Delete Min Procedure, Using the Binary Tree Representation. (a) Beginning heap configuration. (b) Root (minimum) deletion. (c) Pass one. (d) Pass two.

node with minimum key value. Figure 4 illustrates the multipass *delete\_min* in terms of the binary tree representation. The *delete(t)* operation is the same as in twopass except that multipass comparison-linking is used on the children of the deleted node *t*.

While working with these algorithms, we designed two new variations, that we call *auxiliary twopass*

and *auxiliary multipass*. Auxiliary twopass works as follows: In addition to the main tree in the heap, we maintain an auxiliary area that consists of an ordered list of other trees. It is convenient in the implementation to store the auxiliary area as the right subtree (in the binary tree sense) of the root.

The *insert(t, v)* and *decrease\_key(t, d)* operations





FIGURE 4. Multipass Delete Min Procedure, Using the Binary Tree Representation. (a) Beginning heap configuration. (b) Root (minimum) deletion. (c) Pass one. (d) Pass two. (e) Pass three.

function as in the regular twopass algorithm except for one major difference. Rather than comparison-linking node  $t$  with the tree root, the node is added to the end of the list of auxiliary trees. (As usual, in the case of a *decrease\_key* operation, the subtree rooted at  $t$  remains intact.) If the *find\_min* operation is to be implemented, a separate minimum pointer

must be maintained; each *insert* or *decrease\_key* node must be checked against the minimum pointer so that the pointer can be updated if necessary.

The *delete\_min* operation, which is illustrated in Figure 5, begins by “batching” the auxiliary area, that is, by running the multipass pairing procedure on the auxiliary area. (Note that although the



**FIGURE 5. Auxiliary Twopass Delete Min Procedure, Using the Binary Tree Representation. (a) Initial heap configuration. (b) Multipass on auxiliary forest. (c) Link auxiliary root to main root. (d) Root (minimum) deletion. (e) Twopass back to one root.**

method described is called auxiliary twopass, the auxiliary area is linked together using the multipass method.) When this linking is complete, the auxiliary area consists of a single tree. If the auxiliary area originally consists of  $2^k$  singleton trees, for some  $k \geq 0$ , the resulting tree is a binomial tree [2, 9]. The next action is a comparison-link between the main tree root and the new auxiliary root. After this comparison-link, the heap again contains only one root node, that of minimum key value. From this point on, the *delete\_min* operation proceeds exactly as in the twopass algorithm. The root node is removed, and we link the remaining forest of trees via the twopass procedure.

The rationale for maintaining the auxiliary area is that it prevents many comparisons between single nodes from an *insert* and large trees already in the forest. We shall prove in Section 4 that if there are no *decrease\_key* operations, auxiliary twopass achieves the amortized time bounds of  $O(1)$  for *insert* and  $O(\log n)$  for *delete\_min*.

The *delete(t)* operation works as in twopass. If node  $t$  is the current minimum, then a *delete\_min* is performed. Otherwise, the children of  $t$  are recombined into a single tree, which is then comparison-linked to the main tree.

Auxiliary multipass is identical except that the multipass algorithm is used on the regular tree. The auxiliary area is still batched using multipass. Section 4 derives slightly weaker amortized time bounds than for auxiliary twopass, under the assumption that no *decrease\_key* operations are performed:  $O(1)$  per *insert* and  $O((\log n \log \log n)/\log \log n)$  per *delete\_min*.

Lazy variants of these algorithms are also possible, in which the heap consists of a forest of trees rather than a single tree. One possible implementation is described in [5]. However, extra comparisons other than in comparison-link actions are required to implement *find\_min*, since the heap no longer has a single root. As a result, the *find\_min* operation in the lazy variants often does some restructuring of the tree. The *find\_min* operation for the auxiliary variants can be done in constant time, since a pointer to the current minimum node can easily be maintained during *inserts* and *decrease\_keys*; the extra comparison to do this is balanced by the fact that *inserts* and *decrease\_keys* do not perform any comparison-links. To make our simulation results of *insert*, *decrease\_key*, and *delete\_min* fair, we have excluded lazy variants from our study. Their performances are similar.

Although the twopass algorithm has provided the fastest general amortized time bounds so far, our intuition suggested that the multipass variants

should run faster. Although all make roughly the same number of comparison-link actions on a similar heap configuration, the multipass variants tend to build a more “structured” forest configuration over time. All the uppermost nodes that are directly involved in link actions will be formed into a binomial-like tree, which helps limit the number of links during subsequent *delete\_min* operations. The simulation results described in the next section are somewhat surprising; auxiliary twopass consistently outperforms the multipass versions.

### 3. SIMULATIONS

Our test simulations of the pairing heap algorithms consisted of structured sets of *insert*, *decrease\_key*, and *delete\_min* operations. No key values were ever assigned to nodes. Instead, we used a “greedy” heuristic to determine the winners of comparisons, in hopes of causing a worst-case scenario. Every time a comparison-link operation was performed, the node with more children won the comparison; that is, it was judged to have the smaller key value. This node gained one child in the link operation. Our greedy approach allowed us to keep the nodes with many children at the uppermost levels in the heap. Since the number of children at the root level determines how much work a *delete\_min* performs, this greedy approach forced the priority queue to do significantly more work than would have been the case if the key values were assigned randomly. Two methods were used for determining which nodes to use for *decrease\_key* operations: In a *random decrease\_key*, we chose a nonroot node at random. In a *greedy decrease\_key*, we used the greedy heuristic and chose the node with the most children, subject to the constraint that the node could not be the root or a child of the root.

The binary tree representation of a multiway tree was used to implement the pairing heaps. Heap nodes were implemented as record structures with left (first child) and right (next sibling) pointer fields.

In this section, we report on nine simulations of twopass and multipass pairing heap algorithms. Each simulation consisted of several *phases*. A phase consisted of some set of *inserts* and *decrease\_keys* followed by a *delete\_min*. In the descriptions that follow,  $n$  refers to the size of the priority queue at the beginning of the phase. The phases of the nine simulations consisted of, respectively:

- (1)  $\log n$  *inserts*, followed by one *delete\_min*.
- (2)  $0.5 \log n$  (*insert*, random *decrease\_key*) pairs, followed by one *delete\_min*.
- (3)  $0.5 \log n$  (*insert*, greedy *decrease\_key*) pairs, followed by one *delete\_min*.

- (4) one *insert*, then  $x(\log n) - 1$  greedy *decrease\_keys*, followed by one *delete\_min*, for  $x = 0.25$ , with an initial binomial tree of size  $2^{12}$ .
- (5) same for  $x = 1.0$ .
- (6) same for  $x = 4.0$ .
- (7) one *insert* then  $x(\log n) - 1$  greedy *decrease\_keys*, followed by one *delete\_min*, for  $x = 0.22$ , with an initial binomial tree of size  $2^{18}$ .
- (8) same for  $x = 1.0$ .
- (9) same for  $x = 4.0$ .

Our measure of performance compared the actual work done by each algorithm with an allowance for the operations processed. The actual work done was considered to be one unit for an *insert*, one unit for a *decrease\_key*, and one unit for the delete plus one unit for each comparison-link that occurred during a *delete\_min*. Allowances for the operations corresponded to the amortized time bounds sought for them: the *insert* and *decrease\_key* allowances were each one unit, and the *delete\_min* allowance was  $\log n$ , where  $n$  was the heap size at the time of the *delete\_min*.

Phases were grouped into a smaller number of increments to facilitate graphical display of the results; in the first three simulations the size of the heap grew by a fixed amount in each increment, and in Simulations 4–9 (Figures 9–14.) increments consisted of a fixed number of phases. For each increment in a simulation, we calculated its *work ratio*. The work ratio is defined as actual work performed divided by the operations' allowances. By seeing how the work ratio changed over time across these increments, we were able to judge the performances of the algorithms. If the work ratio increased over time, the  $O(1)$ -time *insert* and *decrease\_key* allowances and the  $O(\log n)$ -time *delete\_min* allowance would not be bounding the actual work growth. If instead the work ratio stayed constant or decreased, then the experiments would provide encouragement that the sought for amortized time bounds are possible.

In our simulations, we chose the particular order and frequency of operations so that if any of the conjectured amortized time bounds of  $O(1)$  for *insert* and *decrease\_key* and  $O(\log n)$  for *delete\_min* did not hold, we would detect a discrepancy in the results. The simulations of pairing heaps in [6], on the other hand, were limited for several reasons: First, no *decrease\_key* operations were performed. We will see in Section 4 that if no *decrease\_keys* are done, we can prove the conjectured bounds analytically for auxiliary twopass. More importantly, however, the simulations performed the same number of *inserts* as *delete\_mins*. It is already known from [5] that each

operation can be done in  $O(\log n)$  amortized time. Therefore, it was impossible in Jones's [6] simulations to distinguish between  $O(1)$  time and  $O(\log n)$  time per *insert*. Our simulations, on the other hand, tested the time bound conjectured more effectively by performing  $O(\log n)$  *insert* and *decrease\_key* operations for each *delete\_min*.

Simulation 1 allowed us to examine how the heaps performed when no *decrease\_key* operations were used. Work ratios were calculated over increments of 12,000 nodes of heap growth, with the heap size eventually reaching 1,200,000 nodes. Four initial *inserts* were used to "start-up" the simulation. The results are graphed in Figure 6 (p. 242). Only multipass showed a steady increase in work ratio; however, there was a marked decrease near the simulation's end. Both twopass and auxiliary twopass remained mostly steady, and auxiliary multipass exhibited a clear decrease. Auxiliary twopass was the fastest algorithm, a fact that would continue through most of the following simulations. It is interesting to note that the auxiliary algorithms were not subject to wide fluctuations in work ratio as were the regular algorithms.

Simulation 2 utilized random *decrease\_key* operations primarily to see how random disruptions in the heap structure would affect overall algorithm performances. Again work ratios were calculated in 12,000 node increments, and the total heap size grew to 1,200,000 nodes. Sixteen *inserts* were used to initialize the simulation. This simulation's results, which are shown in Figure 7 (p. 242), were quite similar to those of Simulation 1. No steady work ratio increases were evident, nor was there an appreciable gain in work ratio values from Simulation 1. Both multipass algorithms exhibited work ratio decreases, with regular multipass remaining slightly superior. Auxiliary twopass was again the fastest algorithm, and twopass was the slowest. Curiously, the total work ratio for auxiliary twopass over the entire simulation was slightly less than its total in Simulation 1. In essence, the random *decrease\_keys* helped the algorithm run faster.

Simulation 3 utilized greedy *decrease\_key* operations in which the node with the most children was chosen for the operation. Nodes such as the root and children of the root, whose choice would have no effect on the heap structure, were excluded from being candidates. This simulation's *decrease\_key* operation was intended to move nodes with many children from the central heap up to the top root level, thereby forcing the *delete\_min* operations to do even more work. Because of the extra storage required, work ratios were calculated in increments of 6,000 nodes. The final heap size was 546,000 nodes. A



SIMULATION 1




FIGURE 6.  $\log n$  inserts, 1 delete min per phase

SIMULATION 2




FIGURE 7.  $0.5 \log n$  (insert, random decrease key) pairs, 1 delete min per phase

start-up set of sixteen *insert* operations was used.

After showing small jumps in the work ratios, all four algorithms maintained steady levels at the simulation's end. The results are graphed in Figure 8. Clearly, the greedy *decrease\_key* operations did have an effect, as work ratio values were higher than those in the first two simulations. Auxiliary twopass was again the fastest algorithm, twopass was the slowest, and the multipass algorithms were quite similar. The general multipass algorithm had a definite superiority at smaller heap sizes, however.

Simulations 4–9 primarily examined how the *decrease\_key* operation affected algorithm performances. An initial binomial tree of some size was built. This provided all four algorithms with the same starting point, so no initial bias was introduced. Each phase contained only one *insert* operation; hence, a constant heap size was maintained, that of the initial binomial tree. By varying the number of *decrease\_key* operations between the *insert* and the *delete\_min* in a phase, we could see exactly how this number affected the work ratio.

In Simulations 4–6, we used an initial binomial tree size of  $n = 2^{12} = 4096$  nodes. Simulations 4, 5, and 6 used  $0.25 \log n = 3$ ,  $\log n = 12$ , and  $4 \log n = 48$  *decrease\_key* operations per phase, respectively. We grouped 100 phases into an increment for work ratio calculations, then we ran the simulation for

100 increments. The respective results are shown in Figures 9–11 (pp. 244–245). All three tests showed very steady work ratio rates, which is encouraging. The actual values were quite different, however. The twopass variants had lower work ratios when there were a smaller number of *decrease\_key* operations per phase, whereas the multipass variants exhibited an opposite behavior; they performed better as the number of *decrease\_keys* per phase increased. Auxiliary twopass was overall the best algorithm, but regular twopass exhibited a curious variation on its usual slowest behavior. In the two simulations with more *decrease\_key* operations, twopass was clearly slowest. In fact, in Simulation 6 with 48 *decrease\_key* operations per phase, twopass was blatantly behind the other three algorithms performances. But in Simulation 4 with the fewest (three) *decrease\_keys* per phase, twopass was the fastest algorithm! It appears that twopass performs best when there are relatively few *insert* and *decrease\_key* operations compared to the number of *delete\_min* operations.

In Simulations 7–9, we used a much larger initial binomial tree size of  $n = 2^{18} = 262,144$  nodes. Simulations 7–9 used  $0.22 \log n = 4$ ,  $\log n = 18$ , and  $4 \log n = 72$  *decrease\_key* operations per phase, respectively. We grouped 1000 phases into an increment for work ratio calculations, then we ran the simula-




FIGURE 8.  $0.5 \log n$  (insert, greedy decrease key) pairs, 1 delete min per phase

SIMULATION 4




FIGURE 9. 1 insert,  $(0.25 \log n) - 1$  greedy decrease keys, 1 delete min per phase

SIMULATION 5




FIGURE 10. 1 insert,  $(\log n) - 1$  greedy decrease keys, 1 delete min per phase

## SIMULATION 6




FIGURE 11. 1 insert,  $(4.0 \log n)$  — 1 greedy decrease keys, 1 delete min per phase

tion for 100 increments. The results are shown in Figures 12–14 (pp. 246–247). The relative performances of the four algorithms were quite similar to those of Simulations 4–6. The curious behavior of the twopass algorithm was again evident, as it performed poorly with many *decrease\_key* operations per phase, but improved dramatically with few. With the larger initial heap size, however, it never overtook auxiliary twopass as the fastest algorithm, as it did in Simulation 4.

In the last six simulations with constant heap size, the work ratio during the formation of the initial binomial tree was just under 2.0; we performed  $n$  *insert* operations, followed by  $n - 1$  comparison-links. Therefore, actual work was  $2n - 1$  units, and the *insert* allowance was  $n$  units, giving a work ratio of  $2 - 1/n$ . This amount was included in the total work ratio for the simulation.

Aside from Simulations 1–9, we also performed two randomized simulations to verify that our data were not dependent on the fixed structure of each phase. In the first, we kept the heap size constant as in Simulations 4–9, but the number of *decrease\_keys* per phase was uniformly distributed between 0 and  $2 \log n$ . The second began by performing 64 initial *inserts* which were followed by a random sequence of *insert*, *greedy decrease\_key*, and *delete\_min* operations, all having the same probability of occurrence.

Both results were consistent with those above; work ratio values stayed steady or showed a small decrease, and auxiliary twopass again exhibited the lowest overall work ratios. In the random sequence simulations, however, twopass and auxiliary twopass performed almost identically.

The data from these simulations allowed us to make the following conclusions: First, the  $O(1)$ -time bounds for *insert* and *decrease\_key* and the  $O(\log n)$ -time bound for *delete\_min* appear to hold in the amortized sense. Our data provided no evidence to the contrary. In fact, they provide some clue as to the actual coefficients implicit in the big-oh terms. Let us make the simplifying assumption that the amortized running time for each *insert* and *decrease\_key* operation in the simulations is  $c$  time units and that the amortized time per *delete\_min* is  $d \log n$  units. Solving a set of linear equations obtained from Simulations 4–9 gives  $c \approx 2.9$  and  $d \approx 1.5$  for twopass,  $c \approx 1.8$  and  $d \approx 2.2$  for multipass,  $c \approx 2.0$  and  $d \approx 1.7$  for auxiliary twopass, and  $c \approx 1.9$  and  $d \approx 2.3$  for auxiliary multipass. Second, the auxiliary twopass algorithm was clearly the best overall. It typically exhibited lower work ratios than the other three algorithms. Third, adding the auxiliary area to multipass caused no great improvements to the algorithm. In our tests, the multipass algorithm was almost always superior to its auxiliary variant. Finally,

SIMULATION 7




FIGURE 12. 1 insert,  $(0.22 \log n) - 1$  greedy decrease keys, 1 delete min per phase

SIMULATION 8




FIGURE 13. 1 insert,  $(\log n) - 1$  greedy decrease keys, 1 delete min per phase

the regular twopass algorithm's performance was quite variable. It was often the worst, especially when many *insert* and *decrease\_key* operations were processed. As the number of these operations declined, however, its performance improved to rival that of auxiliary twopass.

#### 4. BATCHED POTENTIAL

The best known amortized time bounds of the *insert*, *decrease\_key*, and *delete\_min* operations for the twopass pairing heap are all  $O(\log n)$ , due to [5]. To equal the time bounds for Fibonacci heaps, the *insert* and *decrease\_key* time bounds must be shown to be  $O(1)$ . In an effort to prove the constant time bounds for pairing heaps, we will use the auxiliary twopass algorithm and introduce *batched potential*. But before that, let us briefly review the concept of potential as it applies to amortized algorithmic analysis.

The potential technique for amortized analysis is discussed in [8]. Each configuration of the pairing heap is assigned some real value  $\Phi$ , known as the "potential" of that configuration. For example, one could define the potential of a pairing heap configuration to be the number of trees it contains. For any sequence of  $n$  operations, the amortized time of the  $i$ th operation is defined to be the actual running

time of the operation plus the change in potential, namely,  $t_i + \Phi(i) - \Phi(i - 1)$ , where  $t_i$  is the actual time of the  $i$ th operation,  $\Phi(i)$  is the potential after the  $i$ th operation, and  $\Phi(i - 1)$  is the potential before the  $i$ th operation. If we start with potential 0 and end up with positive potential, then the total running time is bounded by the total amortized time, via the telescoping effect of the potential changes.

#### THEOREM 1.

*The auxiliary twopass pairing heap algorithm achieves amortized time bounds of  $O(1)$  for insert and find\_min and  $O(\log n)$  for delete\_min and delete if no decrease\_key operations are allowed.*

#### PROOF.

Without loss of generality, we can restrict ourselves to *insert* and *delete\_min* operations exclusively. Let us define the rank of a node to be the binary logarithm of the number of nodes in the subtree rooted there (in the binary tree sense). We use a variant of the potential function used to analyze twopass in [5]. We define the potential  $\Phi$  of a heap configuration to be the sum of the ranks of all nodes in the main tree (that is, not counting the auxiliary area) plus 5 times the number of roots in the auxiliary area.

Each *insert* places a new node into the auxiliary




FIGURE 14. 1 insert,  $(4.0 \log n) - 1$  greedy decrease keys, 1 delete min per phase

area and increases the potential by 5; its amortized time is thus 6. There is no large single change in potential until a *delete\_min* is performed, when the nodes in the auxiliary area are added to the main tree. We refer to this as “batched potential” because in effect changes in potential are not considered until a *delete\_min* is performed.

Let us consider the case in which the auxiliary area is nonempty when a *delete\_min* takes place. We let  $i > 0$  denote the number of (root) nodes in the auxiliary area. The multipass linking spends  $i - 1$  units of work building these  $i$  nodes into a single tree, reducing the potential by  $5i - 5$ . We can show that the sum of the ranks of the nodes in the resulting auxiliary tree is bounded by  $4i - 4$ , as follows: If  $i$  is a power of 2, then the auxiliary tree in the binary sense has a complete left subtree of size  $i - 1$  and no right subtree. Since the number of nodes on descending levels of the binary tree doubles as subtree sizes are roughly halved, the sum of the ranks is

$$\log i + \sum_{0 \leq j < \log i} 2^j \log \left( \frac{i}{2^j} - 1 \right).$$

Now let  $m = \log i$  and simplify. The sum is bounded by

$$\begin{aligned} m + m \sum_{0 \leq j \leq m} 2^j - \sum_{0 \leq j \leq m} j 2^j \\ = m + m(2^{m+1} - 1) - m2^{m+2} + (m + 1)2^{m+1} - 2 \\ = 2^{m+1} - 2 \\ = 2i - 2 \end{aligned}$$

If  $i$  is not a power of 2, then we can append  $\leq i - 2$  extra dummy root nodes to the auxiliary area so that there are  $2^{\lceil \log i \rceil}$  root nodes. It is straightforward to show by induction that the tree resulting from the multipass linking of the auxiliary area without the dummy nodes can be “embedded” in the tree resulting from the multipass linking of the auxiliary area with the dummy nodes. By the analysis given above for the case when  $i$  is a power of 2, the sum of the ranks is bounded by  $2(2i - 2) - 2 + 2\delta_{i=1} = 4i - 6 + 2\delta_{i=1} \leq 4i - 4$ . (The Kronecker delta  $\delta_{i=1}$  denotes 1 if  $i = 1$  and 0 otherwise.) We can get a better bound by considering the contribution of the dummy nodes, but for our purposes this bound is adequate, since we are ignoring constant factors.

When the auxiliary tree is linked with the main tree (so that all the nodes are in the main tree and the auxiliary area is empty), one unit of work is expended, and the change in potential is bounded by  $\log n + 4i - 4 - 5$ , where  $n$  is the number of nodes in the priority queue. Next, the root is deleted via

one unit of work, reducing the potential by  $\log n$ , and the resulting subtrees are recombined via the twopass scheme. If there are  $k$  subtrees, a total of  $k - 1$  units of work are used to recombine; the resulting potential increase is shown in [5] to be bounded by  $2 \log n - k + 3$ . The total amount of work spent during the *delete\_min* and the net change in potential can thus be bounded by  $i + k$  and  $2 \log n - i - k - 1$ , respectively, which bounds the amortized time by  $2 \log n - 1$ .

The other case to consider is when the auxiliary area is empty at the time of a *delete\_min*, that is, when  $i = 0$ . In that case, the reasoning in the last paragraph shows that the total amount of work and the net change in potential are bounded by  $k$  and  $\log n - k + 3$ , respectively. This completes the proof.  $\square$

The same approach combined with the analysis in [5] proves the following:

**THEOREM 2.**

*The auxiliary multipass algorithm achieves amortized time bounds of  $O(1)$  for insert and find\_min and  $O((\log n \log \log n) / \log \log \log n)$  for delete\_min and delete if no decrease\_key operations are allowed.*

Unfortunately, we cannot as yet extend either algorithm’s analysis to include *decrease\_key* nodes in the auxiliary area. The problem lies in the subtrees attached to nodes whose values are decreased. We are hopeful that some variant of this batching technique will prove that the  $O(1)$ -time amortized bounds for *insert* and *decrease\_key* and the  $O(\log n)$ -time amortized bound for *delete\_min* do hold. Note that we can get a weaker result by a slight modification of our algorithms. If the auxiliary area is batched whenever a *decrease\_key* or *delete\_min* is performed, then get the desired bounds, except that *decrease\_key* and *delete\_min* use  $O(\log n)$  time for auxiliary twopass and  $O((\log n \log \log n) / \log \log \log n)$  time for auxiliary multipass.

**5. CONCLUSIONS AND OPEN PROBLEMS**

The experimental data gathered from our simulations provide empirical evidence that the  $O(1)$ -time bounds for *insert*, *decrease\_key*, and *find\_min* and the  $O(\log n)$ -time bounds for *delete\_min* and *delete* do hold in the amortized sense, where  $n$  is the size of the priority queue at the time of the operation. All the pairing heap methods performed well in our simulations. The auxiliary twopass variant clearly did the best. This result is satisfying because we have shown that the auxiliary twopass algorithm achieves the above mentioned bounds, assuming that no *decrease\_key* operations occur. Or if the auxiliary area is batched and merged with the main tree

whenever a *decrease\_key* or *delete\_min* operation is done, then *insert* and *find\_min* have  $O(1)$ -time bounds and all other operations have  $O(\log n)$ -time bounds in the amortized sense.

Proving amortized time bounds of  $O(1)$  for *insert* and *decrease\_key* and  $O(\log n)$  for *delete\_min* for some pairing heap implementation remains the major open problem. [5] has shown that an  $O(\log n)$ -time amortized bound can be proven for all operations of the twopass variant. No multipass variant has been proven to achieve the same  $O(\log n)$ -time amortized bound for all operations; *delete\_min* remains the primary stumbling block. Demonstrating this logarithmic time bound for all multipass operations is another interesting open problem.

## REFERENCES

1. Aho, A.V., Hopcroft, J.E., and Ullman, J.D. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, Reading, Mass., 1974.
2. Brown, M.R. Implementation and analysis of binomial queue algorithms. *SIAM J. Comput.* 7 (Aug. 1978), 298–319.
3. Crane, C.A. Linear lists and priority queues as balanced binary trees. Tech. Rep. STAN-CS-72-259, Dept. of Computer Science, Stanford University, Feb. 1972.
4. Fredman, M.L., and Tarjan, R.E. Fibonacci heaps and their uses in improved network optimization algorithms. In *Proceedings of the 25th Annual Symposium on Foundations of Computer Science*, (West Palm Beach, Fla., Oct. 1984), 338–344.
5. Fredman, M.L., Sedgwick, R., Sleator, D.D., and Tarjan, R.E. The pairing heap: A new form of self-adjusting heap. *Algorithmica* 1 (Mar. 1986), 111–129.
6. Jones, D.W. An empirical comparison of priority queue and event set implementations. *Commun. ACM* 29, 4 (Apr. 1986), 300–311.
7. Knuth, D.E. *The Art of Computer Programming, Vol. 3: Sorting and Searching*. Addison-Wesley, Reading, Mass., 1973.
8. Tarjan, R.E. Amortized computational complexity. *SIAM J. Algorithm Disc. Meth.* 6 (Apr. 1985), 306–318.
9. Vuillemin, J. A data structure for manipulating priority queues. *Commun. ACM* 21, 4 (Apr. 1978), 309–314.
10. Williams, J.W.J. Algorithm 232: Heapsort. *Commun. ACM* 7, 6 (June 1964), 347–348.

**CR Categories and Subject Descriptors:** D.1 [Programming Techniques]; E.1 [Data Structures]—trees; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—sorting and searching; G.2.2 [Discrete Mathematics]: Graph Theory—trees

**General Terms:** Algorithms, Design, Experimentation, Performance, Theory

**Additional Key Words and Phrases:** Amortized analysis, heap, pairing heap, priority queue


Received 4/86; revised 11/86; accepted 11/86

Author's Present Address: John T. Stasko and Jeffrey Scott Vitter, Department of Computer Science, Box 1910, Brown University, Providence, R.I. 02912.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

# Q:

What's Fast, Convenient,  
and Free?



# A:

ACM's "Order Express"  
Service for ACM Publications.

# 1-800-342-6626

Your credit card and our toll free number provide quick fulfillment of your orders.

- Journals
- Conference Proceedings
- SIG Newsletters
- SIGGRAPH VIDEO REVIEW
- "Computers in your Life" (An Introductory Film from ACM)

For Inquiries and other Customer Service  
call: (301) 528-4261

**acm** ASSOCIATION FOR  
COMPUTING MACHINERY