

**stichting
mathematisch
centrum**



AFDELING ZUIVERE WISKUNDE

ZW 60/75

DECEMBER

P. VAN EMDE BOAS, R. KAAS & E. ZIJLSTRA

DESIGN AND IMPLEMENTATION OF AN EFFICIENT
PRIORITY QUEUE

Prepublication

2e boerhaavestraat 49 amsterdam

BIBLIOTHEEK MATHEMATISCH CENTRUM
—AMSTERDAM—

56.000

Printed at the Mathematical Centre, 49, 2e Boerhaavestraat, Amsterdam.

The Mathematical Centre, founded the 11-th of February 1946, is a non-profit institution aiming at the promotion of pure mathematics and its applications. It is sponsored by the Netherlands Government through the Netherlands Organization for the Advancement of Pure Research (Z.W.O), by the Municipality of Amsterdam, by the University of Amsterdam, by the Free University at Amsterdam, and by industries.

AMS(MOS) subject classification scheme (1970): 68A20, 90BXX, 05C05, 90-04,
05-04.

Design and implementation of an efficient Priority Queue ^{*)}

by

P. van Emde Boas, R. Kaas & E. Zijlstra

ABSTRACT

We present a data structure, based upon a hierarchically decomposed tree, which enables us to manipulate on-line a priority queue whose priorities are selected from the interval $1 \dots n$ with an average and worst case processing time of $O(\log \log n)$ per instruction. The structure can be used to obtain a mergeable heap whose time requirements are about as good. Full details are explained based upon an implementation of the structure in a PASCAL program contained in the paper.

KEY WORDS & PHRASES: *priority queue, binary tree, analysis of algorithms, mergeable heap, PASCAL, disjoint set union, insert-extract min problem.*

^{*)} This paper is not for review; it is meant for publication elsewhere.

1. INTRODUCTION

The main problems in the design of efficient algorithms for set-manipulation result from the incompatible requests posed by the distinct operations one likes to execute simultaneously. Instructions for inserting or deleting or for testing membership of elements in sets require a data structure supporting random access. On the other hand instructions for computing the value of the smallest or largest element, or the successor or predecessor of a given element, require an ordered representation. Finally instructions which unite two sets, so far, have only been implemented efficiently using a tree structure.

An example of an efficient algorithm which resolves one of these conflicts is the well-known union-find algorithm; its worst case average processing time per instruction has been shown to be of the order $A(n)$ in case of $O(n)$ instructions on an n -elements universe, where A is the functional inverse of a function with Ackerman-like order of growth (cf. [1], [9]).

The algorithms published until now to resolve the conflicting demands of order and random access all show a worst case processing time of $O(\log n)$ per instruction for a program of $O(n)$ instructions on an n -elements universe which has to be executed on-line. Clearly we should remember that each instruction repertoire which enables us to sort n reals by issuing $O(n)$ instructions needs an $O(\log n)$ processing time for the average instruction in doing so. However if the universe is assumed to consist of the integers $1 \dots n$ only, the information-theoretical lowerbound on the complexity of sorting does not apply; moreover it is known that n integers in the range $1 \dots n$ can be sorted in linear time.

Data structures which have been used to solve the conflict between

order and random access are (among others) the binary heap, AVL trees and 2-3 trees. In AHO, HOPCROFT & ULLMAN [1] 2-3 trees are used to support the instruction repertoire INSERT, DELETE, UNION and MIN with a worst case processing time of order $O(\log n)$ per instruction. The authors introduce the name *mergeable heap* (resp. *priority queue*) for a structure supporting the above operations (excluding UNION).

In the present paper we describe a data structure which represents a priority queue with a worst case processing time of $O(\log \log n)$ per instruction, on a Random Access Machine. The storage requirement is of the order $O(n \log \log n)$. The structure can be used in combination with the tree-structure from the efficient union-find algorithm to produce a mergeable heap with a worst-case processing time of $O((\log \log n) \cdot A(n))$ and a space-requirement of order $O(n^2)$. The possible improvements of the space requirements form a subject of continued research.

The mentioned $O(\log n)$ processing time for manipulating priority queues and mergeable heaps sometimes form the bottleneck in algorithms; in most cases however the priority values involved are assumed to be reals, and consequently our structure can not be used. In other examples like TARJAN's recent algorithm to compute dominators in directed graphs [10], the algorithm based on a mergeable heap has in the mean time been replaced [8] by one which does not use such a structure and which runs in inverse Ackermann time $A(n)$. It seems however reasonable to assume that applications will be described soon; such applications were suggested by HUNT & SZYMANSKI [7] and EVEN & KARIV [4].

Disadvantages of the structure are, beside the impossibility of manipulating real priority values, the large storage requirements of $O(n)$ for

each additional copy of a priority queue, and the theoretical objection that the time is measured using the uniform measure on a RAM; if charged according to the logarithmic measure the improvements of efficiency are lost because of the size of the involved arguments and values.

1.1. Structure of the paper

Section 2 contains some notations and background information. In section 3 we present a "silly" implementation of a priority queue with an $O(\log n)$ processing time per instruction. Reconsidering this implementation we indicate two possible ways to improve its efficiency to $O(\log \log n)$.

In section 4 we describe our stratified trees and their decomposition into canonical subtrees. Next we show how these trees can be used to describe a priority queue with an $O(\log \log n)$ worst and average case processing time per instruction. The algorithms for performing the elementary manipulations on stratified trees are explained in section 5. These algorithms are derived from the PASCAL implementation of our priority queue in section 8. It is explained how the complete stratified tree is initialized using time $O(n \log \log n)$. Section 6 discusses how the structure can be used if more than one priority queue has to be dealt with; the latter situation arises if we use our structure for implementing an efficient mergeable heap. Finally, in section 7, we indicate a few relations with other set-manipulation problems.

Throughout sections, 4, 5 and 6 identifiers typed in this different type font denote the values and meanings of the same identifiers in the PASCAL implementation.

2. GENERAL BACKGROUNDS

2.1. Instructions

Let n be a fixed positive integer. Our universe will consist of the subsets of the set $\{1, \dots, n\}$. For a set S in our universe we consider the following instructions to be executed on S :

MIN	: Compute the least element of S
MAX	: Compute the largest element of S
INSERT(j)	: $S := S \cup \{j\}$
DELETE (j)	: $S := S \setminus \{j\}$
MEMBER (j)	: Compute whether $j \in S$
EXTRACT MIN	: Delete the least element from S
EXTRACT MAX	: Delete the largest element from S
PRECEDECESSOR (j)	: Compute the largest element in $S < j$
SUCCESSOR (j)	: Compute the least element in $S > j$
NEIGHBOUR (j)	: Compute the neighbour of j in S (see definition in section 3)
ALL MIN (j)	: remove from S all elements $\leq j$.
ALL MAX (j)	: remove from S all elements $\geq j$.

If an instruction cannot be executed properly, e.g. MIN if $S = \emptyset$, an appropriate action is taken.

2.2. Priority queues

A *priority queue* is a data structure representing a single set $S \subset \{1, \dots, n\}$ on which the instructions INSERT, DELETE, and MIN can be executed *on-line* (i.e. some arbitrary order and such that each instruction should be

executed before reading the next one). Although the priority queue is our main target, we mention at this point that actually the complete instruction repertoire given above is supported on our data structure with a worst and average case processing time of $O(\log \log n)$ per instruction (except for the last two instructions where the processing time is $O(\log \log n)$ for each element removed).

The complete list of instructions above will be called the *extended repertoire* hereafter.

2.3. Union-find problem

For arbitrary partitions $\Pi = \{A, B, \dots\}$ of $\{1, \dots, n\}$ we consider the following instructions:

FIND (i) : compute the set currently containing i
 UNION (A,B,C) : Form the union of the sets A and B and give the name C to this union.

There is no specific name for a data structure supporting these two instructions; the problem of manipulating such a structure is known as the *union-find problem*.

For the famous union-find algorithm based on Tittertrees, which is described in more details in section 6, TARJAN [9] has proved an upper and lower bound of order $A(n)$.

2.4. Mergeable heaps

A *mergeable heap* is a data structure which supports the instructions INSERT, DELETE, MEMBER and MIN on sets which themselves can be united and searched, i.e. UNION and FIND are also supported. A mergeable heap may be

obtained from the union-find structure by replacing the internal nodes of the Tritter trees by complete priority queues. The resulting structure has an $O(\log \log(n)A(n))$ average processing time per instruction; its space requirements are $O(n^2)$. Details are given in section 6.

3. A "SILLY" PRIORITY QUEUE WITH $O(\log n)$ PROCESSING TIME

3.1. The structure

The scheme described in this section is designed primarily in order to explain the ideas behind the operations to be executed on the much more complicated structure in the next section.

We assume in this section that $n = 2^k$. We consider a fixed binary tree of height k . The $2^k = n$ leaves of this tree will represent the numbers $1 \dots n$ in their natural order from left to right. The leaves thus represent the potential members of the set S . If we had counted from 0 to $n - 1$ this order is nothing but the interpretation of the binary representation of a number as an encoding of the path from the root to the leaf; the binary digits are read from left to right where 0 denotes "go left" and 1 means "go right".

To each node in the tree we associate three pointers, linking the node to its father and its left- and righthand son. Moreover each node has a one-bit mark field.

A subset $S \subseteq [1 \dots n]$ is represented by marking all the leaves corresponding to members of S , together with all nodes on the paths from these leaves to the root of the tree; see diagram 1.

We present the following sketches of algorithms:

INSERT (i) : mark leaf i and all nodes on the path from leaf i to the root, until you encounter a node which was already marked.

DELETE (i) : unmark leaf i and all nodes on the path from leaf i to the root upto but not including the lowest node on this path having two marked sons.

MEMBER (i) : test whether leaf i is marked.

MIN (MAX) : proceed from the root to the leaves selecting always the leftmost (rightmost) present son.

EXTRACT MIN (EXTRACT MAX): min (max) followed by delete

ALLMIN (j) : while MIN \leq j do EXTRACTMIN od

ALLMAX is defined analogously.

PREDECESSOR (j) : proceed from leaf j to the root until a node is encountered having j as a righthand side descendent where the lefthand son is marked. Proceed from this lefthand son to the leaves always taking the rightmost present son.

SUCCESSOR (j) is defined analogously.

Note that all instructions except PREDECESSOR and SUCCESSOR use the lowest marked node on the path from an unmarked leaf to the root or the lowest *branchpoint* (i.e. a node having both sons marked) on the path from a marked leaf to the root. An analogous instruction which does not climb beyond the lowest "interesting" node is the instruction.

NEIGHBOUR (j) : proceed from leaf j to the lowest node such that the "other" son of this node is marked. If this other son is

a lefthand son then proceed from this node to the leaves always selecting the rightmost marked leaf; otherwise select always the leftmost marked leaf.

If we represent the set furthermore using a doubly linked list (for which the two downward pointers at the leaves may be used), a call of NEIGHBOUR followed by one or two steps in the list will be adequate to execute a call of PRECEDECESSOR or SUCCESSOR.

From the description of the instruction NEIGHBOUR it is clear in which sense the neighbour of a number j is the "nearest" present element to j ; the neighbour has the largest possible number of common ancestors and in case this condition does not define the neighbour unambiguously the neighbour is the other descendent of the lowest common ancestor which has developed as near to j as possible.

3.2. Improvements of the time efficiency

It is clear from the above description that our "silly" structure supports the extended repertoire with an $O(\log n)$ processing time per instruction. Using the doubly linked list as an "extra" representation INSERT and DELETE and NEIGHBOUR take time proportional to the distance in the tree traversed upon a call whereas MIN and MAX take constant time.

The remaining instructions are "composite". This observation opens a way to improve the efficiency. The time saved by not climbing high upwards in the tree can be used to perform more work at a single node. For example, if we decide to use at each node a linear list of present sons instead of a fixed number, we can easily accomodate for a tree with branching orders increasing from the root to the leaves without disturbing the $O(k)$ processing

time. Using a tree with branching orders $2, 3, 4, \dots, k$ which contains $k! = n$ leaves, we can maintain a priority queue of size $O(k!)$ in time $O(k)$; the above set up yields therefore an $O(\log n / \log \log n)$ priority queue which is already better than we had before.

There is, however, much more room for improvement. The operations which we like to execute at a single node are themselves priority queue operations. Consequently using a binary heap we can accomodate for the branching orders $2, 4, 8, \dots, 2^k$, which yields a priority queue of size $O(2^{k^2/2})$, and the processing time is reduced to $O(\sqrt{\log n})$.

Note that in both modifications the space requirements remain of $O(n)$, which is not true for the structure described in §4.

According to the "divide and conquer" strategy, we should however use at each node the same efficient structure which we are describing. This suggests the following approach. The universe $[1 \dots n]$ is divided into \sqrt{n} blocks of size \sqrt{n} . Each block is made a priority queue of size \sqrt{n} , whereas the blocks themselves form another priority queue of this size. To execute an INSERT we first test whether the block containing the element to be inserted contains already a present element. If so, the new element is inserted in the block; otherwise the element is inserted as first element in its block and the complete block is inserted in the "hyper-queue". A DELETE instruction can be executed analogously.

Assuming that we can implement the above idea in such a way that inserting a first and deleting the last element in a block takes constant time *independent of the size of the block*, the above description yields for the run-time a recurrence equation of the type $T(n) \leq T(\sqrt{n}) + 1$ which has a solution $T(n) \leq O(\log \log n)$.

Another way to improve the "silly" representation which leads again to the same efficiency is conceived as follows. As indicated the "hard" instructions proceed by traversing the tree upwards upto the lowest "interesting" node (e.g. a branchpoint), and proceeding downwards along a path of present node.

If these traversals could be executed by means of a "binary search on the levels" strategy, the processing time is reduced from $O(k)$ to $O(\log k) = O(\log \log n)$. A similar idea is involved in the efficient solution of a special case of the lowest common ancestor problem given by AHO, HOPCROFT & ULLMAN [2].

The reader should keep both approaches in mind while reading the sequel of this paper.

4. A STRATIFIED-TREE STRUCTURE

4.1. Canonical subtrees and static information

The hierarchical decomposition of a tree in top and bottom trees, as suggested by the divide and conquer strategy, can be performed perfectly, only if the height of the tree k is a perfect power of two, i.e. $k = 2^h$ for some natural number h . Since in approximating the given size of the universe $\{1 \dots n\}$ by numbers of the form 2^{2^h} the size of n might have to be almost squared (with the resulting prohibitive time- and space requirements) we settle for a less perfect decomposition where $n = 2^k$ and where the characteristic processing time is of order $h = \lceil \log \log n \rceil$. Replacing n by the next largest power of two, the overhead for the initializing time and storage requirements is at most a constant factor 2.

The hierarchical decomposition of our tree is based on a precomputed *rank function*. If $k = 2^n$ a function RANK is defined as follows:

For $1 \leq j$ we define $\text{RANK}(j)$ to be the largest number d such that $2^d | j$ and $2^{d+1} \nmid j$. For example $\text{RANK}(12) = 2$ and $\text{RANK}(7) = 0$. By convention we take $\text{RANK}(0) = h + 1$.

Note that for $j > 0$, $\text{RANK}(j) = d$ and $j - 2^d > 0$ we have $\text{RANK}(j) < \text{RANK}(j+2^d)$ and $\text{RANK}(j) < \text{RANK}(j-2^d)$; moreover $\text{RANK}(j+2^d) \neq \text{RANK}(j-2^d)$. Inspired by these properties we define for less perfect k the concept of a rank function as follows:

A function $f : \{i \dots j\} \rightarrow \mathbb{N}$, with $i < j$ is called a *rank function* provided:

- (i) $f(i) \neq f(j)$
- (ii) $f(k) < \min(f(i), f(j))$ for $i < k < j$
- (iii) either $j = i + 1$ or otherwise there exists a unique element k such that $i < k < j$, $f(k) = \min(f(i), f(j)) - 1$ and the restrictions $f|_{\{i \dots k\}}$ and $f|_{\{k \dots j\}}$ are rank functions.

For example if $\{i \dots j\} = \{0 \dots 8\}$ then the function described by the sequence (4 0 1 0 2 0 1 0 3) is a legal rank function and so is the function described by (8 0 1 2 3 4 5 6 7); clearly the first rank function reflects perfect "divide and conquer" whereas the second function is "as bad as possible". If $f: \{i \dots j\} \rightarrow \{0 \dots h+1\}$ is a rank function then clearly $j-i+1 \leq 2^h$.

An "optimal" rank function $\{0 \dots k\} \rightarrow \{0 \dots h+1\}$ may be defined inductively as follows: $\text{RANK}(0) = h+1$; $\text{RANK}(n) = h$; next as long as there exist intervals $\{i, j\}$ such that $\text{RANK}(i)$ and $\text{RANK}(j)$ are already defined and $\text{RANK}(k)$ is not yet defined for $i < k < j$ one defines $\text{RANK}(\lfloor (i+j)/2 \rfloor) = \min(\text{RANK}(i), \text{RANK}(j)) - 1$. For an explicit algorithm the reader is referred to the procedures `fillranks` and `start` in the program in section 8, where it is shown that this function can be precomputed on a RAM in time $O(k)$ (without use of division by two).

In the sequel we consider a fixed binary tree T with n leaves and k levels; $RANK$ denotes a fixed optimal rank function mapping $\{0 \dots k\}$ onto $\{0 \dots h+1\}$.

The *level* of a node v in T is the length of the path from the leaves of T to v ; the *rank* of v is the rank of the level of v . Note that the rank of the leaves equals $h + 1$, and the rank of the top equals h ; all other nodes have lower ranks. The *position* of a leaf is the number in the set $\{1, \dots, n\}$ represented by this leaf. The *position* of an internal node v equals the position of the rightmost descendent leaf of its lefthand son; this number indicates where the borderline lies from the two parts resulting from splitting the tree along the path from v to the root.

A *canonical subtree* (CS) of rank d of T is a binary subtree having as root a node of rank $\geq d$, as leaves all descendents of the root at the nearest level of rank $\geq d$ (and consequently internal nodes of rank $< d$). From the definition of a rank function it follows that for non-trivial CS either the rank of the top or the ranks of the leaves equals d . In the first case the tree is called a *bottom tree* whereas it is called a *toptree* otherwise. A CS of rank d either is a trivial three-element tree, or it is composed in a top and bottom CS's of rank $d-1$. Canonical subtrees of rank 0 are always trivial. The subtree of a CS consisting of its root with all its left (right) hand side descendants is called a *left (right) canonical subtree*.

To any node v of T we associate the following subtrees which are called the canonical subtrees of v . Let $d = RANK(v)$.

$UC(v)$: the unique canonical subtree of rank d having v as a leaf.

$LC(v)$: the unique canonical subtree of rank d having v as a root.

Note that $UC(v)$ is not defined if v is the root whereas $LC(v)$ is not defined if v is a leaf of T . When $d = 0$, $UC(v)$ and $LC(v)$ consist of three nodes. Note moreover that the rank of the root of $UC(v)$ and the rank of the leaves of $LC(v)$ is higher than d .

The left(right) canonical subtree of $LC(v)$ is denoted $LLC(v)$ ($RLC(v)$). $LC(v)$ and the half of $UC(v)$ containing v together form the *reach* of v , denoted $R(v)$. The dynamical information stored at v depends only on what happens within its reach. The reach of the top is the complete tree, whereas the reach of a leaf is the set of leaves. See diagram 2 for an illustration.

Clearly the reach of an internal node v of rank d is a subset of some canonical subtree of rank $d + 1$, denoted $C(v)$. We say that v lies at the *center-level* of $C(v)$; moreover, v is called the *center* of its reach $R(v)$.

For each node v and each $j \leq h$ we denote by $FATHER(v, j)$ the lowest proper ancestor of v having rank $\geq j$. Clearly $FATHER(v, h)$ equals the root t of T , whereas $FATHER(v, 0)$ is the "real" father of v in T (provided $v \neq t$). At each node we have an array of h pointers $father[0 : h-1]$ such that $father[i]$ yields the rank- i father of v . Since $FATHER(v, h)$ always yields the root of the tree this element doesn't need to be included. These pointers enable us to climb along a path in the tree to a predetermined level in $O(h)$ steps. Moreover, given the root of a cs U and one of its leaves, we can proceed in a single step to the center of the smallest reach containing the two which is entirely contained within U .

Actually it will be seen that within our recursive algorithms in section 8 for a node of rank d a father pointer to a father of rank $\geq d$ is never asked for ! Consequently these pointers might be omitted from the structure; the resulting structure however should consist of rank-dependent nodes, whereas the gain in storage efficiency is at most a

constant factor; therefore this improvement is not implemented.

The static information at a node contains moreover its position and if it is an internal node its rank. The static information can be allocated and initialized in time $O(n \log \log n)$; details will be given in the next section.

4.2. Dynamical information

The dynamical information at internal nodes is stored using four pointers $lmin$, $lmax$, $rmin$ and $rmax$ and an indicator field ub , which can assume the values plus, minus and undefined. At leaves the dynamical information consists of two pointers $successor$ and $predecessor$, and a boolean $present$.

Let $S \subset \{1, \dots, n\}$ be a set which has to be represented in our stratified tree. We say that the leaves corresponding to members of S and all their ancestors in the tree are *present*; the present nodes are exactly the nodes which were marked in our silly structure. A present node can become *active* and in this case its information fields contain meaningful information. The values of these fields of a non-active internal node are:

$lmin = nil$, $lmax = nil$, $rmin = nil$, $rmax = nil$ and $ub = undefined$. For a non-active leaf these values are $predecessor = nil$, $successor = nil$,

$present = false$. For an active leaf v the meaning of these fields should be:

$predecessor$: points to the leaf corresponding to the predecessor in S of the number corresponding to v if existent; otherwise
 $predecessor = nil$.

$successor$: analogous for the successor

$present = true$

Remember that a branchpoint is an internal node having two present sons.

Let v be an internal node, and denote the top of $C(v)$ by t . If v is active its dynamical information fields have the following meaning:

$lmin$: points to the leftmost present leaf of $LLC(v)$ if such node exists;
 otherwise $lmin = nil$.
 $lmax$: idem for the rightmost present leaf of $LLC(v)$.
 $rmin$: idem for the leftmost present leaf of $RLC(v)$
 $rmax$: idem for the rightmost present leaf of $RLC(v)$
 ub = plus if there occurs a branchpoint in between v and t , and minus otherwise.

If v is an active internal node it is present and consequently $LC(v)$ contains at least one present leaf; this shows that it is impossible to have an active internal node with four pointers equal to nil .

As suggested in the preceding section the time needed to insert a first or to delete a last element should be independent of the size of the tree. This is realized by preventing present nodes from becoming active unless their activity is needed. This is expressed by the following:

Properness condition: Let v be a present internal node. Then v is active if and only if there exists a branchpoint in the interior of the reach of v (i.e. there exists a branchpoint $u \in R(v)$ which is neither the top nor a leaf of $C(v)$).

A leaf is active if and only if it is present; the root is active iff the set is non-empty.

If the internal node v is non-active but present then there is a unique path of present nodes going from the top t of $R(v)$ to a unique present leaf w of $C(v)$ contained in $R(v)$. In our approach we can proceed from t to w and backwards without ever having to visit v , making it meaningless to store information at v .

If some canonical half-tree has two present leaves then all its present nodes at its center level are active. Also if a node v of rank d is active then $FATHER(v, d)$ is active as well. We leave the verifications of these assertions as an exercise to the reader.

The set $S \subset \{1, \dots, n\}$ is represented as follows. First the leaves corresponding to the element of S and all their ancestors are declared to be present. Next we compute using the properness condition which present nodes become active. Finally the dynamical fields of all active and non-active nodes are given their proper values. The resulting information content is called the representation the set S . We leave it to the reader to convince himself that this representation is unique.

(In our actual program the structure is initialized at $S = \emptyset$, representations of all other sets being the result of execution of a sequence of instructions from the extended repertoire.)

An example of a proper information content is given in diagram 3 (omitting the evident doubly linked list data). The symbol \sim denotes nil resp. undefined.

5. OPERATIONS ON THE STRATIFIED TREE STRUCTURE

Once having described the representation of a set S by assigning values to particular fields in the stratified tree, the next step is to indicate

how the setmanipulation operations mentioned in section 1 can be executed such that

- (i) a processing time of $O(h) = O(\log \log n)$ is realized;
- (ii) the structure of the representation is preserved, i.e. the properness condition should remain valid.

Moreover we must indicate how the static information together with a legitimate initial state for the dynamic information can be created in the proper time and space (i.e. both of order $n \log \log n$).

In this section we pay no attention to the self-evident operations needed to manipulate the doubly linked list structure formed by the leaves of our tree; they are described explicitly in the program in section 8.

5.1. Initialization

Initialization takes place during a single tree-transversal in mixed pre- and in-order. When a node is processed its father-pointers and its position and in case of internal nodes its rank are stored in the appropriate fields. The needed computations are based on the following relations:

- (i) the fathers of the top are nil; the fathers of a direct son v of a node w where $RANK(w) = d$ satisfy

$$FATHER(v, j) = w \quad \text{for } j \leq d$$

$$FATHER(v, j) = FATHER(w, j) \quad \text{for } j > d.$$

The node w is accessible during the processing of v by use of a parameter $fath$ in the recursive procedure which executes the tree-transversal

- (ii) the level of a node is one less than the level of its father
- (iii) the position of an internal node equals the number of leaves encountered after its lefthand subtree has been processed. This number

is accessible by means of a global variable count. (It is this evaluation of the position where in-order traversal is used).

- (iv) The rank of a node depends only on the level and can be stored using a table ranks of size $k+1$.

Once having precomputed the ranks using the local procedure fillranks, the above relations show now the static structure is initialized without having "illegitimate" instructions like multiplications and bit-manipulations, in time $O(\log \log n)$ per node processed. Since there exist $2n - 1$ nodes this shows that the initialization takes time $O(n \log \log n)$. The space requirements $O(n \log \log n)$ follow since the space needed for each node is $O(\log \log n)$.

Since the level is not part of the static information it may be represented as a local variable of the recursive initialization routine.

5.2. Operations

The extended instruction repertoire can be expressed (disregarding the doubly linked list operations) in terms of three primitive operations ins, del and neighb. Each of these operations is described by a linearly recursive procedure. The procedures are called upon the complete tree of rank h . If called upon a canonical subtree the procedures either terminate within constant time independent of the rank, or the procedure executes a single call on a top or bottom canonical subtree of rank one less, preceded and followed by a sequence of instructions taking constant time independent of the rank. A call upon a subtree of rank 0, or any other trivial CS terminates without further recursive calls of the procedure. From the above assertions which can be verified by inspection of the procedure bodies, it follows directly that the run-time of each procedure is of order $h = \lceil \log \log n \rceil$. The procedures ins, del, and neighb all have the property

that their innermost call is a bottom call, where we consider the complete tree to be a bottom tree as well.

In the execution of an algorithm we have frequently the situation that we have a CS with root t and leaf v and that we want to inspect or modify the fields at t in the direction of v , i.e. the left-hand fields at t if v is a left-hand descendant of t etc. To decide whether a certain descendant of t lies in the left- or right-hand subtree it is sufficient to compare the positions of the two nodes. We have in general:

The descendant v of t is a left-hand descendant if the position of v is not greater than the position of t .

Actually the position of a node was introduced to facilitate this easy test on the handiness of a descendant.

The procedures `ins`, `del` and `neighb` use the following primitive operations.

<code>myfields (v,t)</code>	yields	a pointer to the fields at t in the direction of v . This pointer is of the type <code>fieldptr</code> .
<code>mymin (v,t)</code>	yields	the value of the min-field at t in the direction of v (which happens to be a pointer).
<code>mymax (v,t)</code>		analogous for the maxfield.
<code>yourfields (v,t),</code> <code>yourmin (v,t) and</code> <code>yourmax (v,t)</code>		yield the analogous values of the field at t in the other direction.

minof (t) yields the leftmost value of the four pointer fields at t if t is active, and nil otherwise.

maxof (t) yields the rightmost value analogously.

Finally the procedure clear gives the dynamic fields at its argument the values corresponding to the non-active state. The identifiers mentioned in the procedures mostly are of the type "pointer to node" (ptr) where "node" is a record-type containing the fields mentioned in the preceding sections.

5.2.1. The procedure ins

ins is a function procedure yielding as result the value of a pointer to the neighbour of the node being inserted. This neighbour is subsequently used for inserting the node into the doubly linked list. (It should be mentioned that we tacitly have generalized the meaning of neighbour to the case of a CS which is not the complete tree.)

ins has five parameters called by value; its procedure heading reads:

```
function ins (leaf, top, pres: ptr; nobranchpoint: boolean; order: integer):
    ptr;
```

The meaning of the parameters is as follows:

order: the rank of the CS on which the procedure is called

leaf: the node to be inserted

top: the root of the CS on which the procedure is called

pres: a present leaf (actually the minimal one) of the CS on which
 the procedure is called of the same handiness as leaf at top
 nobranchpoint: true iff leaf's side of the CS on which the procedure is
 called contains no branchpoint.

At first glance the parameter pres seems to be unnecessary since its value can be derived from the values of myfields(leaf, top). However in the case where the CS under consideration is a top-CS of a CS of next higher rank the fields at top refer to nodes at a level far below the level of leaf and consequently their values may be misunderstood.

A call of ins terminates without further recursive calls if leaf's side of the CS under consideration does not contain a present leaf (pres = nil). Otherwise the nodes hl = FATHER(leaf, order-1) and hp = FATHER(pres, order-1) are computed. Now if nobranchpoint is true then hp is present without being active and special actions should be undertaken in this case. In this case hl is present iff hl = hp and depending on this equality either the bottom-call

```
ins (leaf, hl, mymin (leaf, hl), true, order-1)
```

or the top-call

```
ins (hl, top, hp, true, order-1)
```

is executed after having "activated" the right fields at hp and hl.

In this situation the procedure delivers pres as its value.

If nobranchpoint is false then hl is present iff it is active which is tested by inspecting its ub-field. If hl is active the bottom-call

```
ins (leaf, hl, mymin (leaf, hl), mymin (leaf, hl) = mymax (leaf, hl),  

    order-1)
```

is executed and its value is yielded as the result of ins. Otherwise, the top-call

ins (hl, top, hp, nobranchpoint, order-1)

is executed after having set nobranchpoint := (hpt.ub = minus) and having activated the fields at hl and hp. This call yields as a result the neighbour of hl in the top-tree in nb, and depending the outcome of a comparison between the positions of hl and nb the value of insert equals minof (nb) or maxof (nb).

After these activities the fields at the top may have to be adjusted if the current call is a call on a bottom-CS, which is the case iff order equals the rank of top. From this point of view the complete tree has to be considered a bottom-CS, which explains why the levels are numbered from the leaves to the top.

The initial call of ins reads:

ins (ej, root, mymin (ej, root), mymin (ej, root) = mymax (ej, root), h)

where it is assumed that ej is not a present leaf. This condition is enforced by the calling routine insert.

For the complete PASCAL text of ins we refer to section 8. The four possible situations are illustrated in diagram 4.

5.2.2. The procedure del

The procedure del yields no value. It has six parameters, the first three of which are called by value, the others being called by reference. The procedure heading reads:

```
procedure del (leaf, top : ptr; inter; var pres1, pres2: ptr; var
    nobranchpoint : boolean);
```

The meaning of the value-parameters is as follows:

leaf : the leaf to be deleted
 top: the root of the CS considered
 order: the rank of the CS considered

The remaining parameters have after a call of del the following meaning:

pres1, pres2: present leaves in the CS considered, one of them being the neighbour of leaf (see explanation below)
 nobranchpoint: true iff there occurs no branchpoint on the path from top to pres 1.

A call of del should make non-present leaf and its ancestors up to the lowest branchpoint but in doing so other nodes on different paths which were active may have to become inactive. As long as this possibility holds nobranchpoint remains true.

Proceeding downwards from the other son of the lowest branchpoint as near as possible we arrive at the neighbour; if we however select always the remotest present node, we arrive at a node which might be called the *extreme* of leaf in the tree. The extreme, as a "binary approximation" of leaf is as good as the neighbour, but in the usual sense it is as far away as possible.

After a call of del, pres1 and pres2 are the neighbour and the extreme of leaf ordered according to their positions (i.e. $\text{pres1} \uparrow \text{pos} \leq \text{pres2} \uparrow \text{pos}$).

del terminates without inner call if the lowest branchpoint equals top; at this time pres1 and pres2 are initialized with the values

yourmin(leaf,top) and yourmax(leaf, top) and nobranchpoint is made true if these two values are equal.

Otherwise h1 becomes the node halfway inbetween leaf and root and depending on whether leaf is the unique present son in its lower canonical half tree a recursive call of del is initialized on UC(h1) or LC(h1). After completion of this inner call the current call of del is completed as follows:

The values pres1 and pres2 are updated depending on whether the inner call just terminated was a top or a bottom call. If the last call was a top call then $\text{pres1} := \text{minof}(\text{pres1})$; $\text{pres2} := \text{maxof}(\text{pres2})$ and their equality is tested again to decide whether nobranchpoint should remain true; if so the node formerly pointed at by pres1 is inactivated.

If the last call was a bottom call the ub field at the former top and the pointers away from pres1 at this node are inspected to decide whether there occurs a branchpoint at or above this node; if not the former top is inactivated.

The fields at the current top are adjusted only when the current call is a bottom call.

The initial call to del reads:

```
del (ej, root, h, p1, p2, nobrpt);
```

For the complete text of del see section 8. The calling routine delete makes sure that pt is a present leaf. For a situation leading to the inactivation of a present node see diagram 5.

5.2.3. The procedure neighb

The function neighb has five parameters which are called by value. Their meaning is about equal to the meaning of the parameters in insert,

however `pres` is replaced by the pair `pmin` and `pmax`.

`neighb` may be called both for present and nonpresent leaves. This is justified by the fact that without expensive bit-manipulation on the positions it is impossible to decide whether the neighbour is the predecessor or the successor of the given argument.

`pmin` and `pmax` are the left and rightmost present leaf on `leaf`'s side of the CS under consideration.

`neighb` terminates without an inner call in the following cases:

- (i) `pmin = nil`; now the neighbour resides on the other side of the tree
- (ii) `pmin = pmax = leaf`; idem
- (iii) `leaf` lies outside the interval `pmin - pmax`; in this case `neighb` yields the nearest of the two in the usual sense without needing to investigate the inner structure of the tree. (See diagram 6).

The short-cut (iii) is unique to the procedure `neighb`. If none of these situations occurs a recursive call is performed. This inner call is a top call if either the node `hl` at the center level is between `leaf` and top is not present (which in these circumstances is equivalent to non-active) or if `leaf` is the unique present descendent of `hl`; otherwise a bottom call is executed.

The initial call of `neighb` reads:

```
neighb (pt, root, mymin(pt, root), mymax(pt, root), h)
```

If called upon an empty tree or on the unique present leaf `neighb` yields `nil` as its result.

6. A TIME EFFICIENT MERGEABLE HEAP

The well known efficient union-find algorithm uses a representation of

sets by means of the so-called Triter trees. Each node in a tree corresponds to a member of a set and contains a pointer which either points to the name of the set if the node happens to be the root of his tree, or to his father in the tree otherwise. A UNION instruction is executed by making the root of the smaller tree a direct son of the larger one (balancing). To execute a FIND instruction the node corresponding to the element asked for is accessed directly, and his pointers are followed until the root of his tree is found; in the mean time all nodes which are encountered during this process are made direct descendants of the root, thus reducing the processing time at subsequent searches (path compression).

It has been established only recently how efficient the above algorithm is. Whereas its average processing time has been estimated originally as $O(\log \log n)$ (FISHER [5]) and $O(\log^* n)$ (HOPCROFT & ULLMAN [6] and independently PATERSON (unpublished)), a final upper and lowerbound $O(A(n))$ has been proved by TARJAN [9]. Remember that $\log^* n$ is the functional inverse of the function $2^{2^{\dots^2}} n$ i.e. $\log^* n$ equals the number of applications of the base-two logarithm needed to reduce n to zero. The function $A(n)$ is a functional inverse of a function of Ackermann-type which is defined as follows:

Define a by: $a(0,x) = 2x$; $a(i,0) = 0$; $a(i,1) = 2$ for $i \geq 0$; and $a(i+1,x+1) = a(i,a(i+1,x))$. Then we let $A(n) = \min\{j \mid a(j,j) \geq n\}$.

(The above definitions differ only inessentially from the ones given by TARJAN).

To obtain a mergeable heap we replace the unordered collection of sons of a certain node by a priority queue where the "value" of a node equals the minimal element in the set formed by this node and its descendants; see diagram 7.

In such a representation the instructions are executed as follows:

UNION: The root-priority queue of the structure containing the least number of elements is inserted in the root-priority queue of the other structure at the place corresponding to its least element.

FIND: First one proceeds from the element itself upwards to find the root-priority queue of the structure to which it belongs. Next, going downwards from this root back to the element the priority queues along this path are disconnected by delete operations. The queues are then inserted in the root priority queue at the position of their (possibly modified) least element.

MIN: By executing a min-instruction at the root-priority queue of a structure its least element will become known; a FIND instruction on this element will yield access to the location where it is stored.

INSERT & DELETE: These operations are reduced to the priority queue insert and delete by first executing a FIND instruction. The same holds for MEMBER.

In doing so the average processing time for an instruction becomes $A(n)$ times the processing time for the priority queue instructions used. As long as the latter time is not reduced below $O(\log n)$ the proposed representation of a mergeable heap should be considered inefficient, since there are $O(\log n)$ structures known for mergeable heaps (2-3-trees with unordered leaves [1]). Using our new efficient priority queue the proposed scheme becomes (as far as time is concerned) more efficient than the traditional ones.

For the above application we have to deal with a large number ($O(n)$) of distinct priority queues simultaneously.

If one has to represent several priority queues it makes sense to separate the static and dynamical information in the nodes. The static information is about equal for each queue. More in particular, using an "address plus displacement" strategy, where the position of a node is used as its address, one has access to each node whose position is known. Since all nodes are accessed by father pointers from below, or by the downward pointers from the dynamical information, it is sufficient to have available a single pre-computed copy of the static information in a stratified tree. For each queue involved in the algorithm a $O(n)$ size block of memory, directly accessible by the position of a node, should be allocated for the dynamical information.

Using the above strategy we arrive at the $O(n \log \log n \cdot A(n))$ -time, $O(n^2)$ -space representation of a mergeable heap promised in the introduction. It is clear that the larger part of the space required is never used, and luckily there is a well-known trick which allows us to use this much space without initializing it [1]. Still it is a reasonable question whether some dynamical storage allocation mechanism can be designed which will cut down the storage requirement to a more reasonable level.

A direct approach should be to allocate storage for a node at the time this node is activated. This method, however, seems to be incorrect. One must be able to give the correct answer to questions of the following type: "Here I am considering a certain CS with root top and some leaves $pres$ and $leaf$, where $pres$ is present and $leaf$ is not. Let hl be the ancestor of $leaf$ at the center level. To decide whether hl is active, and, if so, where it is allocated." Inspection of the ancestor at center level of $pres$ will yield the correct answer only if $pres$ is actually the neighbour of $leaf$; this however is not guaranteed in our algorithm.

The same problem arises if one first tries to compute the neighbour of leaf. Consequently it seems necessary to reserve a predetermined location to store h_1 which can be accessed knowing the position of h_1 in the CS under consideration and having access at its root.

The following approach yields a representation of a mergeable heap in space $O(n\sqrt{n})$ without disturbing the $O(\log \log n)$ processing time. Consider a rank d tree. As long as its left or right-hand side subtree contains not more than one present leaf, all necessary information can be stored at the root of the tree. If at a certain stage a second leaf at the same side must be inserted, the complete storage for the top tree is allocated as a consecutive segment, and a pointer at the root is made to refer to its initial address. In particular the nodes at the center level now have been given fixed addresses which are accessible via the root. The center-level nodes themselves are considered to be the roots of bottom trees of rank $d - 1$ which are treated analogously. In this manner a call of `ins` will allocate not more than $O(\sqrt{n})$ memory cells, whereas `neighb` does not use extra memory and `del` may return the space for a top tree if both sides of its enveloping CS have been exhausted except of a single leaf.

The initial address of the current relevant storage segment is given as a new parameter to the procedures whose value is passed on to an inner call, unless all enveloping calls are bottom calls.

The $O(n\sqrt{n})$ bound on the used memory for the mergeable heap algorithm is obtained by noting that at each intermediate stage the information contents are equal to one obtained by executing not more than n `ins` instructions.

We complete this section by noting that the storage requirements may be further reduced by replacing the binary division of the levels by an

r -ary one for $r > 2$, which might result for each $\epsilon > 0$ in an $O(n \log \log n \cdot A(n))$ -time, $O(n^{1+\epsilon})$ -space representation of a mergeable heap.

7. REDUCIBILITIES AMONG SET-MANIPULATION PROBLEMS

The on-line manipulation of a priority queue, which is also known as the on-line insert-extract min problem, is one out of a multitude of set manipulation problems. Each of these problems has moreover a corresponding off-line variant. In the off-line variant the sequence of instructions is given in advance and the sequence of answers should be produced, the programmer being free to choose the order in which the answers are given.

Clearly, each on-line algorithm can be used to solve the off-line variant, but the converse does not hold.

In [3] we have investigated the reducibilities among the one-line and off-line versions of the insert-extract-min-, union-find- and insert-allmin problems. Here we say that a problem A can be reduced to a problem B if an algorithm for B can be used to design an algorithm for A having the same order of complexity. If moreover A and B are both off-line problems it should be possible to translate an $O(n)$ -size A problem on a $O(n)$ -size structure into an $O(n)$ -size B problem on a $O(n)$ -size structure in time $O(n)$.

It has been shown by HOPCROFT, AHO & ULLMAN that the off-line insert-extract min problem is reducible to the on-line union-find problem [1]. Below we indicate how the off-line union-find problem and the off-line insert-allmin problem are reduced to each other. Together with the "natural" reduction of on-line insert-allmin to online insert-extractmin these reducibilities are represented in diagram 8 (the acronyms denoting the problems discussed).

7.1. The reduction OFIAM \leq OFUF

The reduction below is a modification of the reduction OFIEM \leq ONUF in [1].

Consider a sequence of insert and allmin instructions. For each instruction insert (j) we construct an object x_j . For each instruction allmin(i) we construct a barrier b_i . The value i is called the height of b_i .

Consider the sequence of objects and barriers in their original order. Clearly object x_j will be removed by the allmin instruction corresponding to the first following barrier with height $\geq j$.

We gather all objects x_j between two consecutive barriers b_j and b_i in a set $S(b_i)$ having the top barrier as name. Clearly the union instructions needed to build these sets can be generated off-line in linear time.

Next we generate for $j = 1, 2, \dots, n$ the instruction find(j), followed by instructions to remove all barriers having height j . A barrier b_j is removed by uniting the set having this barrier as a name with the set named by the nearest remaining barrier b_k above b_j . Note that it is illegal to include "computed arguments" in a sequence of instructions which are to be executed off-line. However by keeping the active barriers in a doubly linked list we can compute b_k during the process of translating the insert-allmin instructions into union-find instructions. Clearly the above translation can be executed in time $O(n)$.

It is left to the reader to verify that the result of executing a find(i) instruction in the resulting sequence yields the barrier corresponding to the allmin instruction which removes i .

7.2. The reduction OFUF \leq OFIAM

Consider a sequence of union and find instructions. We first build a binary tree by executing the subsequence of all union instructions as

follows: For each union instruction a new vertex is created having the two sets from which the union is formed as sons. At each vertex we store, moreover the name which is given to the union, and the time (number of the instruction) at which the union is formed. Without loss of generality we may assume that the final result is a single tree (otherwise add a number of extra union instructions).

Clearly construction of this tree takes time $O(n)$.

By running through the sequence of instructions another time we can organize for each element j a linear list of all the times at which instruction $\text{find}(j)$ must be executed.

Next we will traverse the tree in post order (i.e. first visit the subtrees and next visit the root). While traversing the tree we have available a bucket containing a set of find instructions.

Whenever we visit a leaf of the tree (i.e. a vertex which represents a singleton), we throw into the bucket all the find instructions corresponding to this leaf (i.e. all instructions $\text{find}(i)$ where $\{i\}$ is the set represented by the leaf).

If we process an arbitrary vertex v we remember that the set which is represented by v contains the elements represented by the leaves of the subtree with root v , during the period starting with the formation of the union (which time T is stored at vertex v) and the time T' at which this set is united into a larger set (which time is stored at the father of v). Hence we should throw out of the bucket all find instructions having times satisfying $T \leq t < T'$. Since the two sons of v already have been processed, find's with time $< T$ are no longer in the bucket and consequently we should remove from the bucket all find's with time $< T'$. This is an allmin instruc-

tion. It is clear that if $\text{find}(i)$ is removed at vertex v then $\text{find}(i)$ is answered by the name of the set stored at v .

Up to this point everything that we have described can be executed in time $O(n)$ for the complete tree, but this `allmin` instruction makes the algorithm nonlinear (assuming that no linear algorithm is known for the `insert-allmin` problem).

It is clear that, after having constructed the tree, the complete sequence of `insert` and `allmin` instructions which is going to be executed on the bucket can be computed by traversing the tree once in post order. This shows that the off-line union-find problem is reducible to the off-line `insert-allmin` problem.

8. THE PROGRAM

The program follows exactly the syntax of Standard Pascal [11], with the exception that `(*` and `*)` are used to begin and end a comment.

Together with a driver here omitted it has worked correctly to maintain a priority queue of size 2048 in (in octal) 150000 words of central memory.

```
PROGRAM PRIORITYQUEUE(INPUT, OUTPUT);
```

```
CONST N = 512; K = 9; H = 4; (* K H-1 H *)
(* N = 2 ; 2 < K <= 2 *)
```

```
TYPE PTR = ^NODE; FIELDPTR = ^SUBTREE;
SUBTREE = RECORD MIN, MAX: PTR END;
NODE = RECORD POSITION: 1..N; FATHERS: ARRAY[0..H] OF PTR;
CASE (* INTERNAL: *) BOOLEAN OF
TRUE: (UB: (MINUS, UNDEF, PLUS);
LEFT, RIGHT: FIELDPTR;
RANK: 0..H);
FALSE: (PRESENT: BOOLEAN; PRED, SUCC: PTR);
END;
```

```
VAR ROOT: PTR; CARDINALITY: 0..N; ELEMENTS: ARRAY[1..N] OF PTR;
```

```
FUNCTION MYMIN(LEAF, TOP: PTR): PTR;
```

```
BEGIN IF LEAF^.POSITION <= TOP^.POSITION THEN
MYMIN:= TOP^.LEFT^.MIN ELSE MYMIN:= TOP^.RIGHT^.MIN
END;
```

```
FUNCTION MYMAX(LEAF, TOP: PTR): PTR;
```

```
BEGIN IF LEAF^.POSITION <= TOP^.POSITION THEN
MYMAX:= TOP^.LEFT^.MAX ELSE MYMAX:= TOP^.RIGHT^.MAX
END;
```

```
FUNCTION MYFIELDS(LEAF, TOP: PTR): FIELDPTR;
```

```
BEGIN IF LEAF^.POSITION <= TOP^.POSITION THEN
MYFIELDS:= TOP^.LEFT ELSE MYFIELDS:= TOP^.RIGHT
END;
```

```
FUNCTION YOURMIN(LEAF, TOP: PTR): PTR;
```

```
BEGIN IF LEAF^.POSITION > TOP^.POSITION THEN
YOURMIN:= TOP^.LEFT^.MIN ELSE YOURMIN:= TOP^.RIGHT^.MIN
END;
```

```
FUNCTION YOURMAX(LEAF, TOP: PTR): PTR;
```

```
BEGIN IF LEAF^.POSITION > TOP^.POSITION THEN
YOURMAX:= TOP^.LEFT^.MAX ELSE YOURMAX:= TOP^.RIGHT^.MAX
END;
```

```
FUNCTION YOURFIELDS(LEAF, TOP: PTR): FIELDPTR;
```

```
BEGIN IF LEAF^.POSITION > TOP^.POSITION THEN
YOURFIELDS:= TOP^.LEFT ELSE YOURFIELDS:= TOP^.RIGHT
END;
```

```
FUNCTION MINOF(TOP: PTR): PTR;
```

```
  VAR TEMP: PTR;
```

```
  BEGIN TEMP:= TOP^.LEFT^.MIN; IF TEMP = NIL THEN
    MINOF:= TOP^.RIGHT^.MIN ELSE MINOF:= TEMP
  END;
```

```
FUNCTION MAXOF(TOP: PTR): PTR;
```

```
  VAR TEMP: PTR;
```

```
  BEGIN TEMP:= TOP^.RIGHT^.MAX; IF TEMP = NIL THEN
    MAXOF:= TOP^.LEFT^.MAX ELSE MAXOF:= TEMP
  END;
```

```
PROCEDURE INTOLIST(LOW, MED, UPP: PTR);
```

```
  BEGIN IF LOW <> NIL THEN LOW^.SUCC:= MED;
    IF UPP <> NIL THEN UPP^.PRED:= MED;
    MED^.PRED:= LOW; MED^.SUCC:= UPP;
    MED^.PRESENT:= TRUE; CARDINALITY:= SUCC(CARDINALITY)
  END;
```

```
PROCEDURE FROMLIST(LOW, MED, UPP: PTR);
```

```
  BEGIN IF LOW <> NIL THEN LOW^.SUCC:= UPP;
    IF UPP <> NIL THEN UPP^.PRED:= LOW;
    MED^.PRED:= NIL; MED^.SUCC:= NIL;
    MED^.PRESENT:= FALSE; CARDINALITY:= PRED(CARDINALITY)
  END;
```

```
FUNCTION MIN: INTEGER;
```

```
  VAR PT: PTR;
```

```
  BEGIN PT:= MINOF(ROOT); IF PT = NIL THEN MIN:= 0 ELSE
    MIN:= PT^.POSITION
  END;
```

```
FUNCTION MAX: INTEGER;
```

```
  VAR PT: PTR;
```

```
  BEGIN PT:= MAXOF(ROOT); IF PT = NIL THEN MAX:= 0 ELSE
    MAX:= PT^.POSITION
  END;
```

```
FUNCTION MEMBER(J: INTEGER): BOOLEAN;
```

```
  BEGIN IF (J > 0) AND (J <= N) THEN MEMBER:= ELEMENTS[J]^.PRESENT
    ELSE MEMBER:= FALSE
  END;
```

```

PROCEDURE CLEAR(V: PTR);

BEGIN WITH V^ DO
  BEGIN LEFT^.MIN:= NIL; LEFT^.MAX:= NIL;
    RIGHT^:= LEFT^; UB:= UNDEF
  END
END;

PROCEDURE START;

  VAR I, COUNT, LASTDIV2: INTEGER; EVEN: BOOLEAN;
    RANKS, DIV2: ARRAY[0..K] OF INTEGER;

  PROCEDURE FILLRANKS(L, U, R: INTEGER);

    VAR M: INTEGER;

  BEGIN M:= L + DIV2[U - L]; RANKS[M]:= R; IF M > L + 1 THEN
    BEGIN FILLRANKS(L, M, R - 1); FILLRANKS(M, U, R - 1) END
  END;

  PROCEDURE INITIALIZE(FATH: PTR; LEVEL: INTEGER);

    VAR I, RNK: INTEGER; V: PTR;

  BEGIN IF LEVEL > 0 THEN NEW(V, TRUE) ELSE NEW(V, FALSE);
    WITH V^ DO
      BEGIN RNK:= FATH^.RANK;
        FOR I:= 0 TO RNK DO FATHERS[I]:= FATH;
        FOR I:= RNK + 1 TO H DO FATHERS[I]:= FATH^.FATHERS[I];
        IF LEVEL = 0 THEN
          BEGIN PRESENT:= FALSE; SUCC:= NIL; PRED:= NIL;
            COUNT:= COUNT + 1; POSITION:= COUNT; ELEMENTS[POSITION]:= V
          END ELSE
          BEGIN NEW(LEFT); NEW(RIGHT); CLEAR(V); RANK:= RANKS[LEVEL];
            INITIALIZE(V, LEVEL - 1); POSITION:= COUNT;
            INITIALIZE(V, LEVEL - 1)
          END
        END
      END
    END;

  BEGIN DIV2[0]:= 0; LASTDIV2:= 0; EVEN:= TRUE;
    FOR I:= 1 TO K DO
      BEGIN EVEN:= NOT EVEN;
        IF EVEN THEN LASTDIV2:= LASTDIV2 + 1; DIV2[I]:= LASTDIV2
      END;
      RANKS[K]:= H; FILLRANKS(0, K, H - 1); CARDINALITY:= 0;
      NEW(ROOT, TRUE); WITH ROOT^ DO
        BEGIN RANK:= H; NEW(LEFT); NEW(RIGHT);
          CLEAR(ROOT); FOR I:= 0 TO H DO FATHERS[I]:= NIL
        END;
        INITIALIZE(ROOT, K - 1); ROOT^.POSITION:= COUNT;
        INITIALIZE(ROOT, K - 1)
      END;
    END;

```



```
PROCEDURE INSERT(J: INTEGER);
```

```
  VAR EJ, NB, MIN: PTR;
```

```
  FUNCTION INS(LEAF, TOP, PRES: PTR; NOBRANCHPOINT: BOOLEAN;
               ORDER: INTEGER): PTR;
```

```
    VAR HL, HP, NB: PTR; FPTR: FIELDPTR;
```

```
  BEGIN IF PRES = NIL THEN
```

```
    BEGIN FPTR:= MYFIELDS(LEAF, TOP); WITH FPTR^ DO
```

```
      BEGIN MIN:= LEAF; MAX:= LEAF END;
```

```
      IF LEAF^.POSITION <= TOP^.POSITION THEN
```

```
        INS:= TOP^.RIGHT^.MIN ELSE INS:= TOP^.LEFT^.MAX
```

```
      END ELSE
```

```
      BEGIN HL:= LEAF^.FATHERS[ORDER - 1]; HP:= PRES^.FATHERS[ORDER - 1];
```

```
      IF NOBRANCHPOINT THEN
```

```
        IF HP <> HL THEN
```

```
          BEGIN FPTR:= MYFIELDS(LEAF, HL); WITH FPTR^ DO
```

```
            BEGIN MIN:= LEAF; MAX:= LEAF END;
```

```
            FPTR:= MYFIELDS(PRES, HP); WITH FPTR^ DO
```

```
              BEGIN MIN:= PRES; MAX:= PRES END;
```

```
              HL^.UB:= PLUS; HP^.UB:= PLUS;
```

```
              NB:= INS(HL, TOP, HP, TRUE, ORDER - 1);
```

```
              INS:= PRES
```

```
          END ELSE
```

```
          BEGIN FPTR:= MYFIELDS(PRES, HP); WITH FPTR^ DO
```

```
            BEGIN MIN:= PRES; MAX:= PRES END;
```

```
            HP^.UB:= MINUS;
```

```
            INS:= INS(LEAF, HL, MYMIN(LEAF, HL), TRUE, ORDER - 1)
```

```
          END
```

```
        ELSE IF HL^.UB <> UNDEF THEN
```

```
          INS:= INS(LEAF, HL, MYMIN(LEAF, HL),
```

```
                MYMIN(LEAF, HL) = MYMAX(LEAF, HL), ORDER - 1) ELSE
```

```
          BEGIN FPTR:= MYFIELDS(LEAF, HL); WITH FPTR^ DO
```

```
            BEGIN MIN:= LEAF; MAX:= LEAF END;
```

```
            NOBRANCHPOINT:= HP^.UB = MINUS;
```

```
            HL^.UB:= PLUS; HP^.UB:= PLUS;
```

```
            NB:= INS(HL, TOP, HP, NOBRANCHPOINT, ORDER - 1);
```

```
            IF HL^.POSITION <= NB^.POSITION THEN
```

```
              INS:= MINOF(NB) ELSE INS:= MAXOF(NB);
```

```
          END;
```

```
          FPTR:= MYFIELDS(LEAF, TOP);
```

```
          IF TOP^.RANK = ORDER THEN WITH FPTR^ DO
```

```
            IF LEAF^.POSITION < MIN^.POSITION THEN MIN:= LEAF ELSE
```

```
            IF LEAF^.POSITION > MAX^.POSITION THEN MAX:= LEAF
```

```
          END
```

```
        END;
```

```
  BEGIN IF NOT MEMBER(J) THEN
```

```
    BEGIN EJ:= ELEMENTS[J]; MIN:= MYMIN(EJ, ROOT);
```

```
    NB:= INS(EJ, ROOT, MIN, MIN = MYMAX(EJ, ROOT), H);
```

```
    IF NB = NIL THEN INTOLIST(NIL, EJ, NIL) ELSE
```

```
    IF NB^.POSITION < J THEN INTOLIST(NB, EJ, NB^.SUCC) ELSE
```

```
    INTOLIST(NB^.PRED, EJ, NB)
```

```
  END
```

```
END;
```

```
PROCEDURE DELETE(J: INTEGER);
```

```
  VAR EJ, P1, P2: PTR; NOBRPT: BOOLEAN;
```

```
  PROCEDURE DEL(LEAF, TOP: PTR; ORDER: INTEGER; VAR PRES1, PRES2: PTR;
    VAR NOBRANCHPOINT: BOOLEAN);
```

```
    VAR FPTR: FIELDPTR; HL, HP: PTR;
```

```
  BEGIN FPTR:= MYFIELDS(LEAF, TOP); WITH FPTR^ DO IF MIN = MAX THEN
```

```
    BEGIN MIN:= NIL; MAX:= NIL;
```

```
      PRES1:= YOURMIN(LEAF, TOP); PRES2:= YOURMAX(LEAF, TOP);
```

```
      NOBRANCHPOINT:= PRES1 = PRES2
```

```
    END ELSE
```

```
    BEGIN HL:= LEAF^.FATHERS[ORDER - 1];
```

```
      IF MINOF(HL) = MAXOF(HL) THEN
```

```
        BEGIN DEL(HL, TOP, ORDER - 1, PRES1, PRES2, NOBRANCHPOINT);
```

```
          CLEAR(HL); HP:= PRES1;
```

```
          PRES1:= MINOF(PRES1); PRES2:= MAXOF(PRES2);
```

```
          IF NOBRANCHPOINT THEN
```

```
            IF (PRES1 = PRES2) THEN CLEAR(HP)
```

```
            ELSE BEGIN NOBRANCHPOINT:= FALSE; HP^.UB:= MINUS END
```

```
          END ELSE
```

```
          BEGIN DEL(LEAF, HL, ORDER - 1, PRES1, PRES2, NOBRANCHPOINT);
```

```
            IF NOBRANCHPOINT THEN
```

```
              IF (HL^.UB = MINUS) AND (YOURMIN(PRES1, HL) = NIL)
```

```
              THEN CLEAR(HL)
```

```
              ELSE NOBRANCHPOINT:= FALSE
```

```
            END;
```

```
            IF TOP^.RANK = ORDER THEN
```

```
              IF MIN = LEAF THEN MIN:= PRES1 ELSE
```

```
              IF MAX = LEAF THEN MAX:= PRES2
```

```
            END
```

```
  END;
```

```
BEGIN IF MEMBER(J) THEN
```

```
  BEGIN EJ:= ELEMENTS[J]; DEL(EJ, ROOT, H, P1, P2, NOBRPT);
```

```
    FROMLIST(EJ^.PRED, EJ, EJ^.SUCC)
```

```
  END
```

```
END;
```

```
PROCEDURE EXTRACTMIN; BEGIN DELETE(MIN) END;
```

```
PROCEDURE EXTRACTMAX; BEGIN DELETE(MAX) END;
```

```
PROCEDURE ALLMIN(J: INTEGER);
```

```
  VAR M: INTEGER;
```

```
BEGIN M:= MIN; WHILE (M > 0) AND (M <= J) DO
```

```
  BEGIN DELETE(M); M:= MIN END
```

```
END;
```

```
PROCEDURE ALLMAX(J: INTEGER);
```

```
  VAR M: INTEGER;
```

```
BEGIN M:= MAX; IF J < 1 THEN J:= 1; WHILE M >= J DO
```

```
  BEGIN DELETE(M); M:= MAX END
```

```
END;
```

```

FUNCTION NEIGHBOUR(J: INTEGER): INTEGER;

  VAR PT: PTR;

  FUNCTION NEIGHB(LEAF, TOP, PMIN, PMAX: PTR; ORDER: INTEGER): PTR;

    VAR Y, Z, NB, HL: PTR; POS: 1..N;

  BEGIN POS:= LEAF^.POSITION;
    IF (PMIN = NIL) OR ((PMIN = PMAX) AND (PMIN = LEAF)) THEN
      IF POS <= TOP^.POSITION THEN NEIGHB:= YOURMIN(LEAF, TOP)
      ELSE NEIGHB:= YOURMAX(LEAF, TOP)
    ELSE IF PMIN^.POSITION > POS THEN NEIGHB:= PMIN
    ELSE IF PMAX^.POSITION < POS THEN NEIGHB:= PMAX
    ELSE
      BEGIN HL:= LEAF^.FATHERS[ORDER - 1]; Y:= MINOF(HL); Z:= MAXOF(HL);
        IF ((Y = Z) AND (Y = LEAF)) OR (HL^.UB = UNDEF) THEN
          BEGIN NB:= NEIGHB(HL, TOP, PMIN^.FATHERS[ORDER - 1],
            PMAX^.FATHERS[ORDER - 1], ORDER - 1);
            IF HL^.POSITION < NB^.POSITION THEN NEIGHB:= MINOF(NB)
            ELSE NEIGHB:= MAXOF(NB)
          END
        ELSE NEIGHB:= NEIGHB(LEAF, HL, MYMIN(LEAF, HL),
          MYMAX(LEAF, HL), ORDER - 1)
      END
    END
  END;

  BEGIN PT:= ELEMENTS[J];
    PT:= NEIGHB(PT, ROOT, MYMIN(PT, ROOT), MYMAX(PT, ROOT), H);
    IF PT = NIL THEN NEIGHBOUR:= 0 ELSE NEIGHBOUR:= PT^.POSITION
  END;

FUNCTION PREDECESSOR(J: INTEGER): INTEGER;

  VAR PT: PTR; NB: INTEGER;

  BEGIN IF NOT MEMBER(J) THEN NB:= NEIGHBOUR(J) ELSE NB:= J;
    IF NB < J THEN PREDECESSOR:= NB ELSE
      BEGIN PT:= ELEMENTS[NB]^.PRED;
        IF PT = NIL THEN PREDECESSOR:= 0 ELSE PREDECESSOR:= PT^.POSITION
      END
    END;

FUNCTION SUCCESSOR(J: INTEGER): INTEGER;

  VAR PT: PTR; NB: INTEGER;

  BEGIN IF NOT MEMBER(J) THEN NB:= NEIGHBOUR(J) ELSE NB:= J;
    IF NB = 0 THEN SUCCESSOR:= N + 1 ELSE
      IF NB > J THEN SUCCESSOR:= NB ELSE
        BEGIN PT:= ELEMENTS[NB]^.SUCC;
          IF PT = NIL THEN SUCCESSOR:= N + 1 ELSE SUCCESSOR:= PT^.POSITION
        END
      END;

PROCEDURE USERSPROGRAM;

  BEGIN (* TO BE SUPPLIED BY THE USER *) END;

  BEGIN START; USERSPROGRAM END.

```

ACKNOWLEDGEMENTS

The efficient priority queue was designed by the first author during a three months visit to the Cornell dept. of Computer science, supported by ZWO grant CR 62/50. Thanks are due to J. Hopcroft for suggesting the problem and other ideas and to Z. Galil for checking the original argumentations and other usefull suggestions. The mergeable heap was obtained on base of suggestions by R.E. Tarjan.

The PASCAL implementation was designed by the remaining authors at the University of Amsterdam. During the process of implementing the original structure was improved and all bugs in the original desk-written algorithms in [3] were eliminated. Our gratitude goes to the Institute for Appl. Math. of the University of Amsterdam, and to T.J. Dekker in particular for offering the technical facilities to complete our work.

Author's addresses:

P. van Emde Boas & R. Kaas: Mathematical Centre, 2e Boerhaavestr. 49,
Amsterdam. Netherlands.

P. van Emde Boas & E. Zijlstra: Mathematical Institute, University of
Amsterdam. Roetersstraat 15. Amsterdam. Netherlands.

REFERENCES

- [1] AHO, A.V., J.E. HOPCROFT & J.D. ULLMAN, *The design and analysis of computer Algorithms*, Addison Wesley, Reading, Mass. (1974).

- [2] AHO, A.V., J.E. HOPCROFT & J.D. ULLMAN, *On finding lowest common ancestors in trees*, Proc. 5-th ACM symp. Theory of Computing (1973), 253-265.
- [3] EMDE BOAS, P. VAN, *An $O(n \log \log n)$ On-Line Algorithm for the Insert-Extract Min problem*, Rep. TR 74-221 Dept. of Comp. Sci., Cornell Univ., Ithaca 14853, N.Y. Dec. 1974.
- [4] EVEN, S. & O. KARIV. *Oral Commun.*, Berkeley, October 1975.
- [5] FISHER, M.J., *Efficiency of equivalence algorithms*, in: R.E. MILLER & J.W. THATCHER (eds.), *Complexity of Computer Computations*, Plenum Press, New York (1972), 158-168.
- [6] HOPCROFT, J. & J.D. ULLMAN, *Set-merging Algorithms*, SIAM J. Comput. 2 (Dec. 1973), 294-303.
- [7] HUNT, J.W. & T.G. SZYMANSKI, *A fast algorithm for computing longest common subsequences*. Manuscript. Dept. Electr. Eng. Princeton Univ. Princeton, N.J. 08540. Oct. 1975.
- [8] TARJAN, R.E. *Applications of path compression on balanced trees*. Manuscript. Stanford Oct. 75. (Submitted to JACM).
- [9] TARJAN, R.E., *Efficiency of a good but non linear set union algorithm*, J. Assoc. Comput. Mach. 22 (1975), 215-224.
- [10] TARJAN, R.E., *Edge disjoint spanning trees, dominators and depth first search*, Rep. CS-74-455 (Sept. 1974), Stanford.

- [11] WIRTH, N., *The Programming Language PASCAL (revised report)*, in
K. JENSEN & N. WIRTH *PASCAL User Manual and Report*, Lecture Notes
in Computer Science 18, Springer, Berlin (1974).




Diagram 1. Example of a five-element set representation using mark bits.






Diagram 2: The canonical subtrees of v . $R(v)$ is the shaded area.




	LEVEL	RANK	LMIN	LMAX	RMIN	RMAX	UB
t	16	4	a	c	~	~	-
j	12	2	h	h	i	i	-
h	8	3	~	~	a	b	+
i	8	3	~	~	c	c	+
g	6	1	d	d	e	e	-
d	4	2	a	a	~	~	+
e	4	2	~	~	b	b	+
f	4	2	~	~	~	~	~


Diagram 3: Example of a proper information content.




hp inactive; hp = hl
bottom call



hp inactive; hp \neq hl
top call



hp active; hl active
bottom call



hp active; hl inactive
top call

Diagram 4: The four possible inner call situations for ins.




Diagram 5: del inactivates a present node




Diagram 6: The short-cut in neighb

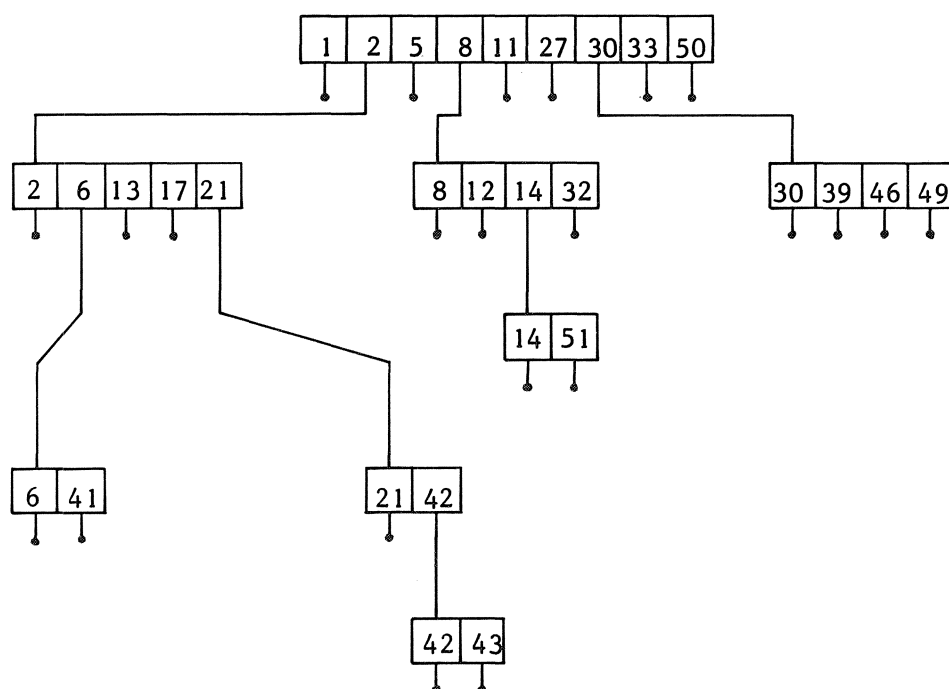


Diagram 7: A Triter tree with priority queues at its internal nodes.

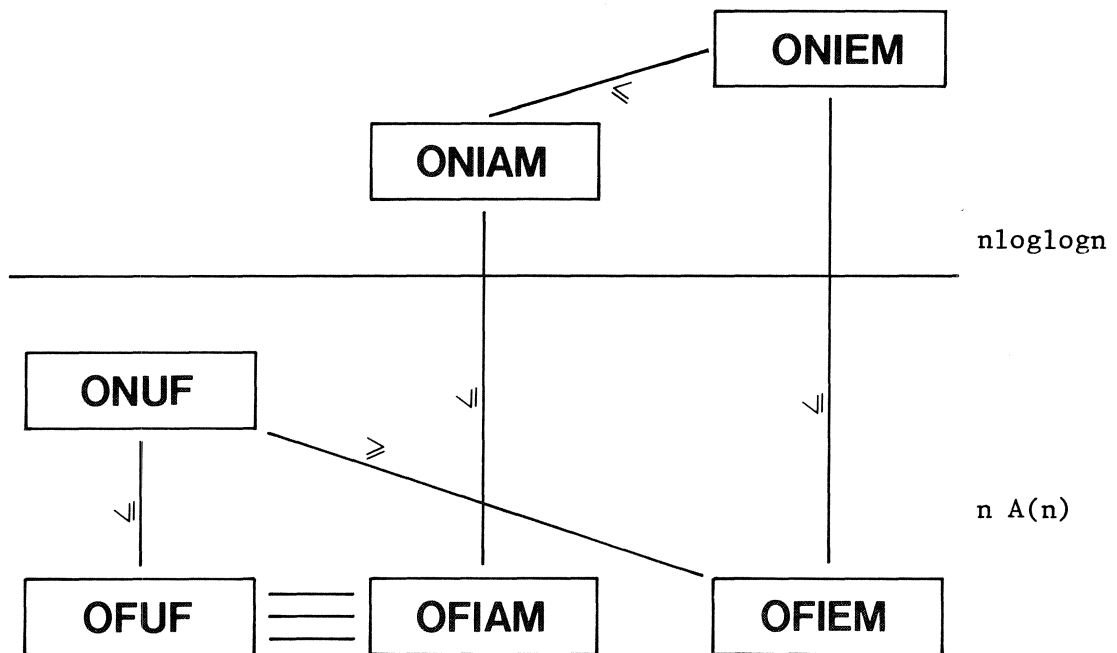


Diagram 8: Reducibilities among set-manipulation problems