

A Simple Implementation Technique for Priority Search Queues

Ralf Hinze

Institute of Information and Computing Sciences, Utrecht University
P.O.Box 80.089, 3508 TB Utrecht, The Netherlands

ralf@cs.uu.nl

ABSTRACT

This paper presents a new implementation technique for priority search queues. This abstract data type is an amazing blend of finite maps and priority queues. Our implementation supports logarithmic access to a binding with a given key and constant access to a binding with the minimum value. Priority search queues can be used, for instance, to give a simple, purely functional implementation of Dijkstra's single-source shortest-paths algorithm.

A non-technical concern of the paper is to foster abstract data types and views. Priority search queues have been largely ignored by the functional programming community and we believe that they deserve to be known better. Views prove their worth both in defining a convenient interface to the abstract data type and in providing a readable implementation.

Categories and Subject Descriptors

D.1.1 [Programming Techniques]: Applicative (Functional) Programming; D.3.2 [Programming Languages]: Language Classifications—*applicative (functional) languages*; D.3.3 [Programming Languages]: Language Constructs and Features—*abstract data types*; E.1 [Data]: Data Structures—*trees*; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*sorting and searching*; I.1.2 [Computing Methodologies]: Algorithms—*analysis of algorithms*

General Terms

Algorithms, design, performance

Keywords

Priority search queues, views, Haskell, tournament

1. INTRODUCTION

The aim of this paper is threefold:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP'01, September 3–5, 2001, Florence, Italy.

Copyright 2001 ACM 1-58113-415-0/01/0009 ...\$5.00.

First, we would like to advertise *priority search queues*, a useful abstract data type that has been largely ignored by the functional programming community and that deserves to be known better. Priority search queues are an amazing blend of finite maps (or dictionaries) and priority queues, that is, they support both dictionary operations (for instance, accessing a binding with a given key) and priority queue operations (for instance, accessing a binding with the minimum value). We give two simple applications that demonstrate their usefulness: a purely functional implementation of Dijkstra's single-source shortest-paths algorithm and an efficient implementation of the first-fit heuristics for the bin packing problem.

Second, we describe a simple implementation technique for the abstract data type. The standard implementation of priority search queues, McCreight's priority search trees [14], combines binary search trees and heaps. Unfortunately, balanced search trees and heaps do not go well together. Rotations that are typically used to maintain balance destroy the heap property and restoring the property takes $\Theta(h)$ time where h is the height of the tree. Consequently, in order to attain overall logarithmic time bounds the underlying balancing scheme must guarantee that the number of rotations per update is bounded by a constant. We show that it is possible to weaken the heap property so that rotations become constant time operations without sacrificing the running time of the priority queue methods. Thus, we can freely choose an underlying balancing scheme—we illustrate our approach using weight-balanced trees [1].

Third, we would like to promote the use of *views*. Views have been introduced by Wadler [23] to relieve the tension between pattern matching and abstraction. Briefly, views allow any type (in particular, any abstract data type) to be viewed as a free data type. We have found views not only useful for providing a convenient interface to an abstract data type but also extremely helpful in the implementation itself. The use of views made the code substantially clearer.

The remainder of this paper is structured as follows. Section 2 briefly reviews the concept of views. Section 3 introduces the abstract data type *priority search queue* and Section 4 illustrates its use. Section 5 provides a simple implementation based on unbalanced trees. Section 6 then shows how to augment the basic implementation by a balancing scheme. Section 7 analyses the running time of so-called range queries. Finally, Section 8 reviews related work and Section 9 concludes.

2. PRELIMINARIES: VIEWS

The code in this paper is given in Haskell 98 [19] augmented by the concept of views [4, 16]. This section briefly reviews Okasaki’s proposal for views [16].

A view allows any type to be viewed as a free data type. A *view declaration* for a type T consists of an anonymous data type, the *view type*, and an anonymous function, the *view transformation*, that shows how to map elements of T to the view type. Here is a simple example that defines a *minimum view* on lists:

```
view (Ord a) => [a] = Empty | Min a [a] where
  []           -> Empty
  a1 : Empty  -> Min a1 []
  a1 : Min a2 as
  | a1 <= a2  -> Min a1 (a2 : as)
  | otherwise -> Min a2 (a1 : as).
```

This declaration introduces two constructors, *Empty* and *Min*, that henceforth can be used to pattern match elements of type $[a]$, where the context ‘(Ord a) =>’ restricts a to instances of *Ord*. The minimum view allows any list to be viewed as an ordered list. The following definition of selection sort nicely illustrates the use of views:

```
selection-sort      :: (Ord a) => [a] -> [a]
selection-sort Empty = []
selection-sort (Min a as) = a : selection-sort as.
```

The view constructors can be freely mixed with ordinary data type constructors. In fact, the view transformation of the minimum view already illustrates nested patterns. A type can even have multiple views. However, view constructors may only appear in patterns—with the notable exception of the view transformation itself.

View declarations can be implemented by a simple source to source translation: each view is expanded into a data type and a function. For the minimum view we obtain:

```
data Min-View a      = Empty | Min a [a]
min-view             :: (Ord a) => [a] -> Min-View a
min-view x1          = case x1 of
  []                 -> Empty
  a1 : x2            -> case min-view x2 of
    Empty           -> Min a1 []
    Min a2 as       ->
      | a1 <= a2    -> Min a1 (a2 : as)
      | otherwise   -> Min a2 (a1 : as).
```

The function is invoked whenever constructors of the view appear in patterns. In our example, the view constructors appear in the view transformation itself. Consequently, it is expanded into a recursive function. Selection sort becomes:

```
selection-sort x = case minView x of
  Empty   -> []
  Min a as -> a : selection-sort as.
```

For a precise definition of the semantics we refer the interested reader to Okasaki’s paper [16]—the proposal is for Standard ML but it can be easily adapted to Haskell 98.

3. PRIORITY SEARCH QUEUES

The abstract data type *priority search queue* is conceptually a finite map that supports efficient access to the binding

with the minimum value, where a *binding* is an argument-value pair and a *finite map* is a finite set of bindings. For emphasis, we call the arguments *keys* and the associated values *priorities*. The functions *key* and *prio* are used to access the key and the priority of a binding.

```
key (k, p) = k
prio (k, p) = p.
```

The abstract data type of priority search queues is parametric in the types of keys and priorities:

```
data PSQ k p.
```

Most operations on priority search queues require that both k and p are totally ordered. This condition is expressed in Haskell by the context ‘(Ord k, Ord p) =>’. However, in the rest of this paper we will omit the context to reduce clutter.

Priority search queues support both finite map and priority queue operations plus so-called *range queries*.

Constructors and insertion.

```
∅           :: PSQ k p
{·}         :: (k, p) -> PSQ k p
insert      :: (k, p) -> PSQ k p -> PSQ k p
from-ord-list :: [(k, p)] -> PSQ k p
```

The constructor \emptyset represents the empty queue; $\{b\}$ creates a queue that contains b as the single binding; *insert* b q inserts binding b into q (if the queue contains a binding with the same key, then the old binding is overwritten); and *from-ord-list* converts a list of bindings into a queue with the precondition that the list is sorted into increasing order by key.

Destructors and deletion.

```
view PSQ k p = Empty | Min (k, p) (PSQ k p)
delete       :: k -> PSQ k p -> PSQ k p
```

A queue is destructed using the patterns *Empty* and *Min b q* introduced by the view declaration. The function *delete* removes a binding with the given key (the queue is left unchanged if it does not contain a binding with the key). The constructors of the *minimum view* have the following meaning: if a queue pattern matches *Empty*, then it is empty; otherwise it matches *Min b q* where b is the binding with the minimum priority and q is the remaining queue. Thus, using the view we can effectively treat a priority search queue as a list of bindings ordered by priority.

Observers.

```
lookup      :: k -> PSQ k p -> Maybe p
to-ord-list :: PSQ k p -> [(k, p)]
at-most     :: p -> PSQ k p -> [(k, p)]
```

The function *lookup* finds the priority associated with a given key: the call *lookup* k q returns *Nothing* if the queue does not contain the key k ; otherwise it yields *Just p* where p is the priority associated with k . The function *to-ord-list* converts a queue into a list of bindings ordered by key. Priority search queues not only support dictionary and priority queue operations. As a little extra they also allow for so-called *range queries*: *at-most* p_i q returns a list of bindings ordered by key whose priorities are at most p_i . In the full

version of the paper [9] we also discuss range queries that additionally take a key range into account.

Modifier.

$adjust :: (p \rightarrow p) \rightarrow k \rightarrow PSQ\ k\ p \rightarrow PSQ\ k\ p$

The function *adjust* changes a binding for the given key by applying the function to its priority (the queue is left unchanged if it does not contain a binding with the key).

4. APPLICATIONS

4.1 Single-source shortest-paths problem

Dijkstra’s algorithm for the single-source shortest-paths problem serves as a nice example for the use of priority search queues. The algorithm maintains a queue that maps each vertex to its estimated distance from the source. The algorithm works by repeatedly removing the vertex with minimal distance and updating the distances of its adjacent vertices. Priority search queues support both operations equally well. The update operation is typically called *decrease*:

$decrease :: (k, p) \rightarrow PSQ\ k\ p \rightarrow PSQ\ k\ p$
 $decrease\ (k, p)\ q = adjust\ (min\ p)\ k\ q$
 $decrease-list :: [(k, p)] \rightarrow PSQ\ k\ p \rightarrow PSQ\ k\ p$
 $decrease-list\ bs\ q = foldr\ decrease\ q\ bs.$

Note that *decrease* $(k, p)\ q$ has no effect if k ’s priority in q is less than p .

To keep the presentation terse we assume that the following functions on graphs are provided from somewhere.

$vertices :: Graph \rightarrow [Vertex]$
 $adjacent :: Graph \rightarrow Vertex \rightarrow [Vertex]$

The function *vertices* returns an *ordered* list of all vertices of a graph; *adjacent* produces a list of vertices adjacent to the given one.

The function *dijkstra* defined below takes three arguments: a directed graph, a weight function, and a source vertex. It returns a list of vertex-distance bindings that determine the minimal distance of each vertex from the source.

type *Weight* = $Vertex \rightarrow Vertex \rightarrow Double$
 $dijkstra :: Graph \rightarrow Weight \rightarrow Vertex \rightarrow [(Vertex, Double)]$
 $dijkstra\ g\ w\ s = loop\ (decrease\ (s, 0)\ q_0)$
where
 $q_0 = from-ord-list\ [(v, +\infty) \mid v \leftarrow vertices\ g]$
 $loop\ Empty = []$
 $loop\ (Min\ (u, d)\ q) = (u, d) : loop\ (decrease-list\ bs\ q)$
where $bs = [(v, d + w\ u\ v) \mid v \leftarrow adjacent\ g\ u]$

The helper function *loop* uses the minimum view to process the queue. Note that the computed list of vertex-distance bindings may contain bindings with priority $+\infty$, which indicates that the given graph was not strongly connected. Now, if we assume that the computation of the view and the *decrease* operation each take $\Theta(\log V)$ time, then the algorithm has a worst-case running time of $\Theta((V + E) \log V)$, which is the best known running time for purely functional implementations.

REMARK 1. *If we modify the computation of the new distances as follows*

\dots **where** $bs = [(v, w\ u\ v) \mid v \leftarrow adjacent\ g\ u],$

we obtain Prim’s algorithm for computing a minimum spanning tree.

4.2 One-dimensional bin packing

As the second example we employ priority search queues to implement the first-fit heuristics for the bin packing problem. Recall that the standard list-based implementation shown below has a worst-case running time of $\Theta(n^2)$ where n is the number of items.

$pack-first-fit :: [Item] \rightarrow [Bin]$
 $pack-first-fit = foldl\ first-fit\ []$
 $first-fit :: [Bin] \rightarrow Item \rightarrow [Bin]$
 $first-fit\ []\ i = [i]$
 $first-fit\ (b : bs)\ i$
 $\quad \mid b + i \leq 1 = b + i : bs$
 $\quad \mid otherwise = b : first-fit\ bs\ i$

The function *pack-first-fit* takes a list of items, each of a certain size, and returns a list of bins that contain the input items. For simplicity, we represent an item by its size and a bin by its total size (each bin has a capacity of 1).

Using priority search queues we can improve the running time of the naïve implementation to $\Theta(n \log n)$. The central idea is to use the function *at-most* to quickly determine the first bin that can accommodate a given item (the bins are numbered consecutively).

type *No* = Int
 $pack-first-fit :: [Item] \rightarrow [Bin]$
 $pack-first-fit\ is = [prio\ b \mid b \leftarrow to-ord-list\ q]$
where $(q, -) = foldl\ first-fit\ (\emptyset, 0)\ is$
 $first-fit :: (PSQ\ No\ Bin, No) \rightarrow Item \rightarrow (PSQ\ No\ Bin, No)$
 $first-fit\ (q, n)\ i = \mathbf{case}\ at-most\ (1 - i)\ q\ \mathbf{of}$
 $\quad [] \rightarrow (insert\ (n, i)\ q, n + 1)$
 $\quad (k, -) : - \rightarrow (adjust\ (+i)\ k\ q, n)$

This is the only place where essential use is made of Haskell’s non-strict semantics as we merely require the first element of the list returned by *at-most*. In a strict language, we would be forced to define a specialized version of *at-most* that computes the first binding only (if any).

5. PRIORITY SEARCH PENNANTS

This section describes an implementation of priority search queues based on unbalanced search trees. Great care has been taken to modularize the code so that a balancing scheme can be added later with ease (Section 6 discusses the necessary amendments). It should be noted, however, that the implementation in this section is perfectly suitable for Dijkstra’s or Prim’s algorithm since both do not require insertions.

The underlying idea of the implementation is best explained using the metaphor of a knockout tournament. Consider the tournament depicted in Figure 1. We have eight participants, so the course of matches forms a complete binary tree. Each external node corresponds to a participant; each internal node corresponds to a winner of a match. To facilitate searching the participants are arranged from left

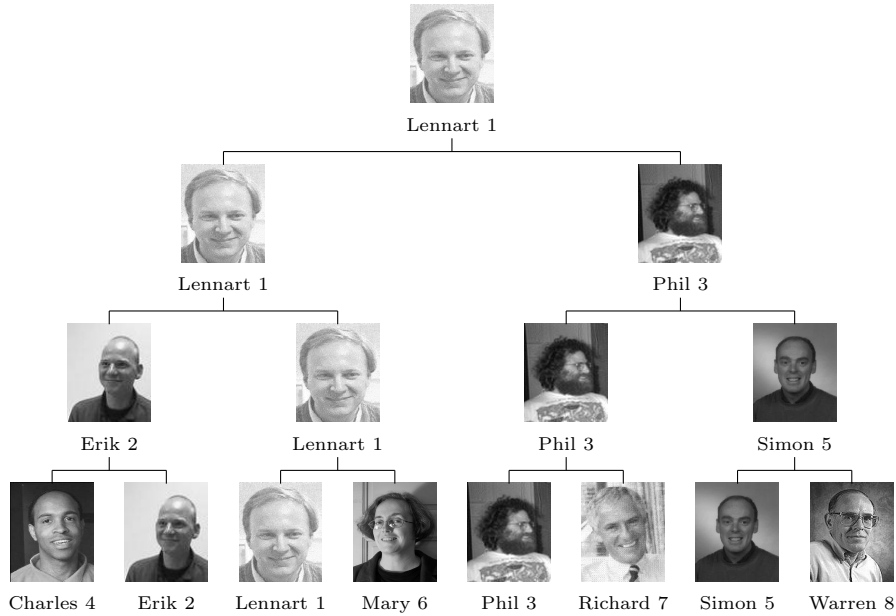


Figure 1: A tournament tree.

to right in increasing order by name. Tournament trees are almost a suitable data structure for priority search queues if it were not for the many repeated entries. The champion, for instance, appears on every level of the tree. Now, there are at least two ways to repair this defect.

One possibility is to promote losers up the tree turning the tournament tree of Figure 1 into the heap-structured tree of Figure 2. This transformation usually involves additional matches. In our example, Erik has to play with Mary to determine the second-best player of the first half of the tournament. Pursuing this idea further leads to a data structure known as a *priority search tree* [14]. We will come back to this data structure in Sections 7 and 8.

An alternative possibility, which we will investigate in this section, is to label each internal node with the loser of the match, instead of the winner, and to drop the external nodes altogether. If we additionally place the champion on top of the tree, we obtain the topped loser tree of Figure 3. We call the resulting data structure *priority search pennant*. Since every participant—with the notable exception of the champion—loses exactly one match, the pennant does not contain repeated entries. It is important to note, however, that the loser tree is not heap-structured. Since the nodes are labelled with losers, they dominate, in general, only one subtree. The node labelled Phil, for instance, dominates its right but not its left subtree. Thus the loser tree constitutes only a so-called *semi-heap*.

The Haskell data type for priority search pennants is a direct implementation of the above ideas except that we additionally introduce split keys (or search keys) to support searching.

```

data PSQ k p = Void
              | Winner (k, p) (LTree k p) k
data LTree k p = Start
              | Loser (k, p) (LTree k p) k (LTree k p)

```

Here, *Void* represents the empty tournament; *Winner b t m* represents a tournament that *b* has won, *t* is the associated loser tree and *m* is the maximum key. Likewise, *Start* is the empty loser tree; *Loser b t_l k t_r* represents a subtournament that *b* has lost, *t_l* is the left subtree, *k* is the split key, and *t_r* is the right subtree. The maximum key is usually accessed using the function *max-key*.

$$\begin{aligned} \text{max-key} & & & \text{:: } \text{PSQ } k \ p \ \rightarrow \ k \\ \text{max-key } (\text{Winner } b \ t \ m) & = \ m \end{aligned}$$

We will see in Section 5.1 why it is useful to keep track of the maximum key.

Priority search pennants combine the features of search trees and semi-heaps. To formulate the invariants, it is convenient to view the top node *Winner b t m* as a binary node with an empty right subtree so that the maximum key becomes an ordinary split key ($\text{Winner } b \ t \ m \cong \text{Loser } b \ t \ m \ \text{Start}$).

Semi-heap conditions: 1) Every priority in the pennant must be greater than or equal to the priority of the winner. 2) For all nodes in the loser tree, the priority of the loser’s binding must be less than or equal to the priorities of the bindings of the subtree from which the loser originates. The loser *originates* from the left subtree if its key is less than or equal to the split key, otherwise it originates from the right subtree.

Search-tree condition: For all nodes, the keys in the left subtree must be less than or equal to the split key and the keys in the right subtree must be greater than the split key.

Key condition: The maximum key and the split keys must also occur as keys of bindings.

Finite map condition: The pennant must not contain two bindings with the same key.

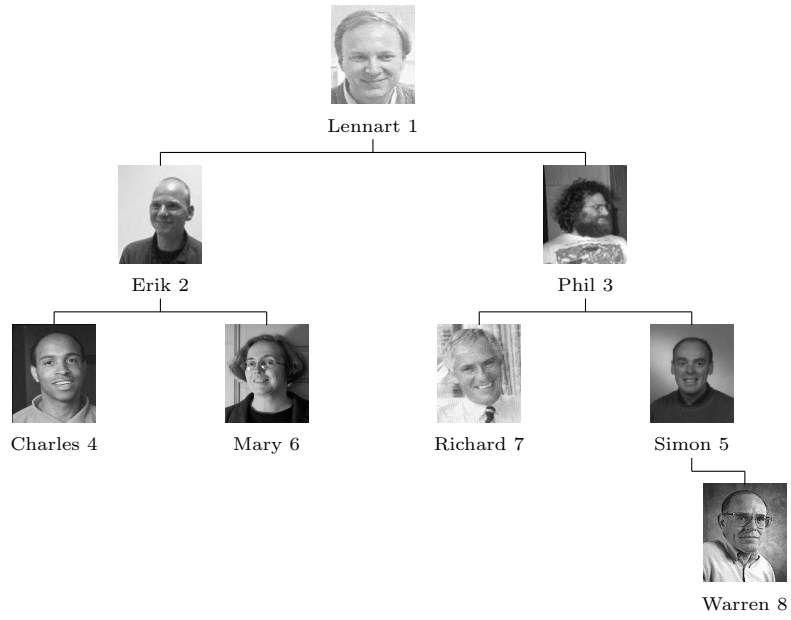


Figure 2: The heap corresponding to the tournament of Figure 1.

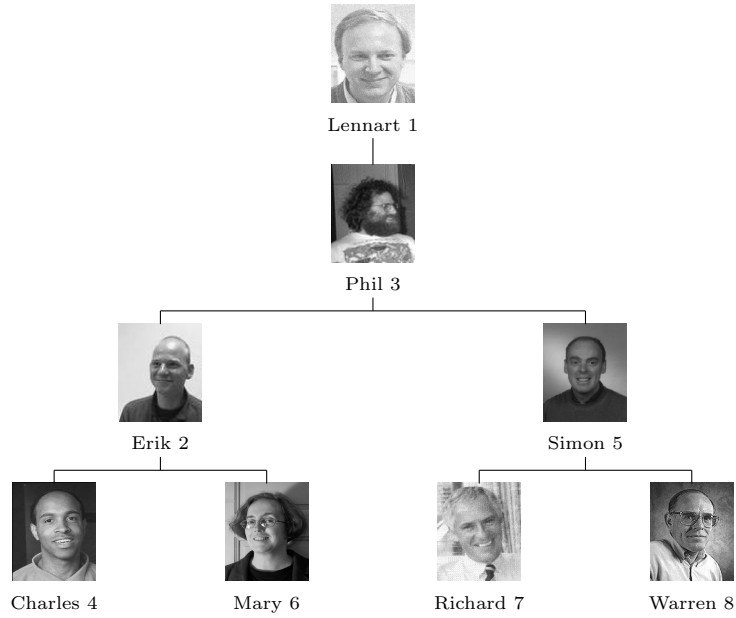


Figure 3: The semi-heap corresponding to the tournament of Figure 1.




Figure 4: The priority search pennant corresponding to the tree of Figure 3.

Two remarks are in order. First, the second semi-heap condition shows that a priority search pennant contains enough information to reconstruct the original tournament tree. This ability is crucial for implementing the priority queue operations. Second, the key condition ensures that every search key originates from a binding in the tree. This means, in particular, that if we delete a binding from a tree, we must also delete the key's second occurrence as a search key. We will see that it is relatively easy to maintain this invariant.

Let us consider an example. If we augment the tree of Figure 3 by split keys, we obtain the priority search pennant depicted in Figure 4. Note that the dotted lines mark the subtrees that are *not* dominated by the loser. As we have remarked before, the semi-heap structure can also be determined by comparing the loser's key to the split key: the node labelled Phil, for instance, dominates its right subtree since $P > M$; the node labelled Erik on the other hand dominates its left subtree since $E \leq E$. The pennant can quite easily be expressed as a Haskell term:

```
Winner (L,1) (
  Loser (P,3) (
    Loser (E,2) (
      Loser (C,4) Start C Start) E (
        Loser (M,6) Start L Start)) M (
      Loser (S,5) (
        Loser (R,7) Start P Start) R (
          Loser (W,8) Start S Start))) W.
```

Note that if we list the search keys from left to right, we obtain the keys of the participants in increasing order.

REMARK 2. *The nodes are decorated with bindings of type (k, p) . While this is convenient for the presentation, it comes at a small run-time cost since every access involves one extra level of indirection. In the production code, which is available from <http://www.cs.uu.nl/~ralf/software>, we speed up the access by storing the keys and the priorities directly in the nodes.*

5.1 Constructors

The empty queue and the singleton queue are defined as follows:

```
∅ :: PSQ k p
∅ = Void
{·} :: (k, p) → PSQ k p
{b} = Winner b Start (key b).
```

The data types *PSQ* and *LTree* have been designed to efficiently support the binary operation (\wedge) , which corresponds to playing a match. This operation, which is used by most of the remaining functions, takes two pennants and returns a new pennant that is the union of the two with the precondition that the keys in the first tree are strictly smaller than the keys in the second tree. The operation is illustrated in Figure 5.

```
(∧) :: PSQ k p → PSQ k p → PSQ k p
Void ∧ t' = t'
t ∧ Void = t
Winner b t m ∧ Winner b' t' m'
| prio b ≤ prio b' = Winner b (Loser b' t m t') m'
| otherwise = Winner b' (Loser b t m t') m'
```

Note that in order to construct the loser tree we require a split key, which is why we keep track of the maximum key in the top node. This makes ' \wedge ' a constant-time operation. It is not hard to see that ' \wedge ' preserves the invariants of priority search pennants. Using ' \wedge ' we can easily define *from-ord-list*.

```
from-ord-list :: [(k, p)] → PSQ k p
from-ord-list = foldm (∧) ∅ · map (λb → {b})
```

The helper function *foldm*, which is listed in the full version of this paper [9], folds a list in a binary-sub-division fashion. For instance,

```
from-ord-list [(A,4), (D,2), (E,1), (J,6),
              (L,3), (N,7), (P,5), (V,8)]
```

reduces to

```
(({A,4} ∧ {D,2}) ∧ ({E,1} ∧ {J,6}))
 ∧ (({L,3} ∧ {N,7}) ∧ ({P,5} ∧ {V,8})),
```

which in turn evaluates to the tree of Figure 4. In general, the expression tree generated by *foldm* takes the form of a leaf-oriented Braun tree [3]. Since ' \wedge ' preserves the shape of the expression tree, the priority search pennant produced by *from-ord-list* corresponds to a topped Braun tree. This means, in particular, that the shape is solely determined by the total number of participants (and not by their priorities).

5.2 Destructors

The minimum view is implemented as follows:

```
view PSQ k p = Empty | Min (k, p) (PSQ k p) where
Void → Empty
Winner b t m → Min b (second-best t m).
```

The function *second-best* used in the second clause determines the second-best player by replaying the tournament without the champion.

```
second-best :: LTree k p → k → PSQ k p
second-best Start m = Void
second-best (Loser b t k u) m
| key b ≤ k = Winner b t k ∧ second-best u m
| otherwise = second-best t k ∧ Winner b u m
```

Note that only those players who lost to the champion are taken into account. The origin of the champion is determined by comparing the loser's key to the split key.

Again, it is straightforward to see that *second-best* preserves the invariants except perhaps for the key condition: does *second-best* also remove the search key of the champion? This is most easily shown if we define *second-best* on




Figure 5: Playing a match ($b_1 \leq b_2$ is shorthand for $\text{prio } b_1 \leq \text{prio } b_2$).

pennants instead of loser trees (we call this variant *del-min*):

```

del-min      :: PSQ k p → PSQ k p
del-min Void = Void
del-min (Winner b Start m)
  = Void
del-min (Winner b (Loser b' t k u) m)
  | key b' ≤ k = Winner b' t k λ del-min (Winner b u m)
  | otherwise  = del-min (Winner b t k) λ Winner b' u m.

```

Since the argument of *del-min* is always a legal pennant, m must equal *key* b in the second equation by virtue of the key condition. Furthermore, we know that b is the champion, since the champion is passed unchanged to the recursive calls. The function *second-best* can now be seen as a simple optimization: we have

$$\text{del-min (Winner } b \ t \ m) = \text{second-best } t \ m.$$

REMARK 3. When we replay a tournament we determine the origin of a loser by comparing the loser's key to the split key (*key* $b \leq k$). Instead of using this perhaps costly comparison, we can alternatively code the information into the constructors when building the tree:

```

data LTree k p = Start
  | LLosser (k, p) (LTree k p) k (LTree k p)
  | RLosser (k, p) (LTree k p) k (LTree k p).

```

This is, in fact, the representation we use in the production code. The original representation, however, is slightly easier to augment by a balancing scheme.

5.3 Observers

Views are not only convenient for the client of an abstract data type. They can also be tremendously helpful when implementing an abstract data type. The following declaration allows us to view a pennant as a tournament tree.

```

view PSQ k p      = ∅ | {k, p} | PSQ k p λ PSQ k p
  where
  Void            → ∅
  Winner b Start m → {b}
  Winner b (Loser b' tl k tr) m
    | key b' ≤ k → Winner b' tl k λ Winner b tr m
    | otherwise  → Winner b tl k λ Winner b' tr m

```

Note that we have taken the liberty of using \emptyset , $\{\cdot\}$ and λ also as constructors. There is little danger of confusion since the constructors of the view may only appear in patterns—with the notable exception of the view transformation itself—while the functions of the same name may only appear in expressions. The view transformation is essentially the inverse of the λ operation. In particular, if a

winner tree matches $t_l \lambda t_r$, then it is guaranteed that the keys in t_l are strictly smaller than the keys in t_r . Furthermore, both t_l and t_r are non-empty.

The function *to-ord-list*, which converts a queue into a list of bindings ordered by key, nicely illustrates the use of the tournament view.¹

```

to-ord-list      :: PSQ k p → [(k, p)]
to-ord-list ∅    = []
to-ord-list {b} = [b]
to-ord-list (tl λ tr) = to-ord-list tl ++ to-ord-list tr

```

In the last clause we rely on the fact that the keys in t_l precede the keys in t_r .

It is instructive to rewrite the definition of *to-ord-list* into a form that does not make use of views. We will see that the resulting code is much harder to read. On the other hand, the rewrite opens the possibility of small improvements (which a good optimizing compiler might be able to perform automatically). As the first step, we fuse the view transformation and the original function:

```

to-ord-list      :: PSQ k p → [(k, p)]
to-ord-list Void = []
to-ord-list (Winner b Start m)
  = [b]
to-ord-list (Winner b (Loser b' tl k tr) m)
  | key b' ≤ k = to-ord-list (Winner b' tl k)
  ++ to-ord-list (Winner b tr m)
  | otherwise  = to-ord-list (Winner b tl k)
  ++ to-ord-list (Winner b' tr m).

```

Note that in each of the recursive calls *to-ord-list* is passed a non-empty winner tree. Furthermore, the maximum key and the split keys are never used. This suggests specializing *to-ord-list (Winner b t m)* to *traverse b t*:

```

to-ord-list      :: PSQ k p → [(k, p)]
to-ord-list Void = []
to-ord-list (Winner b t m) = traverse b t
traverse         :: (k, p) → LTree k p → [(k, p)]
traverse b Start = [b]
traverse b (Loser b' tl k tr)
  | key b' ≤ k = traverse b' tl ++ traverse b tr
  | otherwise  = traverse b tl ++ traverse b' tr.

```

Most of the following functions can be optimized along these lines.

¹Due to the use of $(++)$ the definition of *to-ord-list* exhibits $\Theta(n^2)$ worst-case behaviour. This is, however, easily remedied using standard techniques.

The look-up function is very similar to the look-up function for binary search trees. Again, the tournament view allows for a very natural implementation.

$$\begin{aligned}
\text{lookup} &:: k \rightarrow \text{PSQ } k \ p \rightarrow \text{Maybe } p \\
\text{lookup } k \ \emptyset &= \text{Nothing} \\
\text{lookup } k \ \{b\} & \\
\quad | k == \text{key } b &= \text{Just } (\text{prio } b) \\
\quad | \text{otherwise} &= \text{Nothing} \\
\text{lookup } k \ (t_l \wedge t_r) & \\
\quad | k \leq \text{max-key } t_l &= \text{lookup } k \ t_l \\
\quad | \text{otherwise} &= \text{lookup } k \ t_r
\end{aligned}$$

The running time of *lookup* is proportional to the height of the tree even if we search for a binding that is high up in the tree. This observation suggests to additionally test the bindings on the search path at the cost of one additional comparison per recursive call. Of course, this change neither affects the worst-case nor the average-case running time.

$$\begin{aligned}
\text{lookup}' &:: k \rightarrow \text{PSQ } k \ p \rightarrow \text{Maybe } p \\
\text{lookup}' k \ (\text{Min } b \ q) & \\
\quad | k == \text{key } b &= \text{Just } (\text{prio } b) \\
\text{lookup}' k \ \emptyset &= \text{Nothing} \\
\text{lookup}' k \ \{b\} &= \text{Nothing} \quad \text{-- we know that } k \neq \text{key } b \\
\text{lookup}' k \ (t_l \wedge t_r) & \\
\quad | k \leq \text{max-key } t_l &= \text{lookup}' k \ t_l \\
\quad | \text{otherwise} &= \text{lookup}' k \ t_r
\end{aligned}$$

Note that this version of the look-up function uses both the minimum and the tournament view.

5.4 Modifier, insertion, and deletion

The dictionary functions *adjust*, *insert*, and *delete* can be most easily implemented using the tournament view.

$$\begin{aligned}
\text{adjust} &:: (p \rightarrow p) \rightarrow k \rightarrow \text{PSQ } k \ p \rightarrow \text{PSQ } k \ p \\
\text{adjust } f \ k \ \emptyset &= \emptyset \\
\text{adjust } f \ k \ \{b\} & \\
\quad | k == \text{key } b &= \{k, f \ (\text{prio } b)\} \\
\quad | \text{otherwise} &= \{b\} \\
\text{adjust } f \ k \ (t_l \wedge t_r) & \\
\quad | k \leq \text{max-key } t_l &= \text{adjust } f \ k \ t_l \wedge t_r \\
\quad | \text{otherwise} &= t_l \wedge \text{adjust } f \ k \ t_r
\end{aligned}$$

The modifier *adjust* does not change the shape of the pennant. By contrast, *insert* possibly increases the height of the tree. Since the loser trees are not balanced, there is the annoying possibility that repeated insertions may produce a degenerated tree.

$$\begin{aligned}
\text{insert} &:: (k, p) \rightarrow \text{PSQ } k \ p \rightarrow \text{PSQ } k \ p \\
\text{insert } b \ \emptyset &= \{b\} \\
\text{insert } b \ \{b'\} & \\
\quad | \text{key } b < \text{key } b' &= \{b\} \wedge \{b'\} \\
\quad | \text{key } b == \text{key } b' &= \{b\} \quad \text{-- update} \\
\quad | \text{key } b > \text{key } b' &= \{b'\} \wedge \{b\} \\
\text{insert } b \ (t_l \wedge t_r) & \\
\quad | \text{key } b \leq \text{max-key } t_l &= \text{insert } b \ t_l \wedge t_r \\
\quad | \text{otherwise} &= t_l \wedge \text{insert } b \ t_r
\end{aligned}$$

In the case of search trees deletion is notoriously more difficult to handle than insertion. Perhaps surprisingly, this does not hold for priority search pennants. The reason is simply that using the tournament view all modifications take place

at the fringe of the tree:

$$\begin{aligned}
\text{delete} &:: k \rightarrow \text{PSQ } k \ p \rightarrow \text{PSQ } k \ p \\
\text{delete } k \ \emptyset &= \emptyset \\
\text{delete } k \ \{b\} & \\
\quad | k == \text{key } b &= \emptyset \\
\quad | \text{otherwise} &= \{b\} \\
\text{delete } k \ (t_l \wedge t_r) & \\
\quad | k \leq \text{max-key } t_l &= \text{delete } k \ t_l \wedge t_r \\
\quad | \text{otherwise} &= t_l \wedge \text{delete } k \ t_r.
\end{aligned}$$

Given the hybrid nature of priority search pennants the definition of *delete* is surprisingly attractive.

5.5 Range queries


Like the second version of the look-up function, the query function *at-most* employs two views simultaneously. The minimum view is used to prune the search if a node is encountered whose priority is greater than the given one.

$$\begin{aligned}
\text{at-most} &:: p \rightarrow \text{PSQ } k \ p \rightarrow [(k, p)] \\
\text{at-most } p_t \ (\text{Min } b \ q) & \\
\quad | \text{prio } b > p_t &= [] \\
\text{at-most } p_t \ \emptyset &= [] \\
\text{at-most } p_t \ \{b\} &= [b] \quad \text{-- we know that } \text{prio } b \leq p_t \\
\text{at-most } p_t \ (t_l \wedge t_r) &= \text{at-most } p_t \ t_l \ ++ \ \text{at-most } p_t \ t_r
\end{aligned}$$

The query function is analysed in Section 7.

6. A BALANCED SCHEME

One of the strengths of priority search pennants as compared to priority search trees is that the basic implementation can be easily extended by a balancing scheme. Most schemes use rotations to restore balancing invariants. Now, while rotations preserve the search-tree property, they do not preserve the semi-heap property as the following example shows.



In the original tree, both losers, *D* and *F*, dominate their right subtree. This implies that they have not played against each other and that the winner stems from the leftmost subtree t_1 . Now, if we rotate the loser tree to the right, the new root should dominate its right subtree but it does not. To restore the semi-heap property we have to exchange *D5* and *F2*. We will see that, in general, at most one exchange at the cost of at most one additional comparison is required. In other words, rotations are constant time operations for priority search pennants.

By contrast, in the case of priority search trees we have to preserve the *heap property*, which takes $\Theta(h)$ time where h is the height of the tree. This means, in particular, that in order to ensure an overall logarithmic time bound, the number of rotations per update must be bounded by a constant. Red-black trees [6] or 2-3-4 trees [10] satisfy this constraint. On the other hand, AVL trees [2] or weight-balanced trees

[1] do not guarantee such a bound. Ironically, Okasaki’s elegant functional implementation of red-black trees [17] also fails to meet this condition.

However, for priority search pennants we can freely choose an underlying balancing scheme. We pick Adams’s weight-balanced trees [1] since they support insertions and deletions equally well. A tree is *weight-balanced* if for all nodes *either* both subtrees have at most one element *or* one subtree does not have more than ω times as many elements as the opposite subtree, where ω is some constant > 3.75 . To check and to maintain the invariant, each node in a loser tree is augmented by a size field:

```

type Size = Int
data Tree k p = Lf
                | Nd Size (k, p) (Tree k p) k (Tree k p).

```

Using views and *smart constructors* we can make the computation of the size field totally transparent.

```

leaf          = Lf
node b l k r  = Nd (1 + size l + size r) b l k r
view Tree k p = Leaf
                | Node (k, p) (Tree k p) k (Tree k p)

where
  Lf          → Leaf
  Nd _ b l k r → Node b l k r
  size        :: Tree k p → Size
  size Lf     = 0
  size (Nd s _ _ _ _) = s

```

In the sequel we will use the smart constructors *leaf* and *node* to construct weight-balanced trees, the view constructors *Leaf* and *Node* to pattern match weight-balanced trees, and the function *size* to query the size field.

The *balance* function defined below maintains weight-balance using single and double rotations under the precondition that at most one subtree has changed size by at most one element and the original tree was in balance. The algorithm is described in more detail in Adams [1].

```

balance b l k r
  | size l + size r < 2 = node b l k r
  | size r > ω * size l = balance-left b l k r
  | size l > ω * size r = balance-right b l k r
  | otherwise          = node b l k r

balance-left b l k r@(Node _ rl _ rr)
  | size rl < size rr = single-left b l k r
  | otherwise        = double-left b l k r

balance-right b l@(Node _ ll _ lr) k r
  | size lr < size ll = single-right b l k r
  | otherwise        = double-right b l k r

```

The *balance* operation is essentially the same as for search trees. Only the implementation of the rotations is more elaborate since they have to maintain the semi-heap property. Figure 6 displays the possible cases for a single rotation to the right. Since a single rotation involves two nodes and since each node may dominate one of two subtrees, we must distinguish four different cases. The only problematic case is the last one, where we have to perform one additional match to determine the top binding. In general, b_1 is the new top binding iff $key\ b_2 > k_1$ and $prio\ b_1 \leq prio\ b_2$. The

four cases for the left rotation are symmetric:

```

single-left b1 t1 k1 (Node b2 t2 k2 t3)
  | key b2 ≤ k2 ∧ prio b1 ≤ prio b2
    = node b1 (node b2 t1 k1 t2) k2 t3
  | otherwise = node b2 (node b1 t1 k1 t2) k2 t3
single-right b1 (Node b2 t1 k1 t2) k2 t3
  | key b2 > k1 ∧ prio b1 ≤ prio b2
    = node b1 t1 k1 (node b2 t2 k2 t3)
  | otherwise = node b2 t1 k1 (node b1 t2 k2 t3).

```

Double rotations are implemented in terms of single rotations.

```

double-left b1 t1 k1 (Node b2 t2 k2 t3)
  = single-left b1 t1 k1 (single-right b2 t2 k2 t3)
double-right b1 (Node b2 t1 k1 t2) k2 t3
  = single-right b1 (single-left b2 t1 k1 t2) k2 t3

```

REMARK 4. Since a double rotation is defined in terms of two single rotations, at most two additional matches are required. Perhaps surprisingly, one can show that only one additional match suffices. A direct implementation of the double rotations is left as an exercise to the reader.

It remains to adapt the implementation of Section 5 to balanced trees. This can be done by a simple renaming: occurrences of the constructors *Start* and *Loser* in patterns must be replaced by *Leaf* and *Node*; occurrences in expressions must be replaced by *leaf* and *balance*. The smart constructor *node* can be used instead of *balance* if the shape of the tree has not changed (as in the case of *adjust*) or if the tree is known to be balanced (as in the case of *from-ord-list*).

Let us conclude the section with a brief discussion of the running times of the various operations. For simplicity, we assume that we are working in a strict setting. Weight-balanced trees have a height that is logarithmic in the number of elements. Consequently, the dictionary operations (*lookup*, *insert*, and *delete*) and the priority queue operations (*Min*) have a worst-case running time of $\Theta(\log n)$. The conversion functions *from-ord-list* and *to-ord-list* are both linear in the number of bindings. Finally, the range query *at-most* takes $\Theta(r(\log n - \log r + 1))$ time where r is the length of the output list—the next section contains a detailed analysis. The following table summarizes the running times:

Constructors and insertion		Destructors and deletion	
\emptyset	$\Theta(1)$	<i>Empty</i>	$\Theta(1)$
$\{\cdot\}$	$\Theta(1)$	<i>Min</i>	$\Theta(\log n)$
<i>insert</i>	$\Theta(\log n)$	<i>delete</i>	$\Theta(\log n)$
<i>from-ord-list</i>	$\Theta(n)$		
Observers		Modifier	
<i>lookup</i>	$\Theta(\log n)$	<i>adjust</i>	$\Theta(\log n)$
<i>to-ord-list</i>	$\Theta(n)$		
<i>at-most</i>	$\Theta(r(\log n - \log r + 1))$		

7. ANALYSIS OF RANGE QUERIES

The range query *at-most* is a so-called *output-sensitive* algorithm, that is, its running time is not only governed by the total number of bindings in the tree but also by the number of bindings it returns as a result. To estimate its running time we have to determine the number of nodes that must be inspected to return r outputs. A general observation is




Figure 6: A single rotation to the right ($-\infty$ represents the winner; $b_1 \leq b_2$ is shorthand for *prio* $b_1 \leq \text{prio } b_2$).

that whenever a player enters the output list, we must additionally check all the players who have lost to this particular player. Consider the pennant of Figure 3. If Lennart is selected, we must check Phil, Erik, and Mary. If Phil is also selected, we must additionally check Simon and Richard.

The structure becomes more apparent if we turn the binary semi-heap into a *multiway heap*. The dominated subtrees become children and the non-dominated subtrees become siblings. Figure 7 displays the tree thus obtained. This transformation is an instance of what is known as the *natural correspondence* between binary trees and forests, see Knuth [12].

To simplify the analysis let us assume that the original trees are perfectly balanced as in our example, so that we have a total number of $n = 2^h$ bindings. In this special case we obtain as the result of the transformation a so-called *binomial heap* [21]. Now, in a binomial heap with $n = 2^h$ elements, we have one node with h subtrees (namely the root), 2^0 nodes with $h - 1$ subtrees, 2^1 nodes with $h - 2$ subtrees, \dots , 2^{h-2} nodes with 1 subtree, and 2^{h-1} nodes with 0 subtrees. Summing up and adding one for the root we obtain a total of n nodes:

$$n = 1 + h + 2^0 \cdot (h - 1) + 2^1 \cdot (h - 2) + \dots + 2^{h-2} \cdot 1 + 2^{h-1} \cdot 0.$$

Using the binary logarithm we can rewrite the above identity into the following form:

$$n = 1 + h + \sum_{k=1}^{n-1} (h - 1 - \lfloor \lg k \rfloor).$$

On the right-hand side we have a sum with $n + 1$ summands. Now, if we only sum up the first $r + 1$ summands, we obtain the desired maximum number of successors of r nodes. Consequently, the worst-case running time of *at-most* is proportional to

$$1 + h + \sum_{k=1}^{r-1} (h - 1 - \lfloor \lg k \rfloor),$$

for $1 < r \leq n$. To estimate the asymptotic growth of this function we use the formula

$$\sum_{k=1}^m \lfloor \lg k \rfloor = (m + 1) \lfloor \lg(m + 1) \rfloor - 2^{\lfloor \lg(m+1) \rfloor + 1} + 2$$

and calculate

$$\begin{aligned} & 1 + h + \sum_{k=1}^{r-1} (h - 1 - \lfloor \lg k \rfloor) \\ &= 1 + h + (r - 1)(h - 1) - (r \lfloor \lg r \rfloor - 2^{\lfloor \lg r \rfloor + 1} + 2) \\ &= r \lg n - r - r \lfloor \lg r \rfloor + 2^{\lfloor \lg r \rfloor + 1} \\ &= r(\lg n - \lg r) + \Theta(r). \end{aligned}$$

Thus, if r is small, we have a logarithmic running time. The running time eventually becomes linear as r approaches n .

Let us conclude the section by noting that priority search pennants answer range queries less efficiently than priority search trees, which support them in $\Theta(\log n + r)$ time [5]. The reason is simply that the heap property is stronger than the semi-heap property: in the case of binary heaps at most two additional elements must be checked for every element that enters the output list. As an aside, this also shows that binomial heaps, which are essentially sequences of semi-heaps [8], are less well-suited for answering range queries.

8. RELATED WORK

Priority search queues. We have already commented on the relationship between priority search pennants and McCreight’s priority search trees [14]. Let us briefly summarize the main points. Priority search trees are restricted to balancing schemes where the number of rotations per update is bounded by a constant. By contrast, our methods works with arbitrary balancing schemes. The asymptotic running times of the finite map and the priority queue operations are the same for both approaches. However, priority search trees support range queries more efficiently.

As an aside, priority search trees should not be confused with *cartesian trees* or *treaps*, which are also a combination of search trees and priority queues [22]. In a priority search tree each node is labelled with two keys, the key of the binding and an additional split key, whereas in a treap the key of the binding serves as the split key, which completely determines the structure of the treap.

Tournament trees and pennants. Tournament trees and loser trees already appear in Knuth’s TAOCP series [13]. The term *pennant* was coined by Sack and Strothotte [20] to denote topped, perfectly balanced trees (we do not require the trees to be perfectly balanced though). Pennants are widespread: Sack and Strothotte employ them to design algorithms for splitting and merging heaps in the form of *left-complete* binary trees, Okasaki [15] uses pennants as a fundamental building block for data structures modelled after number systems, pennants underly *binomial heaps* [8], and they are useful for analysing *red-black trees* [7].

Dijkstra’s algorithm. Using priority search queues we were able to implement Dijkstra’s single-source shortest-paths algorithm in a purely functional way. Previous formulations like that of King [11] relied in an essential way on stateful computations. King writes:

... if a purely function solution exists for these algorithms [Dijkstra’s and Kruskal’s] it will probably involve using a state-encapsulating combinator.

Perhaps surprisingly, by using a different abstract data type—priority search queues instead of priority queues—we obviate the need for state. We feel that the resulting code is much clearer than the state-based formulation.

Views. Views have originally been introduced by Wadler [23]. Later the idea was fleshed out into a proposal for an extension to Haskell [4]. Okasaki slightly simplified the proposal and adapted it to Standard ML [16]. A recent paper by the same author [18], where Okasaki strongly advocates the use of views, revived my interest in this language feature.

9. CONCLUSION

Priority search queues are an amazing combination of finite maps and priority queues in that they support both dictionary and priority queue operations. Building upon the metaphor of a knockout tournament we have developed a simple, yet efficient implementation technique for this abstract data type. In developing the code the concept of views was tremendously helpful: views enhanced both the readability and the modularity of the code. We have presented

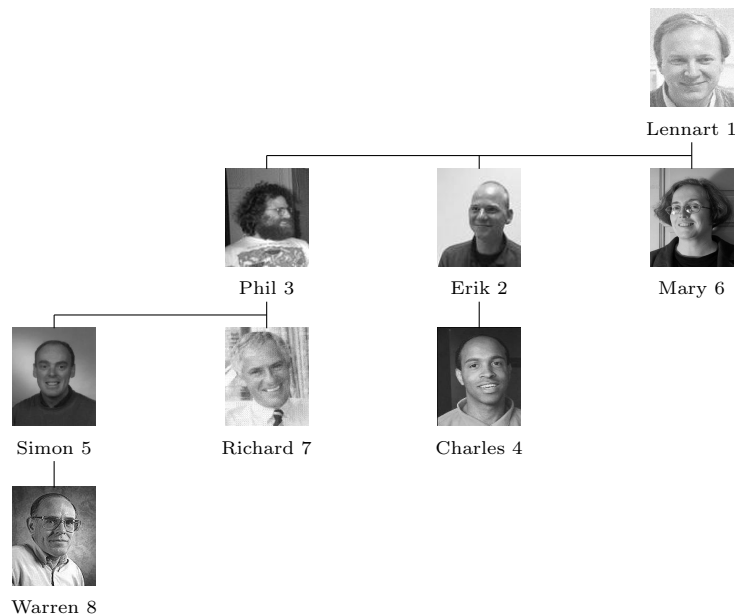


Figure 7: The multi-way heap corresponding to the binary semi-heap of Figure 3.

two applications of priority search queues: a purely functional implementation of Dijkstra’s single-source shortest-paths algorithm and an efficient implementation of the first-fit heuristics for the bin packing problem. We hope to see further applications in the future.

10. REFERENCES

- [1] S. Adams. Functional Pearls: Efficient sets—a balancing act. *J. Functional Programming*, 3(4):553–561, October 1993.
- [2] G. Adel’son-Vel’skiĭ and Y. Landis. An algorithm for the organization of information. *Doklady Akademii Nauk SSSR*, 146:263–266, 1962. English translation in Soviet Math. Dokl. 3, pp. 1259–1263.
- [3] W. Braun and M. Rem. A logarithmic implementation of flexible arrays. Memorandum MR83/4, Eindhoven University of Technology, 1983.
- [4] W. Burton, E. Meijer, P. Sansom, S. Thompson, and P. Wadler. Views: An extension to Haskell pattern matching. Available from <http://www.haskell.org/development/views.html>, 1996.
- [5] O. Fries, K. Mehlhorn, S. Näher, and A. Tsakalidis. A log log n data structure for three-sided range queries. *Information Processing Letters*, 25(4):269–273, June 1987.
- [6] L. J. Guibas and R. Sedgwick. A dichromatic framework for balanced trees. In *Proceedings of the 19th Annual Symposium on Foundations of Computer Science*, pages 8–21. IEEE Computer Society, 1978.
- [7] R. Hinze. Constructing red-black trees. In C. Okasaki, editor, *Proceedings of the Workshop on Algorithmic Aspects of Advanced Programming Languages, WAAAPL’99, Paris, France*, pages 89–99, September 1999. The proceedings appeared as a technical report of Columbia University, CUCS-023-99, also available from <http://www.cs.columbia.edu/~cdo/waaapl.html>.
- [8] R. Hinze. Functional Pearl: Explaining binomial heaps. *J. Functional Programming*, 9(1):93–104, January 1999.
- [9] R. Hinze. A simple implementation technique for priority search queues. Technical report, UU-CS-2001-09, Universiteit Utrecht, March 2001.
- [10] S. Huddleston and K. Mehlhorn. A new data structure for representing sorted lists. *Acta Informatica*, 17:157–184, 1982.
- [11] D. King. *Functional Programming and Graph Algorithms*. Ph.d. thesis, Department of Computer Science, University of Glasgow, March 1996.
- [12] D. E. Knuth. *The Art of Computer Programming, Volume 1: Fundamental Algorithms*. Addison-Wesley Publishing Company, 3rd edition, 1997.
- [13] D. E. Knuth. *The Art of Computer Programming, Volume 3: Sorting and Searching*. Addison-Wesley Publishing Company, 2nd edition, 1998.
- [14] E. M. McCreight. Priority search trees. *SIAM Journal on Computing*, 14(2):257–276, May 1985.
- [15] C. Okasaki. *Purely Functional Data Structures*. Cambridge University Press, 1998.
- [16] C. Okasaki. Views for Standard ML. In *The 1998 ACM SIGPLAN Workshop on ML, Baltimore, Maryland*, pages 14–23, September 1998.
- [17] C. Okasaki. Functional Pearl: Red-Black trees in a functional setting. *J. Functional Programming*, 9(4):471–477, July 1999.
- [18] C. Okasaki. Breadth-first numbering: lessons from a small exercise in algorithm design. *ACM SIGPLAN Notices*, 35(9):131–136, September 2000.
- [19] S. Peyton Jones and J. Hughes, editors. *Haskell 98 — A Non-strict, Purely Functional Language*, February 1999. Available from <http://www.haskell.org/definition/>.
- [20] J.-R. Sack and T. Strothotte. A characterization of heaps and its applications. *Information and Computation*, 86(1):69–86, May 1990.
- [21] J. Vuillemin. A data structure for manipulating priority queues. *Communications of the ACM*, 21(4):309–315, 1978.
- [22] J. Vuillemin. A unifying look at data structures. *Communications of the ACM*, 23:229–239, 1980.
- [23] P. Wadler. Views: a way for pattern matching to cohabit with data abstraction. In *Fourteenth Annual ACM SIGACT-SIGPLAN Symposium on Principles of programming languages*, 1987, pages 307–313. ACM Press.