New application of failure functions

Authors
Hirschberg, D. S.
Larmore, L. L.

Publication Date
1984

Peer reviewed

New Applications of Failure Functions

D. S. Hirschberg and *L. L. Larmore*

October, 1984

Technical Report 236

Abstract. Several algorithms are presented whose operations are governed by a principle of failure functions: when searching for an extremal value within a sequence, it suffices to consider only the subsequence of items each of which is the first possible improvement of its predecessor. These algorithms are more efficient than their more traditional counterparts.

This research was supported in part by National Science Foundation Grant MCS-82-00362. Authors' address: Department of Information and Computer Science, University of California, Irvine, CA 92717.

New Applications of Failure Functions

D. S. Hirschberg and *L. L. Larmore*

1. Introduction

The notion of failure functions is often associated with the linear-time substring recognition algorithm [1,6]. The principle of failure functions is disarmingly simple: when searching for an extremal value within a sequence, it suffices to consider only the subsequence of items each of which is the first feasible alternative of its predecessor. The value of the failure function is a pointer to that first feasible alternative. Implementing this function in isolation will not yield any advantage since the effort required to determine what is the first feasible alternative is equal to the hoped for savings, which is not having to consider many losing alternatives. The preprocessing costs negate the run-time savings. (In practice, the preprocessing may be chronologically interspersed with the processing.) However, if many such searches are contemplated and they are closely related then the preprocessing costs may be spread over the multiple searches with some additional intersearch fix-up costs. The net effect may be some real savings. This was the case for the pattern matching algorithm, and is also the case for the algorithms given in this paper.

We first consider the problem of determining the optimum way to break a paragraph (scroll of words) into lines, provided the penalty function (for a line being too long or too short) is linear. Our algorithm for this problem is linear-time. We then exhibit an algorithm for more general penalty functions which is linear-time in the case of a piecewise quadratic function.

We also consider the problem of finding the minimum sum of key length pagination of a scroll of n items. We present a linear-time algorithm, improving on the $O(n \log n)$ result of Diehr and Faaland.

2. Breaking a Paragraph into Lines

We are given a paragraph consisting of a scroll of n words, where the i^{th} word has length $w_i > 0$, and a non-negative valued function *penalty*(x) which is defined over the closed interval $[lmin,lmax]$, where $0 < lmin < lmax$. We assume that there is an optimum line length $lopt \in [lmin,lmax]$ for which $penalty(lopt) = 0$. We define a ***break sequence*** of the paragraph to be a monotone increasing sequence $1 = b_1, b_2, \dots b_m \le n$ of integers. The break sequence defines lines, where the k^{th} line begins with the $b_k{}^{th}$ word, and its length is $length_k = w_{b_k} + \dots + w_{b_{k+1}-1}$ (where b_{m+1} is taken to be $n+1$). We say that a break sequence is ***admissable*** if $length_k \in [lmin,lmax]$ for all $k < m$, and $length_m \le lmax$.

The *total penalty* of a given break sequence is defined to be the sum of the penalties of the lines, but where the last line is not penalized for being too short. The problem is to find an admissable break sequence with minimum total penalty.

2.1 The Traditional Line-Breaking Algorithm

For each $i \le n+1$, define $f[i]$ to be the lowest total penalty of any break sequence of the subscroll $w_i \dots w_n$. We let $f[n+1] = 0$ by default. The traditional algorithm (see, for example, [3]) uses dynamic programming.

For any $1 \le i \le j \le n+1$, let $Line(i,j) = w_i + \dots + w_{j-1}$, and let $Legal(i,j)$ be the boolean function which is true if and only if $Line(i,j) \in [lmin,lmax]$.

Algorithm 1 — Traditional Algorithm

```
        f[n+1] ← 0
Loop:   for i from n downto 1 do
           if Line(i,n+1)≤lopt then
              begin
                 f[i] ← 0
                 nextbreak[i] ← n+1
              end
           else if Legal(i,j) for some j then
              begin
Choose:          Choose r such that Legal(i,r) and f[r] + penalty(Line(i,r)) is minimized
```

```
            f[i] ← f[r] + penalty(Line(i,r))
            nextbreak[i] ← r
        end
    else
        f[i] ← ∞
if f[1] < ∞ then Define_break_sequence
```

The subroutine Define_break_sequence recovers the breakpoint vector b from the array *nextbreak*.

Subroutine Define_break_sequence

```
s ← 1
t ← 1
while t ≤ n do
    begin
        b[s] ← t
        s ← s+1
        t ← nextbreak[t]
    end
```

The bottleneck in the Traditional Algorithm is the Choose step, since all other steps can be done in time $O(n)$. The total time for all executions of the Choose step is $O(nM)$, where M is the largest number of words that could possibly occur in a line (we could set $M = lmax/W$, where W is the minimum value of w_i). If M is considered to be bounded, then the Traditional Algorithm is linear. However, if we consider a class of problems in which M, as well as n, grows then the Traditional Algorithm is no longer linear.

2.2 Linear Penalty Function

We consider the case that *penalty* is linear, i.e., for all $x \in [lmin, lmax]$, $penalty(x) = C(x - lopt)$, for some constant C which may be positive, negative, or zero. We define $penalty(x) = \infty$ if $x \notin [lmin, lmax]$.

We use dynamic arrays *leftlow* and *rightlow*, which have pointer (actually index) values, and dynamic arrays f and g, which have penalty values. The array f is identical to the f in the traditional algorithm, and g is a modified array which always satisfies the equation $g[k] = f[k] + CLine(1,k)$. At any given time, *rightlow*$[k]$ is the smallest $l > k$ such that $g[l] \leq g[k]$, and *leftlow*$[k]$ is the largest $l < k$ such that $g[l] < g[k]$. *leftlow* is used as a failure function for choosing the previous breakpoint (beginning of a line) corresponding to a current end of line, and *rightlow* is used as a failure function for updating the *leftlow* values.

Algorithm 2 — Linear Penalty Algorithm

```
        g[n+2] ← -∞
        f[n+1] ← 0
        g[n+1] ← CLine(1,n+1)
        r ← n+1
        rightlow[n+1] ← n+2
Loop:   for i from n downto 1 do
           begin
              if Line(i,n+1)≤lopt then
                 begin
                    f[i] ← 0
                    nextbreak[i] ← n+1
                 end
              else
                 begin
Choose1:            while Line(i,r) > lmax do    r ← r-1
Choose2:            while leftlow[r] defined and Legal(i,leftlow[r]) do    r ← leftlow[r]
                    f[i] ← f[r] + penalty(Line(i,r))
                    if f[i] < ∞ then nextbreak[i] ← r
                 end
              g[i] ← f[i] + CLine(1,i)
              k ← i+1
Update:       while g[k] > g[i] do
                 begin
                    leftlow[k] ← i
                    k ← rightlow[k]
                 end
              rightlow[i] ← k
           end (of Loop)
        if f[1] < ∞ then Define_break_sequence
```

We can prove the correctness of the Linear Penalty Algorithm by showing that it simulates the Traditional Algorithm. We need three lemmas.

Lemma A. Fix i. Among the set of all j such that $Legal(i,j)$, that j which minimizes $g[j]$ also minimizes $f[j] + penalty(Line(i,j))$.

Proof. $f[j] + penalty(Line(i,j)) - g[j] = penalty(Line(i,j)) - CLine(1,j)$ $= -C(lopt + Line(1,i))$, which is constant for fixed i. □

Lemma B. The following loop invariant holds after each iteration of the main loop of the Linear Penalty Algorithm.

LB1: For all $i \le s \le n+1$, $rightlow[s] = t$, where $t > s$ is the smallest value such that $g[t] \le g[s]$.

LB2: For all $i < s \le n+1$, $leftlow[s] = t$, where $i \le t < s$ is the largest value such that $g[t] < g[s]$, provided such a t exists. Thus, in this case, for all $leftlow[s] < j < s$, $g[j] \ge g[s]$. Otherwise, $leftlow[s]$ is undefined.

Proof. By induction on i. Initially, i.e., before the main loop iterates at all, we can take $i = n+1$. Then, LB2 holds vacuously, while LB1 holds by initial assignment.

Our inductive hypothesis is that LB holds for all values larger then i. Thus, before execution of the Update loop, for all $i+1 \le s \le n+1$:

(a) $rightlow[s] = t$, where $t > s$ is the smallest value such that $g[t] \le g[s]$.

(b) $leftlow[s] = t$, where $i+1 \le t < s$ is the largest value such that $g[t] < g[s]$, provided such a t exists. Otherwise, $leftlow[s]$ is undefined.

Thus, we need to prove only that, at the end of an iteration of the main loop, $rightlow[i]$ has the correct value, and that $leftlow[s] = i$ if $i < s \le n+1$ and $g[i] < g[s]$ and $leftlow[s]$ was undefined before the Update loop.

Let k_0, ... k_m be the sequence of values of k produced in the Update loop, i.e., $k_0 = i+1$, and $k_{l+1} = rightlow[k_l]$ for $0 \le l < m$. Note that $g[k_m] \le g[i]$, and $g[k_l] >$

$g[_i]$ for all $0 \leq l < m$.

Sublemma. For all $0 \leq l \leq m$, and for all $i < s \leq k_l$, $g[s] \geq g[k_l]$. Furthermore, if $l \geq 1$ and $s < k_l$, $g[s] \geq g[k_{l-1}]$.

Proof. By induction on l. For $l = 0$, the sublemma holds since we must have $s = k_0 = i+1$. For the inductive step, assume the sublemma holds for $l-1$. If $i < s \leq k_l$ then, because $k_l = rightlow[k_{l-1}] > k_{l-1}$ by LB1, either (i) $s \leq k_{l-1}$, or (ii) $k_{l-1} < s < k_l$, or (iii) $s = k_l$.

If (i) then, by the inductive hypothesis, $g[s] \geq g[k_{l-1}]$.

If (ii) then $k_l = rightlow[k_{l-1}]$ by the second assignment of the Update loop. By LB1, $g[k_l] \leq g[k_{l-1}]$ and for any s between k_{l-1} and k_l, $g[s] > g[k_{l-1}]$.

In either case we are done, since $g[k_l] \leq g[k_{l-1}]$. Case (iii) is trivial. □

Proof of Lemma B, continued. By the sublemma, all $i < s \leq k_l$ are unsuitable values for $leftlow[k_l]$ since $g[s] \geq g[k_l]$. Therefore, for any $l < m$, $leftlow[k_l]$ should be assigned the value i if $g[i] < g[k_l]$. The first assignment of the Update loop does exactly that. We now show that $rightlow[i]$ should be assigned the value k_m. That is, we need to show that, for $i < s < k_m$, $g[s] > g[i]$ and that $g[k_m] \leq g[i]$.

Since the Update loop no longer iterates when $k = k_m$, $g[k_m] \leq g[i]$. We now show that $g[s] > g[i]$ for all $i < s < k_m$. If $m = 0$, this is vacuously true since $k_0 = i+1$. Otherwise, $g[s] \geq g[k_{m-1}]$ by the sublemma, and $g[k_{m-1}] > g[i]$ because the Update loop continues to iterate when $k = k_{m-1}$.

Thus, the Update loop makes the correct assignment to $rightlow[i]$. □

Lemma C. After execution of the Choose loops, one of the following two conditions holds.

(I) There is no $j \leq n+1$ such that $Legal(i,j)$, and r is the largest possible value of $j \leq n+1$ such that $Line(i,j) \leq lmax$.

or (II) Among all j such that $Legal(i,j)$, r is the choice of j which minimizes $g[j]$.

Proof. By induction on i. Condition (I) holds before the first iteration of the main loop (consider $i = n+1$). We define a loop invariant.

LC1: For all $r < j \leq n+1$ such that $Line(i,j) \leq lmax$, $g[j] \geq g[r]$.

LC2: If $r < n+1$, $Line(i,r+1) \geq lmin$.

We will establish that the loop invariant holds initially i.e., before the main loop (consider $i = n+1$), and that it is preserved by each execution of the while loops of the choose block, as well as when i is decremented in the main loop.

Initially, LC holds vacuously. Decrementing i cannot cause LC to fail, because $Line(i,j)$ is monotone decreasing on the first argument. An iteration of Choose1 preserves LC2 because, immediately after any such iteration, $Line(i,r+1) > lmax \geq lmin$. Also, if Choose1 iterated one or more times, LC1 holds vacuously.

We now show that an iteration of Choose2 preserves LC. r' denotes the value of r after the iteration.

If LC holds before an iteration then, by LC1, $g[j] \geq g[r]$ for all $r < j \leq n+1$ such that $Line(i,j) \leq lmax$. In order for Choose2 to iterate, $leftlow[r]$ is defined and $Line(i,leftlow[r]) \geq lmin$.

LC2 will hold after the iteration, since $Line(i,leftlow[r]+1) > Line(i,leftlow[r]) \geq lmin$ and thus $Line(i,r'+1) \geq lmin$.

To prove that LC1 is preserved, consider j such that $r' = leftlow[r] < j \leq n+1$ and $Line(i,j) \leq lmax$. We need to show that $g[j] \geq g[leftlow[r]] = g[r']$. If $r < j \leq n+1$ then $g[j] \geq g[r]$ by LC1 and $g[r] > g[r']$ by LB2. If $r' = leftlow[r] < j \leq r$ then

$g[j] \geq g[r] > g[leftlow[r]]$ by LB2.

Thus, LC is loop invariant.

Consider the case when $Legal(i,r)$ after Choose1 has executed. By LC1, $g[j] \geq g[r]$ for any $j > r$ such that $Legal(i,j)$. If $j < r$ and $Legal(i,j)$ then, since Choose2 has ceased iterating, either $leftlow[r]$ is undefined, or $Line(i,leftlow[r]) < lmin$ which implies by monotonicity of *Line* that $leftlow[r] < j < r$. In either case, by LB2, $g[j] \geq g[r]$. Thus case (II) of Lemma C holds.

Otherwise, not $Legal(i,r)$ after Choose1 has executed. In this case, $Line(i,r) \leq lmax$ since Choose1 did not iterate again, and therefore $Line(i,r) < lmin$ since otherwise $Legal(i,r)$. Thus, Choose2 cannot iterate by the monotonicity of *Line* and the fact that if $leftlow[r]$ is defined then $leftlow[r] < r$. Now, either Choose1 iterated once or more, or it did not iterate. If Choose1 iterated, Case (I) of Lemma C holds, since then $Line(i,r+1) > lmax$ (and Choose2 did not iterate). If Choose1 did not iterate, r has the same value as it did after execution of Choose2 of the previous iteration of the main loop. By the inductive hypothesis, Lemma C held for $i+1$. Case (II) cannot have held, because then $Line(i,r) > Line(i+1,r) \geq lmin$, which would imply that $Legal(i,r)$. Thus Case (I) of Lemma C held for $i+1$. Since *Line* is monotone decreasing on its first parameter, Case (I) of Lemma C then holds for i. $\square$

We now prove the correctness of the Linear Penalty Algorithm, assuming that the traditional algorithm is correct. We need to show that the Choose steps of the Linear Penalty Algorithm are equivalent to the Choose step of the traditional algorithm. By Lemmas A and C (only case C(II) applies since we are within the *Else if* clause), the value of r after execution of Choose2 is the same as the value of r chosen in the Choose step of the traditional algorithm, provided such a legal r exists. $\square$

Time complexity. We can use $O(n)$ preprocessing time to compute $Line(1,i)$ for all i. Then $Line(i,j)$ (and hence $Legal(i,j)$) can be computed in $O(1)$ time by the formula $Line(i,j) = Line(1,j) - Line(1,i)$.

The Choose loops appear to iterate $O(n)$ times within each iteration of the main loop. However, r decreases with each iteration of each of those loops. Thus, the total number of such iterations cannot exceed n.

The Update loop also appears to iterate $O(n)$ times within each iteration of the main loop. Note, by LB2, the value of $leftlow[k]$, once defined, is never redefined. It follows that the total number of iterations of the Update loop, over all iterations of the main loop, is at most n.

2.3 General Concave Penalty Function

We say that a function $p(x)$ is *concave* if, for any $x < y < z$ in its domain, $(z-x)p(y) \leq (y-x)p(z) + (z-y)p(x)$. For example, any quadratic function with non-negative leading coefficient is concave.

We now consider the breaksequence problem where $penalty(x)$ is non-negative and concave for $x \in [lmin,lmax]$. As before, $penalty(x) = \infty$ for $x \notin [lmin,lmax]$, and there is no penalty for the last line if its length does not exceed $lopt$.

The time bottleneck in the General Concave Algorithm (GCA) given below is the evaluation of the Boolean function *Bridge*. All other parts of the algorithm run in linear time, and *Bridge* needs to be evaluated $O(n)$ times. If *penalty* is quadratic, *Bridge* can be evaluated in $O(1)$ time, and hence the entire algorithm is linear. Generally, *Bridge* can be evaluated in $O(\log M)$ time by binary search, making the entire algorithm $O(n \log M)$. We leave open the possibility that a faster general algorithm exists.

Notation. For convenience, we let $F(i,j) = f[j]+penalty(Line(i,j))$, the least cost of a paragraph beginning at the i^{th} word whose *second* line begins at the j^{th} word.

The Boolean function Bridge. $Bridge(j,k,l)$ is defined for $1 < j < k < l \leq n+1$. If *true*, it means that k need not be considered as a choice for $nextbreak[i]$ for any "future" i (i.e., $i < j$), since either j or l is always (i.e., for any i) at least as good a choice as k. Formally, for the algorithm to run correctly, it suffices that *Bridge* satisfy the following two conditions.

Br1: If $1 \leq i < j < k < l \leq n+1$ such that $Legal(i,k)$ and $Bridge(j,k,l)$, then $F(i,j) \leq F(i,k)$ or $F(i,l) \leq F(i,k)$.

Br2: If $1 \leq i < j < k < l \leq n+1$ such that $Legal(i,k)$ and not $Bridge(j,k,l)$, then $F(i,k) \leq F(i,j)$ or $F(i,k) \leq F(i,l)$.

There is an allowed ambiguity in the definition of *Bridge*. Any function that satisfies Br1 and Br2 will work. We note that, for example, one possible *Bridge* function is such that it is *false* if and only if there exists some i such that $F(i,k)$ is less than both $F(i,j)$ and $F(i,l)$. To compute this particular function, we can determine whether such an i exists by binary search since, by concavity of *penalty*, $F(i,j) \leq F(i,k)$ implies that i is too low and $F(i,l) \leq F(i,k)$ implies that i is too high. Because we can restrict our initial search domain to no more than M possible values of i, *Bridge* can be computed in $O(\log M)$ time.

Quadratic Case. Suppose that $penalty(x) = ax^2 + bx + c$ for $x \in [lmin,lmax]$, where $a \geq 0$. Then for any $j < k < l$, let $Bridge(j,k,l)$ be *true* if and only if the following two conditions hold.

Q1: $f[k] + penalty(lmax-Line(k,l)) \geq f[j] + penalty(lmax-Line(j,l))$

Q2: $Line(j,l)f[k] \geq Line(j,k)f[l] + Line(k,l)f[j] + aLine(j,k)Line(j,l)Line(k,l)$

Q1 and Q2 can both be computed in $O(1)$ time. Thus, we have a linear time algorithm for the case of a quadratic penalty function.

Data structure. We make use of an input-restricted deque S of integers. Integers can be deleted from both the top and bottom ends of S, but can only be inserted to the top end. Deque S is used to choose r, similar in function to the *leftlow* pointer array in the Linear Penalty Algorithm (LPA). The chosen value of r will be at the bottom of S.

Let us define time i to be the point in an algorithm when the main loop variable has value i (smaller values of i are later).

After it is completely evaluated, the *leftlow* pointer array is a failure forest that can be thought of as being rooted at 0. At any time in the LPA, the *leftlow* failure tree is only partially constructed. During each loop, the LPA progressively develops the failure tree (in Update) and eliminates from consideration some candidates by consideration of *lmax* (in Choose1) and by following a chain in the failure tree (in Choose2). In the GCA, deque S at time i corresponds to the frontier of the developing failure tree in the LPA at the latest time j when $Line(i,j) \geq lmin$. The following operators on S are used.

Functions:	$\lvert S \rvert$	= current cardinality of S
	Top	= value of the top element of S
	Bottom	= value of the bottom element of S
	2Top	= value of the second from the top element of S
	2Bottom	= value of the second from the bottom element of S

Procedures:	*Pop*	delete the top element of S
	Drop	delete the bottom element of S
	Push(x)	insert x at the top of S

Algorithm 3 — General Concave Algorithm

```
          f[n+1] ← 0
          S ← Λ (empty list)
          eol ← n+1
Loop:     for i from n downto 1 do
            begin
Choose1:      while S nonempty and Line(i,Bottom) > lmax do   Drop
Update:       while Line(i,eol) ≥ lmin do
                begin
                  while S nonempty and F(i,eol) ≤ F(i,Top) do   Pop
                  while |S| ≥ 2 and Bridge(eol,Top,2Top) do   Pop
                  if Line(i,eol) ≤ lmax then Push(eol)
                  eol ← eol-1
```

```
            end (of Update)
Choose2:    while |S| ≥ 2 and F(i,2Bottom) ≤ F(i,Bottom) do   Drop
            if Line(i,n+1) ≤ lopt then
                begin
                    nextbreak[i] ← n+1
                    f[i] ← 0
                end
            else if S nonempty then
                begin
                    nextbreak[i] ← Bottom
                    f[i] ← F(i,Bottom)
                end
            else (i.e., S = Λ)
                f[i] ← ∞
        end (of Loop)
    if f[1] < ∞ then Define_break_sequence
```

2.4 Piecewise Concave Penalty Function

If the penalty function is piecewise concave, the algorithm can be generalized, using one deque for each concave piece. The running times are simply added. If there are Γ concave pieces, the running time for the combined algorithm is $O(\ n\Gamma(1+\log(M/\Gamma))\ )$. In the case where the function is piecewise linear or piecewise quadratic, the running time is $O(n\Gamma)$.

The method is essentially to use independent copies of the general concave algorithm, one for each concave piece of the penalty function. These procedures meet once during each iteration of the main loop to exchange information and decide which one has the best value for *nextbreak*[*i*].

3. Pagination of Scrolls

A *boundary sequence* for a scroll is a sequence $0 = s_0 < s_1 < \ldots < s_{v+1} = n+1$ such that $\Sigma_{s_{k-1}<i<s_k} w_i \in [lmin, lmax]$ for all $1 \le k \le v+1$, where $0 \le lmin < lmax$ are fixed. The *length* of that boundary sequence is defined to be $\Sigma_{1\le k\le v} w_{s_k}$. McCreight [5] asks whether we can "quickly" find a boundary sequence of minimum length.

Diehr and Faaland [2] develop an algorithm which finds the minimum length boundary sequence in $O(n \lg n)$ time. We present a linear-time algorithm.

For convenience, assign any positive value, say 1, to w_{n+1} and w_0.

Define $Gap(a,b)$ as the sum of the lengths of the scroll items, w_i, ***strictly between*** the a^{th} and the b^{th} items. Note that $Gap(a,a+1) = 0$. Define $Gap(a,a) = -w_a$.

Define boolean function $Page(a,b)$ to be ***true*** iff $Gap(a,b) \in [lmin, lmax]$.

For any $0 \leq a \leq b \leq n+1$, we define an ***admissable path*** from a to b to be a sequence $s_0, s_1, \ldots s_v$ such that $Page(s_{k-1}, s_k)$ for each $0 < k \leq v$. The ***length*** of that path is $\Sigma_{1 \leq k \leq v} w_{s_k}$. If there exists an admissable path from j to $n+1$, we say that j is *accessable.*

For any $0 \leq i \leq n+1$, define $f(i)$ to be the minimum length of all paths from i to $n+1$. If i is inaccessable, let $f(i) = \infty$.

For each $0 \leq i < n+1$ such that $Page(i,k)$ for some k, define $\rho(i)$ to be the unique number which satisfies the following three conditions:

(i) $Page(i, \rho(i))$

(ii) $f(\rho(i))$ is minimized subject to (i)

(iii) $\rho(i)$ is maximized subject to (i) and (ii)

If there is no k for which $Page(i,k)$ is ***true***, then $\rho(i)$ is undefined. Also, $\rho(n+1)$ is undefined.

Computation of f and ρ clearly suffices to find the minimum length boundary sequence. A boundary sequence exists if and only if $f(0) < \infty$, and the minimum length boundary sequence can be found by using ρ.

Algorithm 4 — Scroll Pagination

```
          Compute Sum[i] = Σ_{k≤i} w[k], 0 ≤ i ≤ n+1
          leftlow[i] ← -1, 0 ≤ i ≤ n+1
          f[n+1] ← 0
          r ← n+1
Loop:     for i from n downto 0 do
              begin
Choose1:         while Gap(i,r) > lmax do   r ← r-1
Choose2:         while r < leftlow[r] and Gap(i,leftlow[r]) ≥ lmin do   r ← leftlow[r]
                 if Page(i,r) then
                     begin
                         f[i] ← f[r] + w[i]
                         ρ[i] ← r
                     end
                 else
                     f[i] ← ∞
                 k ← i+1
Update:          while f[k] > f[i] do
                     begin
                         leftlow[k] ← i
                         k ← rightlow[j]
                     end
                 rightlow[i] ← k
              end (of Loop)
```

It is important to distinguish between the functions $f(i)$ and $\rho(i)$ on the one hand, which are defined abstractly, and the arrays $f[i]$ and $\rho[i]$, whose values are assigned dynamically during execution of the algorithm. Also, we remind the reader that, for all $0 < i \leq n+1$, either $f(i) = \infty$ or $f(i) = f(\rho(i)) + w_i$.

Intuitively, the algorithm works as follows. r is a running temporary $\rho(i)$, which never decreases. When r is too large because $Gap(i,r) > lmax$, r is decremented by 1 until Gap is small enough. We then need to decrease r, minimizing the f value, thus obtaining $\rho(i)$. In [2], a heap of possible values is maintained, and it takes $\Theta(\lg n)$ time to find $\rho(i)$. In Algorithm 4, the pointer *leftlow* tells us where to look next. Even though it might take $\Theta(n)$ time to find $\rho(i)$ for a particular i, the total time for these searches over all i is still only $O(n)$, since r never increases. Thus, *leftlow* is a failure function. The pointer array *rightlow* is used for updating *leftlow*, and also for updating itself. It too is used as a failure function.

Loop invariant. For any $0 \leq i \leq n+1$, the following conditions hold after $n+1-i$ iterations of the loop of Main:

L1(i): If $\rho(i)$ is defined, $r = \rho(i)$. Otherwise, r is the largest j such that $Gap(i,j) \leq lmax$.

L2(i): For all $i \leq j \leq n+1$, $f[j] = f(j)$.

L3(i): For all $i \leq j \leq n+1$, if $\rho(j)$ is defined, $\rho[j] = \rho(j)$. Otherwise, $\rho[j]$ is undefined.

L4(i): For all $i \leq j \leq n+1$, $leftlow[j]$ is the largest $i \leq k < j$ such that $f(k) < f(j)$, provided there is such a k. Otherwise, $leftlow[j] = -1$.

L5(i): For all $i \leq j < n+1$, $rightlow[j]$ is the smallest $j < k \leq n+1$ such that $f(k) \leq f(j)$.

The reader is referred to [4] for a complete proof of the loop invariants.

References

[1] Aho, A.V., J.E. Hopcroft, and J.D. Ullman. *The Design and Analysis of Computer Algorithms.* Addison-Wesley, 1974, pp.329-335.

[2] Diehr, G. and B. Faaland. Optimal pagination of B-trees with variable-length items. *Comm. ACM 27,* 3 (March 1984), 241-247.

[3] Knuth, D.E. and M.F. Plass. Breaking paragraphs into lines. *Software – Practice & Experience* (1981), 1119-1184.

[4] Larmore, L.L. and D.S. Hirschberg. Efficient optimal pagination of scrolls. Technical Report 224, ICS Dept., Univ. of Calif. Irvine, April 1984.

[5] McCreight, E.M. Pagination of B*-trees with variable-length records. *Comm. ACM 20,* 9 (Sept. 1977), 670-674.

[6] Morris, J.H., Jr. and V.R. Pratt. A linear pattern matching algorithm. Technical Report No. 40, Computing Center, Univ. of Calif. Berkeley, 1970.