# FUNCTIONAL PEARL <br> Inverting the Burrows-Wheeler transform 

RICHARD S. BIRD and SHIN-CHENG MU<br>Programming Research Group, Oxford University, Wolfson Building, Parks Road, Oxford OX1 3QD, UK<br>(e-mail: richard.bird@comlab.ox.ac.uk)


#### Abstract

The Burrows-Wheeler Transform is a string-to-string transform which, when used as a preprocessing phase in compression, significantly enhances the compression rate. However, it often puzzles people how the inverse transform is carried out. In this pearl we to exploit simple equational reasoning to derive the inverse of the Burrows-Wheeler transform from its specification. We also outline how to derive the inverse of two more general versions of the transform, one proposed by Schindler and the other by Chapin and Tate.


## 1 Introduction

The Burrows-Wheeler Transform (Burrows \& Wheeler, 1994) is a method for permuting a string with the aim of bringing repeated characters together. As a consequence, the permuted string can be compressed effectively using simple techniques such as move-to-front or run-length encoding. In Nelson (1996), the article that brought the BWT to the world's attention, it was shown that the resulting compression algorithm could outperform many commercial programs available at the time. The BWT has now been integrated into the high-performance utility bzip2, available from Seward (2000).

Clearly, the best way of bringing repeated characters together is just to sort the string. But this idea has a fatal flaw as a preliminary to compression: there is no way to recover the original string unless the complete sorting permutation is produced as part of the output. Instead, the BWT achieves a more modest permutation, one that aims to bring some but not all repeated characters into adjacent positions. Moreover, the transform can be inverted using a single additional piece of information, namely an integer $b$ in the range $0 \leqslant b<n$, where $n$ is the length of the output (or input) string.

It often puzzles people, at least on a first encounter, as to why the BWT is invertible and how the inversion is actually carried out. We identify the fundamental reason why inversion is possible and use it to derive the inverse transform from its specification. As a bonus, we can further derive the inverse of two variations of the BWT transform, one proposed in Schindler (1997), another in Chapin \& Tate (1998).

## 2 Defining the BWT

The BWT is specified by two functions: bwp :: String $\rightarrow$ String, which permutes the string and $b w n::$ String $\rightarrow$ Int, which computes the supplementary integer. The restriction to strings is not essential to the transform, and we can take $b w p$ to have the Haskell type Ord $a \Rightarrow[a] \rightarrow[a]$, so lists of any type will do provided there is a total ordering relation on the elements. The function $b w p$ is defined by

$$
\begin{equation*}
b w p=\text { map last } \cdot \text { lexsort } \cdot \text { rots } \tag{1}
\end{equation*}
$$

The function lexsort :: Ord $a \Rightarrow[[a]] \rightarrow[[a]]$ sorts a list of lists into lexicographic order and is considered in greater detail in the following section. The function rots returns the rotations of a list and is defined by

$$
\begin{array}{lll}
\text { rots } & :: & {[a] \rightarrow[[a]]} \\
\text { rots } x s & = & \text { take }(\text { length } x s)(\text { iterate lrot } x s) \\
\text { lrot } & :: & {[a] \rightarrow[a]} \\
\text { lrot } x s & = & \text { tail } x s+[\text { head } x s]
\end{array}
$$

The function lrot performs a single left rotation. In words, (1) reads: take the last column in the lexicographically sorted matrix of rotations of the input. The definition of $b w p$ is constructive, but we won't go into details - at least, not in this pearl - as to how the program can be made more efficient.

The function $b w n$ is specified by
lexsort (rots xs) !!bwn xs =xs
where $y s!!k$ returns the element of $y s$ in position $k$. In words, bwn xs returns some position at which $x s$ occurs in the sorted list of rotations of $x s$. If $x s$ is a repeated string, then rots $x s$ will contain duplicates, so $b w n x s$ is not defined uniquely by (2).

As an illustration, consider the string yokohama. The rotations and the lexicographically sorted rotations are as follows:

| ○ $\mathrm{k} \circ \mathrm{h} \mathrm{a} \mathrm{ma}$ | 0 a m |
| :---: | :---: |
| o k o h a m a y | 1 a y ○ $\mathrm{k} \circ \mathrm{ha}$ |
| $\mathrm{k} \circ \mathrm{hamayo}$ | 2 hamayok |
| ○ h a m a y ok | $3 \mathrm{k} \circ \mathrm{hama}$ |
| hamayoko | 4 mayok 0 h |
| a mayokoh | 5 |
| mayokoh a | 6 ○ $\mathrm{k} \circ \mathrm{hamay}$ |
| y $0 \mathrm{k} \circ \mathrm{ham}$ | 7 y ○ $\mathrm{k} \circ \mathrm{hama}$ |

The output of $b w p$ is the string hmooakya, the last column of the second matrix, and bwn "yokohama" $=7$ because row number 7 in the sorted matrix of rotations is the input string.

The BWT helps compression because it brings together characters with a common context. To give a brief illustration, an English text may contain many occurrences of words such as "this", "the", "that" and some occurrences of "where", "when", "she", "he" (with a space), etc. Consequently, many of the rotations beginning with
" $h$ " will end with a "t", some with a " $w$ ", an " $s$ " or a space. The chance is smaller that a rotation beginning with " $h$ " would end in a " $x$ ", a " $q$ ", or an " $u$ ", etc. Thus the BWT brings together a subset of the letters, say, those " $t$ "s, " $w$ "s and " $s$ "s. A move-to-front encoding phase is then able to convert the characters into a series of small-numbered indexes, which improves the effectiveness of entropybased compression techniques such as Huffman or arithmetic coding. For a fuller understanding of the role of the BWT in data compression, consult Burrows \& Wheeler (1994) and Nelson (1996).

The inverse transform unbwt :: Ord $a \Rightarrow$ Int $\rightarrow[a] \rightarrow[a]$ is specified by

$$
\begin{equation*}
u n b w t(b w n x s)(b w p x s)=x s \tag{3}
\end{equation*}
$$

To compute unbwt we have to show how the matrix of lexicographically sorted rotations, or at least row number $t$, where $t=b w n x s$, can be recreated solely from the knowledge of its last column. To do so we need to examine lexicographic sorting in more detail.

## 3 Lexicographic sorting

Let $(\leqslant):: a \rightarrow a \rightarrow$ Bool be a linear ordering on $a$. Define $\left(\leqslant_{k}\right)::[a] \rightarrow[a] \rightarrow$ Bool inductively by

$$
\begin{array}{ll}
x s \leqslant_{0} y s & =\text { True } \\
(x: x s) \leqslant_{k+1}(y: y s) & =x<y \vee\left(x=y \wedge x s \leqslant_{k} y s\right)
\end{array}
$$

The value $x s \leqslant_{k} y s$ is defined whenever the lengths of $x s$ and $y s$ are both no smaller than $k$.

Now, let sort $\left(\leqslant_{k}\right)::[[a]] \rightarrow[[a]]$ be a stable sorting algorithm that sorts an $n \times n$ matrix, given as a list of lists, according to the ordering $\leqslant_{k}$. Thus sort $\left(\leqslant_{k}\right)$, which we henceforth abbreviate to sort $k$, sorts a matrix on its first $k$ columns. Stability means that rows with the same first $k$ elements appear in their original order in the output matrix. By definition, lexsort $=$ sort $n$.

Define cols $j=\operatorname{map}(t a k e j)$, so cols $j$ returns the first $j$ columns of a matrix. Our aim in this section is to establish the following fundamental relationship, which is the key property for establishing the existence of an algorithm for inverting the BWT. Provided $1 \leqslant j \leqslant k$, we have

$$
\begin{equation*}
\text { cols } j \cdot \text { sort } k \cdot \text { rots }=\text { sort } 1 \cdot \operatorname{cols} j \cdot \text { map rrot } \cdot \text { sort } k \cdot \text { rots } \tag{4}
\end{equation*}
$$

The function rrot denotes a single right rotation, defined by:

$$
\text { rrot } x s=\text { last } x s: \text { init } x s
$$

Equation (4) looks daunting, but take $j=n$ (so cols $j$ is the identity) and $k=n$ (so sort $k$ is a complete lexicographic sorting). Then (4) states that the following transformation on the sorted rotations is the identity: move the last column to the front and re-sort the rows on the new first column. As we will see, this implies that the permutation that produces the first column from the last column is the same as that which produces the second from the first, and so on.

To prove (4) we will need some basic properties of rotations and sorting. For rotations, one identity suffices:

$$
\begin{equation*}
\text { map rrot } \cdot \text { rots }=\text { rrot } \cdot \text { rots } \tag{5}
\end{equation*}
$$

More generally, applying a rotation to the columns of a matrix of rotations has the same effect as applying the same rotation to the rows.

For sorting we will need

$$
\begin{equation*}
\text { sort } k \cdot \text { map rrot }^{k}=(\text { sort } 1 \cdot \text { map rrot })^{k} \tag{6}
\end{equation*}
$$

where $f^{k}$ is the composition of $f$ with itself $k$ times. Equivalently, equation (6) can be read as sort $k=(\text { sort } 1 \cdot \text { map rrot })^{k} \cdot$ map $_{\text {lrot }}{ }^{k}$. This identity formalises the fact that one can sort a matrix on its first $k$ columns by first rotating the matrix to bring these columns into the last $k$ positions, and then repeating $k$ times the process of rotating the last column into first position and stably sorting according to the first column only. Since map rrot ${ }^{n}=i d$, the initial processing is omitted in the case $k=n$, and we have the standard definition of radix sort. In this context see Gibbons (1999) which deals with the derivation of radix sorting in a more general setting.

Substituting $k+1$ for $k$ in (6) and expanding the right-hand side, we obtain

$$
\operatorname{sort}(k+1) \cdot \text { map rrot }^{k+1}=\text { sort } 1 \cdot \text { map rrot } \cdot \text { sort } k \cdot \text { map rrot }^{k}
$$

Since $r^{r o t}{ }^{k} \cdot r r o t^{n-k}=r r o t^{n}=i d$ we can compose both sides with maprrot ${ }^{n-k}$ to obtain

$$
\begin{equation*}
\operatorname{sort}(k+1) \cdot \text { map rrot }=\text { sort } 1 \cdot \text { map rrot } \cdot \text { sort } k \tag{7}
\end{equation*}
$$

Finally, we shall need the following two properties of columns. First, for arbitrary $j$ and $k$ :

$$
\begin{equation*}
\operatorname{cols} j \cdot \operatorname{sort} k=\operatorname{cols} j \cdot \operatorname{sort}(j \min k)=\operatorname{sort}(j \min k) \cdot \operatorname{cols} j \tag{8}
\end{equation*}
$$

In particular, cols $j \cdot$ sort $k=$ cols $j \cdot$ sort $j$ whenever $j \leqslant k$. Furthermore, since sort $k$ sorts the list of strings by the first $k$ characters only, we have:

$$
\begin{equation*}
\text { cols } j \cdot \text { sort } k \cdot \text { perm }=\text { cols } j \cdot \text { sort } k \tag{9}
\end{equation*}
$$

whenever $j \leqslant k$ and perm is any function that permutes its argument.
Having introduced the fundamental properties (5), (7), (8) and (9), we can now prove (4). With $1 \leqslant j \leqslant k$ we reason:

$$
\begin{aligned}
& \text { sort } 1 \cdot \text { cols } j \cdot \text { map rrot } \cdot \text { sort } k \cdot \text { rots } \\
= & \{\text { by }(8)\} \\
& \text { cols } j \cdot \text { sort } 1 \cdot \text { map rrot } \cdot \text { sort } k \cdot \text { rots } \\
= & \{\text { by }(7)\} \\
& \text { cols } j \cdot \operatorname{sort}(k+1) \cdot \text { map rrot } \cdot \text { rots } \\
= & \quad\{\text { by }(8)\} \\
& \text { cols } j \cdot \text { sort } k \cdot \text { map rrot } \cdot \text { rots } \\
=\quad & \{\text { by }(5)\}
\end{aligned}
$$

$$
\begin{aligned}
& \text { cols } j \cdot \text { sort } k \cdot \text { rrot } \cdot \text { rots } \\
=\quad & \{\text { by }(9)\} \\
& \text { cols } j \cdot \text { sort } k \cdot \text { rots }
\end{aligned}
$$

Thus, (4) is established.

## 4 The derivation

Our aim is to develop a program that reconstructs the sorted matrix from its last column. In other words, we aim to construct sort $n \cdot$ rots $\cdot$ unbwt $t$. In fact, we will try to construct a more general expression cols $j \cdot \operatorname{sort} k \cdot$ rots $\cdot$ unbwt $t$ (of which the former expression is the case $j=k=n$ ) because the more general expression is used in the two variants of the BWT described in sections 5 and 6.

First, observe that for $0 \leqslant j$,

$$
\begin{equation*}
\operatorname{cols}(j+1) \cdot \text { map rrot }=\text { join } \cdot \text { fork }(\text { map last }, \text { cols } j) \tag{10}
\end{equation*}
$$

where $\operatorname{join}(x s, x s s)=z i p W i t h(:) x s x s s$ is the matrix $x s s$ with $x s$ adjoined as a new first column, and fork $(f, g) x=(f x, g x)$. Hence

$$
\left.\begin{array}{rl} 
& \text { cols }(j+1) \cdot \text { sort } k \cdot \text { rots } \cdot \text { unbwt } t \\
= & \{\text { by }(4)\}
\end{array}\right\}
$$

In particular, consider $t=b w n x s$ for an input $x s$ and $k=n$, the length of $x s$. Since $b w p=$ map last $\cdot$ sort $n \cdot$ rots, and $b w p($ unbwt $t x s)=x s$, the equality shown above reduces to:

$$
\begin{aligned}
& (\text { cols }(j+1) \cdot \text { sort } n \cdot \text { rots } \cdot \text { unbwt } t) x s \\
& \quad=(\text { sort } 1 \cdot \text { join } \cdot \text { fork }(i d, \text { cols } j \cdot \text { sort } n \cdot \text { rots } \cdot \text { unbwt } t)) x s
\end{aligned}
$$

Setting recreate $j=$ cols $j \cdot$ sort $n \cdot$ rots $\cdot$ unbwt $t$, we have just constructed a recursive definition for recreate:

$$
\begin{array}{ll}
\text { recreate } 0 & =\operatorname{map}(\text { const }[]) \\
\text { recreate }(j+1) & =\text { sort } 1 \cdot \text { join } \cdot \text { fork }(\text { id, recreate } j)
\end{array}
$$

The Haskell code for recreate is given in Figure 1. The function sortby :: $(a \rightarrow$ $a \rightarrow$ Bool $) \rightarrow[a] \rightarrow[a]$ is a stable sorting algorithm. It is identical to the standard function sortBy except for a slightly different type.

Now we know that recreate reconstructs the entire matrix, we just need to pick a particular row. Taking $j=n$, we have unbwt $t=(!!t) \cdot$ recreate $n$. The problem is that this implementation of unbwt involves computing sort 1 a total of $n$ times. To

```
recreate :: Ord a => Int -> [a] -> [[a]]
recreate \(0 \quad=\operatorname{map}\) (const [])
recreate ( \(j+1\) ) = sortby leq . join . fork (id, recreate j)
    where leq us vs = head us <= head vs
        join = uncurry (zipWith (:))
        fork (f,g) x = (f x, g x)
```

Fig. 1. Computation of recreate.
avoid repeated sorting, observe that recreate $1 y s=$ sort $y s$, where the function sort now sorts a list rather than a matrix of one column. Furthermore, for some suitable permutation $s p$ we have

$$
\text { sort ys }=\text { permby sp ys }
$$

where permby $::($ Int $\rightarrow$ Int $) \rightarrow[a] \rightarrow[a]$ applies a permutation to a list:

$$
\text { permby } p\left[x_{0}, \ldots, x_{n-1}\right]=\left[x_{p(0)}, \ldots, x_{p(n-1)}\right]
$$

It follows that

$$
\text { recreate }(j+1) y s=\text { join (permby sp ys, permby sp (recreate } j y s))
$$

Equivalently,

$$
\begin{equation*}
\text { recreate } n y s=\text { transpose }(\text { take } n(\text { iterate } 1(\text { permby sp) ys)) } \tag{11}
\end{equation*}
$$

where transpose $::[[a]] \rightarrow[[a]]$ is the standard Haskell function for transposing a matrix and iterate $1=$ tail $\cdot$ iterate . The $t$ th row of a matrix is the $t$ th column of the transposed matrix, that is, $(!!t) \cdot$ transpose $=\operatorname{map}(!!t)$, so we can use the naturality of take $n$ to obtain

$$
\text { unbwt tys }=\text { take } n(\text { map }(!!t)(\text { iterate } 1(\text { permby sp) ys }))
$$

Suppose we define spl $::$ Ord $a \Rightarrow[a] \rightarrow$ Int $\rightarrow(a$, Int $)$ by

$$
\text { spl ys } i=\operatorname{sort}(z i p y s[1 . .])!!i
$$

where sort now sorts a list of pairs. Then

$$
s p l y s j=(y s!!s p j, s p j)
$$

Hence

$$
\operatorname{map}(!!k)(\text { iterate } 1(\text { permby } s p) y s)=\text { thread }(s p l y s k)
$$

where thread $(x, j)=x$ : thread $($ spl $y s j)$.
The final algorithm, written as a Haskell program, is given in Figure 2. In a real implementation, the sorting in $s p l$ would be performed by counting the histogram of the input, which can be done in linear time using a mutable array. The "threading" part can be performed in linear time, assuming constant-time array look-up. In Seward (2001) it was observed that the main inefficiency with the algorithm lies in the cache misses involved in the threading, arising as a result of accessing a large array in non-sequential order.

```
unbwt :: Ord a => Int -> [a] -> [a]
unbwt t ys = take (length ys) (thread t)
    where spl i = sort (zip ys [0..]) !! i
        thread i = x : thread j
            where (x,j) = spl i
```

Fig. 2. Computation of unbwt.

## 5 Schindler's variation

The main variation of BWT is to exploit the general form of (4) rather than the special case $k=n$. Suppose we define

$$
\text { bwpS } k=\text { map last } \cdot \text { sort } k \cdot \text { rots }
$$

This version, which sorts only on the first $k$ columns of the rotations of a list, was considered in Schindler (1997). The derivation of the previous section shows how we can recreate the first $k$ columns of the sorted rotations from $y s=b w p k x s$, namely by computing recreate $k y$ s.

The remaining columns cannot be computed in the same way. However, we can reconstruct the $t$ th row, where $t=b w n k x s$ and

$$
\text { sort } k(\text { rots } x s)!!t=x s
$$

The first $k$ elements of $x s$ are given by create $k y s!!t$, and the last element of $x s$ is $y s!!t$. Certainly we know

$$
\operatorname{take} k(\text { rrot } x s)=\left[x_{n}, x_{1}, \ldots, x_{k-1}\right]
$$

This list begins with the last row of the unsorted matrix, and consequently, since sorting is stable, will be the last occurrence of the list in create $k y s$. If this occurrence is at position $p$, then $y s!!p=x_{n-1}$. Having discovered $x_{n-1}$, we know take $k\left(r r o t^{2} x s\right)$. This list begins the penultimate row of the unsorted matrix, and will be either the last occurrence of the list in the sorted matrix, or the penultimate one if it is equal to the previous list. We can continue this process to discover all of $\left[x_{k+1}, \ldots, x_{n}\right]$ in reverse order. Efficient implementation of this phase of the algorithm requires building an appropriate data structure for repeatedly looking up elements in reverse order in the list zip (recreate $k y s$ ) ys and removing them when found. A simple implementation is given in Figure 3.

## 6 Chapin and Tate's variation

Primarily for the purpose of showing that the pattern of derivation in this paper can be adapted to other cases, we shall consider another variation. Define the following alternative of BWT:

$$
\text { bwpCT } k=\text { map last } \cdot \text { twist }^{k} \cdot \text { lexsort } \cdot \text { rots }
$$

where the function twist rearranges the rows of the matrix. One possible choice

```
unbwt :: Ord a => Int -> Int -> [a] -> [a]
unbwt k p ys = us ++ reverse (take (length ys - k) vs)
    where us = yss !! p
        yss = recreate k ys
        vs = u:search k (reverse (zip yss ys)) (take k (u:us))
        u = ys !! p
search :: Eq a => Int -> [([a],a)] -> [a] -> [a]
search k table xs = x:search k table' (take k (x:xs))
                        where (x,table') = dlookup table xs
dlookup :: Eq a => [(a,b)] >> a -> (b,[(a,b)])
dlookup ((a,b):abs) x = if a==x then (b,abs)
    else (c,(a,b):cds)
    where (c,cds) = dlookup abs x
```

Fig. 3. Computation of Schindler's variation.

```
twist :: Eq a => Int -> [[a]] -> [[a]]
twist k = concat . mapEven (map reverse). groupby (take k)
mapEven, map0dd :: (a->a) -> [a] -> [a]
mapEven f [] = []
mapEven f (x:xs) = f x : mapOdd f xs
mapOdd f [] = []
mapOdd f (x:xs) = x : mapEven f xs
```

Fig. 4. One possible choice of twist.
of twist is shown in Figure 4. As an example, consider the rotations of the string aabab:

| aabab | ababa | abaab |
| :--- | :--- | :--- |
| abaab | abaab | ababa |
| ababa | aabab | aabab |
| baaba | baaba | babaa |
| babaa | babaa | baaba |

Shown on the left is the sorted matrix of rotations. The matrix in the middle is the result of applying twist. The rows are first partitioned into groups by groupby according to their first characters. The even numbered groups (counting from zero) are then reversed. In the example, the group starting with a is reversed. Shown on the right is the result of applying twist ${ }^{2}$ to the matrix in the middle. The rows are partitioned into three groups, starting with ab , aa, and ba respectively. The noughth and the second group are reversed.

The idea of twisting the matrix of sorted rotations was proposed in Chapin \& Tate (1998), where a similar but slightly more complicated version of twist was considered based on the Gray code. Chapin and Tate's generalisation can marginally improve the compression ratio of the transformed text.

For $t$ wist $^{k}$ to be invertible, however, we need only the property that for $0<j \leqslant k$,

In words, further twisting (twist ${ }^{k}$ where $j \leqslant k$ ) does not change the first $j$ columns after they have been set by twist ${ }^{j-1}$. In the example above, for instance, the call to twist ${ }^{2}$ does not change the first two columns of the matrix in the middle, nor do successive calls to $t w i s t^{k}$ where $k \geqslant 2$. Any twist satisfying (12) suffices to make $b w t C T$ invertible. This separation of concerns on compression rate and invertibility means that one can try many possible choices satisfying (12) and experiment with the effect on compression.

To derive an algorithm for the reverse transform we need the following analogue of (4):

$$
\begin{align*}
& \text { cols } j \cdot \text { twist }^{k} \cdot \text { lexsort } \cdot \text { rots } \\
= & \text { twist }^{k} \cdot \text { sort } 1 \cdot \text { untwist }^{k} \cdot \text { cols } j \cdot \text { map rrot } \cdot \text { twist }^{k} \cdot \text { lexsort } \cdot \text { rots } \tag{13}
\end{align*}
$$

where untwist is inverse to twist. The proof of (13) follows a similar path to the derivation in section 3. When $k=0$ (so $t$ wist ${ }^{k}=i d$ ), equation (13) reduces to a special case of (4). In words, (13) means that the following operation is an identity on a matrix generated by $t$ wist ${ }^{k}$. lexsort • rots: move the last column to the first, untwist it, sort it by the first character, and twist it again.

Based on (13) one can now derive an algorithm similar to that of section 4. Defining
we can construct a recursive definition for recreate $C T$ which is similar to (11), but with the permutation $s p$ simulating twist ${ }^{i}$ sort $1 \cdot$ untwist $^{i}$ for appropriate $i$, rather than just sort 1 . The details are more complicated than for the corresponding definition of recreate (which builds one column in each step) because in recreateCT the permutation $s p$ changes each time a new column is built. So the algorithm has to construct a new permutation as well as a new column at each step. The resulting algorithm will thus return a pair whose first component is the reconstructed matrix and the second component is a permutation representating $s p$. In the first step we build the first column and a permutation simulating twist $\cdot$ sort $1 \cdot$ untwist ; in the second step we build the second column and a permutation for twist $^{2} \cdot$ sort $1 \cdot$ untwist $^{2}$, and so on. Further details are omitted in this pearl.

## 7 Conclusions

We have shown how the inverse Burrows-Wheeler transform can be derived by equational reasoning. The derivation can be re-used to invert the more general versions proposed by Schindler, and by Chapin and Tate.

Other aspects of the BWT also make interesting topics. In Manber \& Myers (1993) it is shown how to sort the rotations of a given string in $O(n \log n)$ time using suffix arrays, where $n$ is the length of the string. How efficiently it can be done in
a functional setting remains unanswered, though we conjecture that $O\left(n(\log n)^{2}\right)$ is the best possible.

## Acknowledgements

We would like to thank Ian Bayley, Jeremy Gibbons, Ralf Hinze, Geraint Jones and Barney Stratford for constructive criticism of earlier drafts of this pearl, and to Julian Seward who made some useful comments regarding the practical aspects of the Burrows-Wheeler transform. Also, we would like to thank two anonymous referees for their detailed and useful comments, and for identifying many typos and infelicities in an earlier draft of this paper.

## References

Burrows, M. and Wheeler, D. J. (1994) A block-sorting lossless data compression algorithm. Technical report, Digital Systems Research Center. Research Report 124. (Available online at http://gatekeeper.dec.com/pub/DEC/SRC/research-reports/abstracts/ src-rr-124.html.)
Chapin, B. K. and Tate, S. (1998) Higher compression from the Burrows-Wheeler transform by modified sorting. Data Compression Conference 1998, p. 532. IEEE Press.
Gibbons, J. (1999) A pointless derivation of radixsort. J. Functional Program. 9(3), 339-346.
Manber, U. and Myers, G. (1993) Suffix arrays: A new method for on-line string searches. SIAM J. Comput. 22(5), 935-948.
Nelson, M. (1996) Data compression with the Burrows-Wheeler transform. Dr. Dobb's J. September.
Schindler, M. (1997) A fast block-sorting algorithm for lossless data compression. Data Compression Conference 1997, p. 469. IEEE Press.
Seward, J. (2000) bzip2. http://sources.redhat.com/bzip2/.
Seward, J. (2001) Space-time tradeoffs in the inverse B-W transform. Data Compression Conference 2001, pp. 439-448.

