

Attacking Scrambled Burrows-Wheeler Transform

Martin Stanek

Department of Computer Science, Comenius University
Mlynská dolina, 842 48 Bratislava, Slovak Republic
`stanek@dcs.fmph.uniba.sk`

Abstract

Scrambled Burrows-Wheeler transform [6] is an attempt to combine privacy (encryption) and data compression. We show that the proposed approach is insecure. We present chosen plaintext and known plaintext attacks and estimate their complexity in various scenarios.

1 Introduction

The Burrows-Wheeler transform (BWT) [2] is a commonly used transform in lossless compression algorithms. The BWT does not compress the data itself, instead it is usually the first step in a sequence of algorithms transforming an input data into compressed data. The most prominent example of BWT-based compression is bzip2 program [8], which uses basically the following sequence of algorithms: the BWT, the move-to-front transform (MTF), and Huffman coding.

In practice, there are many situations with the simultaneous requirements of data compression and privacy (encryption). A common approach is compress-then-encrypt paradigm with independent algorithms used for compression and encryption. While widely preferred, the algorithms should be properly combined in order to avoid possible attacks [1, 5]. Another approach is to unify compression and encryption. However, designing a secure, “encrypting” variant of compression method is not an easy task. Many attempts were broken successfully [3, 7, 9].

Recently, Oğuzhan Külekci [6] proposed a novel approach – scrambled BWT – to combine data compression with privacy requirement. The scrambled BWT uses a secret lexicographic order of underlying alphabet as a secret key. In order to thwart some weaknesses, the author proposed to accompany the scrambled BWT with modified MTF that uses the secret lexicographic order as well.

Our contribution We show that the proposed scrambled BWT with MTF is completely insecure and can be attacked easily (in the sense of chosen plaintext as well as known plaintext attacks). In case of known plaintext attacks we estimate experimentally the bit security of the scrambled BWT with MTF in various scenarios.

We briefly introduce the “standard” BWT, the MTF, and the scrambled BWT in Section 2. Section 3 contains our analysis of the proposal, and shows chosen-plaintext and known-plaintext attacks on the scrambled BWT with MTF.

2 Preliminaries

Let \mathcal{A} be an alphabet with size $L = |\mathcal{A}|$. Let N denotes a block length. A cyclic rotation of string/block $x = x_0x_1 \dots x_{N-1} \in \mathcal{A}^N$ with offset k is string/block $x^{(k)} = x_kx_{k+1} \dots x_{k+N}$ where all the indices are computed modulo N . The w -th symbol of $x^{(k)}$ is denoted as $x_w^{(k)}$, for $0 \leq w < N$.

For real-world scenarios one can expect $L = 256$ (using bytes as an alphabet) and the block size N several hundreds kilobytes (e.g. bzip2 uses default block size 900kB).

2.1 “Standard” BWT

Let $x = x_0x_1 \dots x_{N-1} \in \mathcal{A}^N$ be an input block. The BWT sorts all cyclic rotations $x^{(0)}, x^{(1)}, \dots, x^{(N-1)}$ of input block x in lexicographic order. Let j_0, j_1, \dots, j_{N-1} be a permutation of $\{0, 1, \dots, N-1\}$ such that

$$x^{(j_0)} < x^{(j_1)} < \dots < x^{(j_{N-1})}.$$

Then the result of the BWT is a string consisting of the last symbols from each of the sorted cyclic rotations:

$$x_{N-1}^{(j_0)}, x_{N-1}^{(j_1)}, \dots, x_{N-1}^{(j_{N-1})}.$$

In order to facilitate the inversion transformation, an additional pointer is used to remember the position of the original string $x = x^{(0)}$, i.e. t such that $j_t = 0$. Since the inversion transformation is not important for our exposition, we will not describe it here.

Example 1. Let $\mathcal{A} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$ and $N = 10$. Let $x = \mathbf{cdababcdaa}$ be an input block. Applying BWT yields $\mathbf{ddbaaabacc}$, and $t = 7$.

The paper [6] uses a variant of the BWT with special symbol denoting the end of block. In that case, there is no need to remember the position of the original block among sorted cyclic rotations. The analysis done in Section 3 is valid for this variant as well.

2.2 MTF

The MTF transforms an input string $x = x_0x_1 \dots x_{N-1} \in \mathcal{A}^N$ into sequence of N numbers $\{p_i\}_{i=0}^{N-1}$, where $p_i \in \{0, 1, \dots, L-1\}$. The algorithm maintains a table $T[0, \dots, L-1]$ of all symbols from \mathcal{A} , initially sorted in lexicographic order. For each $i = 0, 1, \dots, N-1$ the symbol x_i is processed as follows:

1. p_i is the position of x_i in table T ;
2. T is modified: x_i is moved to the front/top of the table.

It is easy to see that the MTF is invertible. The main idea of MTF is that the recently used symbols are encoded as small integers. This makes its output a suitable data for subsequent compression by simple entropy coders such as Huffman or arithmetic coding.

Example 2. Let $x = \text{ddbaaabacc}$ be an input string. The output of MTF is 3022001130.

2.3 Scrambled BWT with MTF

Oğuzhan Külekci [6] proposed scrambled BWT, where a secret (encryption) key is a secret lexicographic order of symbols in \mathcal{A} . One of the claimed advantages is a large key space, for $L = 256$ it is $256!$ keys. The author observed, that using scrambled BWT is not secure enough, and can be attacked by exploiting known statistical relationships among plaintext symbols (e.g. digram frequencies). Therefore he accompanied the scrambled BWT (sBWT) with MTF, where the secret lexicographic order is applied as well (i.e. the initialization of T depends on this order):

“... Thus, initializing the alphabet ordering in MTF with the secret lexicographic order used in sBWT provides protection against that statistical attack.”

Example 3. Let $x = \text{cdababcdaa}$ be an input block. We show the result of sBWT and MTF with various lexicographic orders:

key k	sBWT with k	MTF with k
$c < a < d < b$	\longrightarrow baadbdccea ($t = 1$)	\longrightarrow 3203213030
$d < c < b < a$	\longrightarrow ccabaaabdd ($t = 2$)	\longrightarrow 1033100130
$a < c < b < d$	\longrightarrow dadbbaaacc ($t = 5$)	\longrightarrow 3113020030

3 Security analysis

Let us denote the secret lexicographic order as k and the corresponding scrambled BWT as sBWT_k . Similarly, the MTF with secret lexicographic order is denoted as MTF_k . Let x be an input block. The author proposes [6] the same secret key (secret lexicographic order) for both transformations: $\text{MTF}_k(\text{sBWT}_k(x))$.

Let us note that attacks described in this sections will work even for situation with two independent secret keys: $\text{MTF}_{k_2}(\text{sBWT}_{k_1}(x))$. The attack will “undo” the MTF (revealing the value of $\text{sBWT}_{k_1}(x)$), and the sBWT part can be attacked by exploiting the statistical properties of plaintext as suggested in [6].

Remark 1. We can ignore other pre- and post-processing steps in the compression algorithm, since they do not depend on the key.

We base our analysis on the following observations:

1. The scrambled BWT, with secret lexicographic order as a key, keeps the frequencies of symbols intact, i.e. when symbol ‘a’ appears l times in an input block, then ‘a’ will appear exactly l times in the output block. One can view sBWT_k as a plaintext-dependent permutation cipher.
2. For any string z and any two lexicographic orders k_1, k_2 , performing $y = \text{MTF}_{k_2}^{-1}(\text{MTF}_{k_1}(z))$ can be viewed as a monoalphabetic substitution, i.e. it change the symbols but it does not change the frequencies. For example, ‘a’ can become ‘w’, ‘b’ can become ‘f’ ...but in such case the number of a’s in z is the same as the number of w’s in y , the number of b’s in z is the same as the number of f’s in y etc.

For the rest of the section we denote an input block x and the resulting block $y = \text{MTF}_k(\text{sBWT}_k(x))$. Our analysis is done primarily with single data block (it is sufficient for most scenarios). However, it can be extended to multiple blocks in a straightforward manner, see Section 3.3.

3.1 Chosen plaintext attack

The goal of the attack is to identify a secret key (lexicographic order):

1. Construct input block x , where symbols from \mathcal{A} have unique frequencies.
2. Compute $z^* = \text{MTF}_{k'}^{-1}(y)$ for an arbitrary lexicographic order k' .
3. Pair symbols in x and z^* according their frequencies, and recover the correct “middle” value $z = \text{sBWT}_k(x) = \text{MTF}_k^{-1}(y)$.

4. Reconstruct the key from z and y .

Let $\mathcal{A} = \{a_1, \dots, a_L\}$. The chosen input block can be, for example, $a_1 a_2 a_2 a_3 a_3 a_3 \dots$ finishing with L symbols a_L . Then we can reconstruct the key from this single block as long as the block size is at least $\binom{L+1}{2}$ (for $L = 256$ it is 32 896). The complexity of the attack is $\Theta(L^2)$ (counting all its steps).

Example 4. Let us illustrate the attack with the following toy example. Let $\mathcal{A} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$ and $N = 10$. We choose the input block $x = \mathbf{abbccddddd}$, and observe the output $y = \mathbf{3133002202}$. We apply inverse MTF with natural lexicographic order ($\mathbf{a} < \mathbf{b} < \mathbf{c} < \mathbf{d}$): $z^* = \mathbf{dacbbbdccb}$. Pairing symbols with equal frequencies yields $z = \mathbf{abcdddbccd}$. Knowing “middle” value z and the result of $\text{MTF}_k(z)$, i.e. y , we can easily reconstruct the secret lexicographic order: $\mathbf{b} < \mathbf{d} < \mathbf{c} < \mathbf{a}$.

3.2 Known plaintext attack (single block)

Known plaintext attack extends the previous attack assuming that the attacker cannot control the frequencies of particular symbols in the input data. However, considering the block sizes used in practice (e.g. default 900kB block size in bzip2), the probability of equal frequencies of symbols is rather low. Moreover, longer blocks usually lead to better compression.

Remark 2. For short blocks, one can assume that the attacker will know the input data for multiple blocks. Therefore he can combine results from these blocks and significantly reduce the complexity of the attack further, see Section 3.3.

Let us denote by $\#_v(x)$ the number of symbols $v \in \mathcal{A}$ in the string x . Let $C(x) = \{\#_v(x) \mid v \in \mathcal{A}\}$ be the set of all distinct values of $\#_v(x)$. For each value $r \in C(x)$ we define $\text{size}(x, r)$ to be the number of symbols having exactly r occurrences in x , i.e. $\text{size}(x, r) = |\{v \in \mathcal{A} \mid \#_v(x) = r\}|$.

The attacker proceeds similarly to the chosen plaintext attack, computing $z^* = \text{MTF}_{k'}^{-1}(y)$ for an arbitrary lexicographic order k' . Then he tries to pair symbols in x and z^* according their frequencies to recover the correct “middle” value $z = \text{sBWT}_k(x) = \text{MTF}_k^{-1}(y)$. However, this time there is no guarantee of unique frequencies, therefore the attacker can perform an exhaustive search for all assignments of symbols in groups with equal frequencies. For each assignment, the attacker computes a corresponding lexicographic order and performs an inverse BWT with this order. Comparing the result with the original plaintext gives a confirmation/rejection of particular assignment. The size of the search space is

$$\text{SS}(x) = \prod_{r \in C(x)} \text{size}(x, r)!$$

We measure the complexity of our known plaintext attack as bit security of the cipher, i.e. binary logarithm of the corresponding search space size: $\log_2 \text{SS}(x)$. In order to estimate the bit security, we performed the following experiments:

RandBytes We generate the input block as a stream of randomly and independently generated bytes (i.e. $L = 256$) with uniform distribution. Since in real-world one can expect much more “compressible” input block (with greater variation of symbols frequencies), our model simulates the worst situation for the attacker. Therefore, the estimations obtained in this experiment can be viewed as upper bounds for bit security of the cipher.

RandReduced This is a similar experiment as the previous one, but this time we restrict possible symbols to the set of $L = 100$ symbols (generated with uniform distribution). The rest of the symbols are not generated.

RandText This and the last experiment (RandKernel) are probably the most realistic of our experiments. In RandText we model the input block as a stream of independently generated bytes, where the probabilities of individual symbols correspond to the probabilities of symbols in a novel *Crime and Punishment* [4].

RandKernel Similar experiment to RandText. However, this time the probabilities of individual symbols (bytes) correspond to the probabilities of symbols in `kernel132.dll` file in Windows 7.

We gradually increased N in each experiment and computed the average bit security of the cipher. The average value was computed from 1000 samples. It is interesting to see the level of bit security degradation from the theoretical level: $\log_2(256!) \sim 1684$ bits. The values for RandReduced, RandText and RandKernel experiments show that the scrambled BWT with MTF offers practically no security. The results are shown in Table 1.

N	RandBytes	RandReduced	RandText	RandKernel
50 000	378.7	28.6	10.1	65.3
100 000	304.2	20.9	8.4	49.0
150 000	263.7	17.2	8.8	41.4
200 000	236.4	15.2	9.2	36.5

Table 1: Bit security for single block KPA (with block size N)

3.3 Known plaintext attack (multiple blocks)

The known plaintext attack from previous section can be easily extended for multiple blocks, with improved performance. The majority of “ties” (groups of symbols with equal frequencies) in one block can be broken by considering the frequencies in other blocks. In order to illustrate this effect, we performed experiments similar to those in the previous section (RandBytes, RandReduced, RandText, and RandKernel). However, this time we used two blocks of known plaintext and 10 times smaller block sizes. The results are shown in Table 2 (again, 1000 samples were used to estimate individual values). Even with such artificially small block sizes one can observe further decrease in the bit security of the scrambled BWT with MTF.

N	RandBytes	RandReduced	RandText	RandKernel
5 000	113.3	4.0	8.4	44.6
10 000	61.0	1.9	5.0	21.4
15 000	41.8	1.3	3.6	13.9
20 000	31.6	1.0	3.2	10.4

Table 2: Bit security for two blocks KPA (with block size N)

4 Conclusion

Usually, providing privacy (encryption) by modifying the data compression techniques is not a good idea in practice. We demonstrated the security problems of the scrambled BWT with MTF. Moreover, it seems that these problems cannot be easily fixed. Further extensions of our attacks can be aimed at known plaintext attack with only partial block knowledge, or even ciphertext only attacks. However, the already identified weaknesses of the cipher allow us to conclude that the cipher is insecure.

Acknowledgment This work was supported by P12/1.

References

- [1] Eli Biham and Paul C. Kocher. A Known Plaintext Attack on the PKZIP Stream Cipher. *Fast Software Encryption, FSE'94*, Lecture Notes in Computer Science vol. 1008, Springer, pp. 144–153, 1995.
- [2] M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical Report 124, Digital Equipment Corporation, 1994.

- [3] John G. Cleary, Sean A. Irvine and Ingrid Rinsma-Melchert. On the insecurity of arithmetic coding. *Computers & Security*, 14(2):167–180, 1995.
- [4] Fyodor M. Dostoevsky. *Crime and Punishment*. Project Gutenberg, vol. 2556, English translation by C. Garnett, 2006. <http://www.gutenberg.org/ebooks/2554>
- [5] Tadayoshi Kohno. Attacking and repairing the winZip encryption scheme, *ACM Conference on Computer and Communications Security*, CCS'04, pp. 72–81, 2004.
- [6] M. Oğuzhan Külekci. On scrambling the Burrows-Wheeler transform to provide privacy in lossless compression, *Computers & Security*, 31(1):26–32, 2012.
- [7] Jen Lim, Colin Boyd and Ed Dawson. Cryptanalysis of Adaptive Arithmetic Coding Encryption Schemes. *Information Security and Privacy, Second Australasian Conference, ACISP'97*, Lecture Notes in Computer Science vol. 1270, Springer, pp. 216–227, 1997.
- [8] J. Seward. bzip2, <http://www.bzip.org>, 1996-2012.
- [9] Jiantao Zhou, Oscar C. Au and Peter Hon-Wah Wong. Adaptive chosen-ciphertext attack on secure arithmetic coding. *IEEE Transactions on Signal Processing*, 57(5):1825–1838, 2009.