

## Fountain Codes for Lossless Data Compression

Giuseppe Caire, Shlomo Shamai, Amin Shokrollahi, and Sergio Verdú

**ABSTRACT.** This paper proposes a universal variable-length lossless compression algorithm based on fountain codes. The compressor concatenates the Burrows-Wheeler block sorting transform (BWT) with a fountain encoder, together with the closed-loop iterative doping algorithm. The decompressor uses a Belief Propagation algorithm in conjunction with the iterative doping algorithm and the inverse BWT. Linear-time compression/decompression complexity and competitive performance with respect to state-of-the-art compression algorithms are achieved.

### 1. Introduction

Noiseless data compression is a key information technology used in innumerable data storage and transmission applications ranging from computer operating systems to modems to lossy compression standards. Although the state-of-the-art has reached a certain level of maturity with data compression algorithms that asymptotically attain the fundamental information theoretic limits with linear computational complexity, they suffer from several shortcomings when used in packetized noisy channels. Mainly for this reason, no payload compression is currently implemented in third-generation standards for high-speed wireless data transmission.

It is known that *linear* fixed-length encoding can achieve for asymptotically large blocklength the minimum achievable compression rate for memoryless sources [1] and for arbitrary (not necessarily stationary/ergodic) sources [2].

After initial attempts [3, 4, 5] to construct linear lossless codes were nonuniversal, limited to memoryless sources and failed to reach competitive performance with standard data compression algorithms, the interest in linear data compression waned. Recently [6, 7, 8] came up with a universal lossless data compression algorithm based on irregular low-density parity-check codes which has linear encoding and decoding complexity, can exploit source memory and in the experiments for binary sources presented in [6, 7, 8] showed competitive performance with respect to standard compressors such as `gzip`, PPM and `bzip`.

The scheme of [6, 7, 8] was based on the important class of sparse-graph error correcting codes called low-density parity check (LDPC) codes. The block-sorting

---

1991 *Mathematics Subject Classification.* Primary 68P30, 94A29; Secondary 94A45, 62B10.

*Key words and phrases.* Noiseless Data Compression, Universal algorithms, Error Correcting Codes, Source Coding, Sources with Memory, Block Sorting Transform.

This research was partially supported by NSF Grant CCR-0312879.

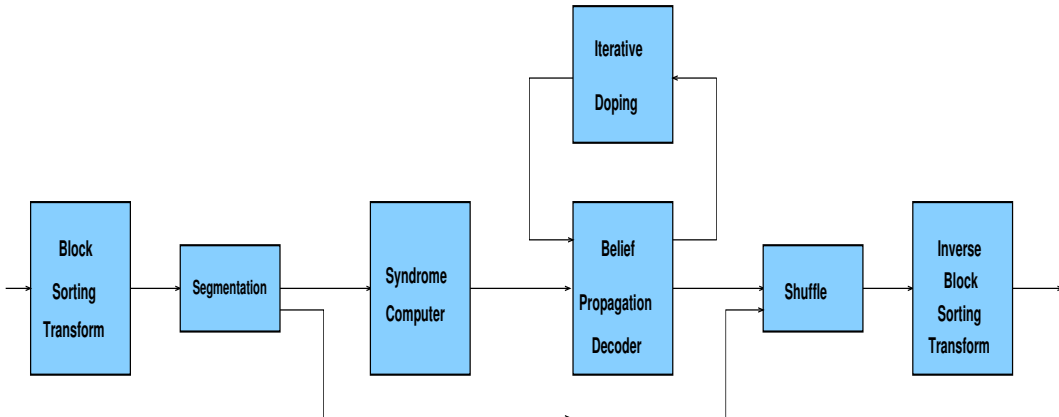


FIGURE 1. Compression/Decompression Scheme for Noiseless Channel

transform (or Burrows-Wheeler transform (BWT)) [9] shown in Figure 1 is a one-to-one transformation, which performs the following operation: it generates all cyclic shifts of the given data string and sorts them lexicographically. The last column of the resulting matrix is the BWT output from which the original data string can be recovered, knowing the *BWT index* which is the location of the original sequence in the matrix. The BWT shifts redundancy in the memory to redundancy in the marginal distributions. The redundancy in the marginal distributions is then much easier to exploit at the decoder as the decoding complexity is independent of the complexity of the source model (in particular, the number of states for Markov sources). The output of the BWT (as the blocklength grows) is asymptotically piecewise i.i.d. for stationary ergodic tree sources. The length, location, and distribution of the i.i.d. segments depend on the statistics of the source. The existing universal BWT-based methods for data compression generally hinge on the idea of compression for a memoryless source with an adaptive procedure that learns implicitly the local distribution of the piecewise i.i.d. segments, while forgetting the effect of distant symbols.

In the data compression algorithm of [6, 7], the compression is carried out by multiplication of the Burrows-Wheeler Transform of the source string with the parity-check matrix of an error correcting code. Of particular interest are LDPC codes since the belief propagation (BP) decoder is able to incorporate the time-varying marginals at the output of the BWT in a very natural way. The nonidentical marginals produced at the output of the BWT have a synergistic effect with the BP algorithm which is able to iteratively exploit imbalances in the reliability of variable nodes. The universal implementation of the algorithm where the encoder identifies the source segmentation and describes it to the decompressor is discussed in [8].

An important ingredient in the compression scheme of [6, 7, 8] is the ability to do decompression at the compressor. This enables to tune the choice of the codebook to the source realization and more importantly it enables the use of the Closed-Loop Iterative Doping (CLID) algorithm of [6, 2]. This is an efficient algorithm which enables zero-error variable-length data compression with performance which is quite competitive with that of standard data compression algorithms.

In this paper, instead of adopting irregular low-density parity check codes of a given rate approximately matched to the source we adopt a different approach based on rateless *fountain codes*. This class of codes turns out to be more natural for variable-length data compression applications than standard block codes and achieves in general comparable performance to the LDPC-based scheme of [6, 7, 8].

The rest of the paper is organized as follows. Section 2 reviews the main features of fountain codes for channel coding. Section 3 gives a brief summary of the principle of belief propagation decoding which is common to both channel and source decoding. Our scheme for data compression with fountain codes is explained in detail in Section 4 in the setting of nonuniversal compression of binary sources. For further background on linear codes for data compression and the closed-loop iterative doping algorithm the reader is referred to [2]. The issues related to implementation with nonbinary alphabets are discussed in Section 5, while the modelling module necessary for universal application is discussed in Section 6. Section 7 shows several experiments and comparisons of redundancy with off-the-shelf data compression algorithms run with synthetic sources.

## 2. Fountain Codes for Channel Coding

Fountain codes [10, 11] form a new class of sparse-graph codes designed for protection of data against noise in environments where the noise level is not known a-priori. To achieve this, a fountain code produces a potentially limitless stream of output symbols for a given vector of input symbols. In practical applications, each output symbol is a linear function of the input symbols, and the output symbols are generated independently and randomly, according to a distribution which is chosen by the designer. The sequence of linear combinations of input symbols is known at the decoder. For example, they can be generated by pseudorandom generators with identical seeds.

The *overhead* of the decoding algorithm for a fountain code over a given communication channel is measured as the fraction of additional output symbols necessary to achieve the desired residual error rate. Here, the phrase “additional” is meant to be with respect to the Shannon limit of the underlying channel.

The decoding process for fountain codes is as follows: the receiver collects output symbols and estimates for each received output symbol the amount of information in that symbol. For example, when output symbols are bits, this measure of information can be obtained from the log-likelihood ratio (LLR) of the received bit. If this ratio is  $\lambda$ , then the empirical mutual information between the information symbol and its LLR equals  $1 - h(p)$ , where  $p = 1/(1 + \exp(\lambda))$ . The decoder stops collecting output bits as soon as the accumulated information carried by the observed channel outputs exceeds  $(1 + \omega)k$ , where  $\omega$  is the overhead associated with the fountain code, and  $k$  is the number of input symbols. Then the decoder uses its BP decoding algorithm to recover the input symbols from the information contained in the output symbols.

If the amount of information in the received output symbols is unknown but the channel is known to be memoryless stationary with capacity  $C(\lambda)$  parameterized in the single output LLR  $\lambda$ , then decoding of fountain codes can be accomplished as follows: the decoder starts with an optimistic guess  $\lambda_1$  of the channel parameter, collects an appropriate number of output symbols  $n_1$  such that  $k/n_1 = C(\lambda_1) - \delta$ , where  $\delta > 0$  is some positive rate margin that enforces successful decoding with

high probability, and applies BP decoding based on the guess  $\lambda_1$ . In case the BP decoder is unsuccessful, a predetermined number of additional output symbols is collected such that the total number of output symbols is  $n_2$ , and the BP decoding is applied to the new graph using LLR  $\lambda_2 = C^{-1}(k/n_2 + \delta)$ . In case the BP decoder is unsuccessful, the same process is repeated using  $n_3 > n_2$  output symbols and so on, until decoding is successful. In practice, instead of resetting the BP decoder at each new decoding attempt and especially if the initial guess  $\lambda_1$  is likely to be not far from the true channel parameter, it might be more convenient to keep the same BP decoder running while incorporating the additional collected output symbols.

Discovered by Michael Luby [10], LT-codes are one of the first classes of efficient fountain codes for the erasure channel. An extension of this class of codes is the class of Raptor codes [12]. These classes of codes are very well suited for solving the compression problem, because the compression algorithm translates to a channel coding problem for a discrete memoryless channel with time-varying transition probability matrix which depends on the source. In applications, it is undesirable to tune the choice of the code to the sequence of transition probabilities, as in universal applications they are not known beforehand. In this case a fountain code is more amenable for universal implementation than other classes of efficient channel codes, such as irregular LDPC codes [13], mainly because a single code can be used regardless of the source rate.

For ease of exposition we concentrate on binary fountain codes. Let  $k$  be a positive integer, and let  $\mathcal{D}$  be a distribution on  $\mathbb{F}_2^k$ . A Fountain Code ensemble with parameters  $(k, \mathcal{D})$  has as its domain the space  $\mathbb{F}_2^k$  of binary strings of length  $k$ , and as its target space the set of all sequences over  $\mathbb{F}_2$ , denoted by  $\mathbb{F}_2^{\mathbb{N}}$ . Formally, a Fountain Code ensemble with parameters  $(k, \mathcal{D})$  is a linear map  $\mathbb{F}_2^k \rightarrow \mathbb{F}_2^{\mathbb{N}}$  represented by an  $\infty \times k$  matrix where rows are chosen independently with identical distribution  $\mathcal{D}$  over  $\mathbb{F}_2^k$ . The blocklength of a Fountain Code is potentially infinite, but in our data compression applications we will solely consider truncated Fountain Codes with a number of coordinates which is tailored to the source realization.

The symbols produced by a Fountain Code are called *output symbols*, and the  $k$  symbols from which these output symbols are calculated are called *input symbols*. The input and output symbols could be elements of  $\mathbb{F}_2$ , or more generally the elements of any finite dimensional vector space over  $\mathbb{F}_2$ . In this paper, we are primarily concerned with Fountain Codes over the field  $\mathbb{F}_2$ , in which symbols are bits.

A special class of Fountain Codes is furnished by LT-Codes [10]. In this class, the distribution  $\mathcal{D}$  has a special form described in the following. Let  $(\Omega_1, \dots, \Omega_k)$  be a distribution on  $\{1, \dots, k\}$  so that  $\Omega_i$  denotes the probability that the value  $i$  is chosen. Often we will denote this distribution by its generator polynomial  $\Omega(x) = \sum_{i=1}^k \Omega_i x^i$ . The distribution  $\Omega(x)$  induces a distribution on  $\mathbb{F}_2^k$  in the following way: For any vector  $v \in \mathbb{F}_2^k$  the probability of  $v$  is  $\Omega_w / \binom{k}{w}$ , where  $w$  is the Hamming weight of  $v$ . Abusing notation, we will denote this distribution in the following by  $\Omega(x)$  again. An LT-code is a Fountain code with parameters  $(k, \Omega(x))$ .

The decoding graph of length  $n$  of a fountain code with parameters  $(k, \Omega(x))$  is a bipartite graph with  $k$  nodes on one side (called *input nodes*, corresponding to the input symbols) and  $n$  nodes on the other side (called *output nodes*). The output nodes correspond to  $n$  output symbols collected at the output of the channel. The decoding graph is the *Tanner graph* of the linear encoder  $\mathbb{F}_2^k \rightarrow \mathbb{F}_2^n$  obtained by

restricting the fountain encoder mapping to those  $n$  components actually observed at the output.

A raptor code [12] performs a pre-coding operation with a linear code (e.g., an LDPC code) prior to using an LT-code. Without a pre-coder the average degree of an LT-coder needs to grow at least logarithmically in the length of the input to guarantee a small error probability, since otherwise there would exist input symbols that are not present in any of the linear equations generating the output symbols. The pre-code lowers the error floor present in an LT-code with small average degree.

### 3. Belief Propagation Decoding

In this section we give a description of the BP algorithm that is used in the decoding process of LT codes. The algorithm proceeds in rounds. In every round messages are passed from input nodes to output nodes, and then from output nodes back to input nodes. The message sent from the input node  $\iota$  to the output node  $\omega$  in the  $\ell$ th round of the algorithm is denoted by  $m_{\iota,\omega}^{(\ell)}$ , and similarly the message sent from an output node  $\omega$  to an input node  $\iota$  is denoted by  $m_{\omega,\iota}^{(\ell)}$ . These messages are scalars in  $\overline{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$ . We will perform additions in this set according to the following rules:  $a + \infty = \infty$  for all  $a \neq -\infty$ , and  $a - \infty = -\infty$  for all  $a \neq \infty$ . The values of  $\infty - \infty$  and  $-\infty + \infty$  are undefined. Moreover,  $\tanh(\infty/2) := 1$ , and  $\tanh(-\infty/2) := -1$ .

Every value  $Y$  received from the channel has a log-likelihood ratio defined as  $\ln(\Pr[Y = 0]/\Pr[Y = 1])$ .

For every output node  $\omega$ , we denote by  $Z_\omega$  the corresponding log-likelihood ratio.

In round 0 of the BP algorithm, the output nodes send the value 0 to all their adjacent input nodes. Thereafter, the following update rules are used to obtain the messages passed at later rounds:

$$(3.1) \quad \tanh\left(\frac{m_{\omega,\iota}^{(\ell)}}{2}\right) := \tanh\left(\frac{Z_\omega}{2}\right) \cdot \prod_{\iota' \neq \iota} \tanh\left(\frac{m_{\iota',\omega}^{(\ell)}}{2}\right),$$

$$(3.2) \quad m_{\iota,\omega}^{(\ell+1)} = \sum_{\omega' \neq \omega} m_{\omega',\iota}^{(\ell)},$$

where the product is over all input nodes adjacent to  $\omega$  other than  $\iota$ , and the sum is over all output nodes adjacent to  $\iota$  other than  $\omega$ , and  $\ell \geq 0$ .

After running the BP-algorithm for  $m$  rounds, the log-likelihood of each input symbol associated to node  $\iota$  can be calculated as the sum  $\sum_{\omega} m_{\omega,\iota}^{(m)}$ , where the sum is over all the output nodes  $\omega$  adjacent to  $\iota$ . In cases where the pre-code of the Raptor code is decoded separately from its LT-code, one may gather the log-likelihood of the input symbols, and run a decoding algorithm for the pre-code (which may itself be the BP algorithm). In that case, the prior log-likelihoods of the inputs are set to be equal to the calculated log-likelihoods according to (3.1).

In data compression applications, the CLID algorithm introduced in [6] runs BP at the encoder and supplies to the decoder the value of the lowest reliability symbol at certain iterations until successful decoding is achieved. Since the decoder runs an identical copy of the BP iterations, it knows the identity of the lowest reliability symbol, without the need to communicate this information explicitly. The symbols supplied by the CLID algorithm are referred to as doped symbols.

#### 4. Data Compression with Fountain Codes

Fountain Codes can be applied to the problem of data compression using the methods described in [2, 6, 7, 8]. In this section we discuss three variants of this approach that use LT-Codes in the setting of nonuniversal compression of binary sources.

**4.1. Raptor-code fountain approach.** In this approach we use as data compressor the encoder of a raptor code, and as decompressor the decoder of a raptor code incorporating the statistics of the source. For a fountain application where a plurality of receivers obtain erased versions of the broadcast data with erasure rates that are unknown at the transmitter, this approach is able to carry out joint source-channel encoding-decoding effectively, both taking into account the redundancy of the source and offering protection against channel erasures. If the source has memory and its statistics are known to the receiver, it is easy to adapt the BWT-based method for fountain applications. In the universal case it is necessary to convey reliably the model (Section 6) to the receivers (using a raptor code for example). In applications where a systematic fountain code is used to send uncompressed redundant data, backward compatibility with receivers that do not do any data decompression is achievable using the method of [14].

In conventional (non-fountain) applications with one encoder and one decoder, the approach of this subsection can also be used. In this case, it is necessary to truncate the semiinfinite sequence of compressed bits at the compressor. In a pure compression application where no robustness against channel noise is required, the decompressor can be run at the encoder and the output sequence is truncated as soon as the BP converges to the true source realization. In those unlikely source realizations in which this method offers no compression, the source realization can be sent uncompressed. As it can be expected (and we will see below) this method is not competitive in terms of performance with the CLID-based approach explained in Section 4.4. An alternative to using a raptor code in this setting is to puncture a given code increasingly until the decoder is unable to reconstruct the original source sequence. This is the “decremental redundancy” approach of [15].

In noisy channel applications, the truncation is not carried out as soon as the BP converges. Instead, a number of extra parity check bits are produced in order to cope with the channel noise. For a given output rate, the interesting performance measure is the block error rate as a function of the channel noise level. If the code were a true linear random code and the decoder were maximum a posteriori, then asymptotically it would be possible to achieve an overall coding rate (channel bits/source bits) equal to the source entropy divided by the channel capacity.

**4.2. LT-CLID direct approach.** For simplicity we limit here our discussion to the compression of a block of  $k$  binary symbols  $(x_1, \dots, x_k)$  independently drawn from a binary source with bias  $p$ .

As an alternative to the pre-code of the raptor code, we could envision reliance on CLID to provide good performance in conjunction with an LT code used in much the same way as the LDPC code used in [2, 8]:  $m$  output bits are generated by the encoder of a  $(k, \Omega(x))$  LT code, and  $d$  bits are generated by the CLID algorithm. The number of bits generated by the CLID algorithm is sufficient to result in successful decoding of the  $k$  input symbols under the assumption that every symbol is 1 with probability  $p$ . This method resembles the method of [6] with the difference

that the check nodes of the LDPC code have a degree distribution equal to  $\Omega(x)$ , and the variable nodes have a binomial degree distribution. If the number of edges in the graph is large, this distribution is close to the Poisson distribution with mean  $\alpha m/k$ , where  $\alpha$  is the average degree of the distribution  $\Omega(x)$ . In particular, the fraction of variable nodes not covered by any edge in the graph is close to  $e^{-\alpha m/k}$ . This means that the number of doped symbols in this construction is on average more than  $ke^{-\alpha m/k}$ , and that the number of output symbols of the compressor is on average more than  $k(e^{-\alpha m/k} + m/k)^1$ . A necessary condition to achieve an output length within distance  $\delta$  from the Shannon limit is

$$(4.1) \quad m/k \leq h(p) + \delta - e^{-\alpha m/k}$$

On the other hand, in order for decoding to be successful the subgraph containing all  $k(1 - e^{-\alpha m/k})$  covered variable nodes and the  $m$  output nodes must correspond to a linear code of rate not smaller than  $h(p)$ . Hence, a second necessary condition is given by

$$(4.2) \quad m/k \geq h(p)(1 - e^{-\alpha m/k})$$

Letting  $m/k = z \in [0, 1]$ , the condition imposed on  $\alpha$  and  $\delta$  for given  $h(p)$  by (4.1) and (4.2) is that the root of the equation  $h(p)(1 - e^{-\alpha z}) = z$  must be smaller than the largest root of the equation  $h(p) + \delta - e^{-\alpha z} = z$ .

A potential shortcoming of this method (which motivates the consideration of the alternative method in Section 4.3) is that sparse-graph codes with check degree distribution  $\Omega(x)$  and variable degree Poisson distribution are known to perform poorly under iterative decoding on the erasure channel and other symmetric channels.

**4.3. Two-stage approach.** In this subsection we describe the principle of a different method of using LT-codes that is much better suited for the data compression application. The detailed description of the algorithm can be found in Section 4.4.

We calculate from the input symbols  $(x_1, \dots, x_k)$  a vector of *intermediate symbols*  $(y_1, \dots, y_k)$  through a linear invertible  $k \times k$  matrix  $\mathbf{G}$ :

$$(4.3) \quad (y_1, \dots, y_k)^T = \mathbf{G}^{-1}(x_1, \dots, x_k)^T$$

Next, using the intermediate symbols we calculate  $m$  output symbols  $(x_{k+1}, \dots, x_{k+m})$  according to a distribution  $\Omega(x)$ .<sup>2</sup> These  $m$  output symbols, together with the doped symbols obtained from the closed-loop iterative doping algorithm constitute the output of the compressor; hence, their total number should be as close as possible to  $kh(p)$ .

The Tanner graph on which the BP is run takes the form shown in Figure 2: a bipartite graph with  $k$  *intermediate nodes* (corresponding to the intermediate symbols) on one side and  $k$  input and  $m$  output nodes on the other side, corresponding to the  $k$  input and  $m$  output symbols, respectively. The variable nodes in that graph are indicated by circles in Figure 2. The invertible matrix  $\mathbf{G}$  is chosen as

<sup>1</sup>All these numbers are only estimates due to the approximation of the binomial distribution with the Poisson distribution; the real numbers are very close to these estimates.

<sup>2</sup>The choice of  $\Omega(x)$  is crucial for performance and its specific expression is given at the end of the section.

a random realization so that the degree distribution of the input nodes is  $\Omega(x)$ .<sup>3</sup> In this way, the degree distribution of both the input and output nodes (i.e., all nodes on the left in the graph of Figure 2) is equal to  $\Omega(x)$ . Notice that the resulting graph can be *interpreted* as the decoding graph of a  $(k, \Omega(x))$  LT code with input symbols  $(y_1, \dots, y_k)$  and output symbols  $(x_1, \dots, x_{k+m})$ , observed through a non-stationary BSC with transition probability  $p$  over the first  $k$  components and 0 over the second  $m$  components. The initial reliabilities (absolute value of the initial LLRs) of  $(y_1, \dots, y_k)$ ,  $(x_1, \dots, x_k)$  and  $(x_{k+1}, \dots, x_{k+m})$  are 0,  $|\log((1-p)/p)|$  and  $+\infty$ , respectively.

Even though the matrix  $\mathbf{G}$  is sparse, its inverse is generally dense. This has two consequences: the intermediate symbols behave like random coin flips, and the computation of (4.3) is, in principle, quadratic in  $k$ . In order to compute (4.3) in linear time,  $\mathbf{G}$  should be such that the linear system  $\mathbf{G}(y_1, \dots, y_k)^T = (x_1, \dots, x_k)^T$  can be solved by direct elimination of one unknown at a time. Equivalently, the BEC message-passing decoding algorithm applied to the Tanner graph of the parity equation  $\mathbf{G}(y_1, \dots, y_k)^T + (x_1, \dots, x_k)^T = 0$  must terminate (i.e., it recovers all intermediate symbols  $(y_1, \dots, y_k)$ ).

**4.4. Compressor and Decompressor.** The basic scheme of Figure 1 is used where the syndrome computing block consists of the scheme presented in Section 4.3. The modeling algorithm is described in Section 6.

For simplicity we first describe the special case of the algorithm for binary data.

- (1) After block sorting the original  $k$ -data vector, we obtain a block of symbols  $(x_1, \dots, x_k)$ .
- (2) An intermediate block  $(y_1, y_2, \dots, y_k)$  of symbols is calculated by (4.3) (in practice, this is accomplished via message-passing).
- (3) A vector of  $m$  symbols  $(x_{k+1}, \dots, x_{k+m})$  is generated from  $(y_1, \dots, y_k)$  through encoding with an LT-code with parameters  $(k, \Omega(x))$ . Together with the doped bits generated as indicated below,  $(x_{k+1}, \dots, x_{k+m})$  forms the payload of the compressed data.
- (4) A bipartite graph (Figure 2) is set up between the nodes corresponding to  $(y_1, \dots, y_k)$ , and the nodes corresponding to  $(x_1, \dots, x_{k+m})$ . The edges in this graph are defined as follows: for all  $i$ , there is an edge from  $x_i$  to all the bits among  $(y_1, \dots, y_k)$  of which  $x_i$  is the addition.
- (5) The BP algorithm is applied to the graph created in the previous step. The objective of this BP algorithm is to decode the symbols  $(y_1, \dots, y_k)$  using the full knowledge of the symbols  $(x_{k+1}, \dots, x_{k+m})$  and the a priori marginal source probabilities  $\{P(x_i = 1) = p_i : i = 1, \dots, k\}$ . For  $i = 1, \dots, k$ , the LLR of the bit  $x_i$  (unavailable to the decoder) is set to  $\log((1-p_i)/p_i)$ . The marginal probabilities  $p_i$  are either known (in a non-universal setting) or estimated from the source sequence itself by the source modeler illustrated in Section 6. The nodes corresponding to  $(y_1, \dots, y_k)$  have initially zero reliability.
- (6) During the BP-algorithm, the CLID algorithm of [6, 2] is applied. In every  $\ell$ -th round of the iteration the intermediate bit  $y_{i_\ell}$  with the smallest reliability is marked, included in the payload, and its log-likelihood is set

---

<sup>3</sup>For example,  $\mathbf{G}$  can be constructed row-by-row, such that every row is sampled from  $\Omega(x)$  subject to the constraint that all the rows created so far are linearly independent.




FIGURE 2. Tanner graph for BP decoding.

to  $+\infty$  or  $-\infty$  depending on whether its value is 0 or 1. The combined BP-decoding with iterative doping is continued until the intermediate bits satisfy all the parity check equations of the LT code.

- (7) The payload output by the compressor consists of the bits  $(x_{k+1}, \dots, x_{k+m})$ , followed by the values of the doped intermediate bits.

The choice of the number of parity checks  $m$  is driven by the source entropy  $H(S)$ . For not too-small blocklength  $k$ , it can be observed that the CLID algorithm yields a small number of doped symbols to converge if  $m/k \geq H(S) + \delta$ , for some rate margin  $\delta > 0$  that generally depends on the LT code, on the source statistics and on the blocklength  $k$ , while it yields a very large number of doped symbols if the coding rate  $m/k$  is too close or above the source entropy. This is due to the fact that, as one may expect, the CLID algorithm is just a “perturbation” of the BP decoder to force it to recover the source sequence with probability 1, but is not a good encoding strategy in itself. Therefore, in order to have good compression performance we should choose  $m = k(H(S) + \delta)$ . The higher  $\delta$  is, the lower the number of required doped bits, and the higher the resilience against channel error and/or erasures. Thus a tradeoff of average vs variance in redundancy is possible by tuning the degree of backoff from entropy. If the entropy of the source is not known, we compute  $m$  based on the empirical entropy of the individual source sequence estimated by the universal source modeler.

The compressor performance, and in particular the output length variance, can be improved by trying several random realizations of the LT code for a given source sequence and choosing the one for which the number of doped symbols is smallest. Using LT codes makes this very simple in practice. In fact, any encoding matrix realization can be identified by a single integer number, seed of a pseudo-random number generator, that can be communicated to the decompressor at the cost of a constant (i.e., independent of the blocklength  $k$ ) number of additional bits.

The decompressor also proceeds in several steps which closely mimic the compression steps using the closed-loop iterative doping algorithm:

- (1) Using  $(x_{k+1}, \dots, x_{k+m})$  and the marked intermediate symbols  $y_{i_\ell}$ , all the intermediate bits  $(y_1, \dots, y_k)$  are reconstructed using a mirror image of the iterations of the BP algorithm used at the compressor.
- (2) The bits  $(x_1, \dots, x_k)$  are obtained applying the encoder for the raptor code to the intermediate bits  $(y_1, \dots, y_k)$ .
- (3) An inverse block sorting transform recovers the original data sequence.

In instances where the data compression algorithm is run in a channel with a low rate of erasures it is still possible to run the CLID algorithm using the Quenched Belief Propagation algorithm introduced in [16]. In general, when the rate of erasures is not low it is preferable not to use CLID and use instead a standard Raptor code for compression together with a BP decompressor that takes into account the probability of the data.

The choice of the output distribution  $\Omega(x)$  is crucial for the performance of the algorithm. The experiments reported in this paper are all based on the following degree distribution:

$$\begin{aligned} \Omega(x) = & 0.008x + 0.494x^2 + 0.166x^3 + 0.073x^4 + 0.083x^5 + \\ & 0.056x^8 + 0.037x^9 + 0.056x^{19} + 0.025x^{65} + 0.003x^{66}. \end{aligned}$$

This distribution was optimized for the binary erasure channel [12], but it produces surprisingly good results for binary symmetric channels [17, 18].

## 5. Nonbinary alphabets

Although the fountain-code based compression scheme applies immediately to non-binary alphabets, the direct application of the scheme given above to  $q$ -ary sources is difficult for the following reasons:

- (1) The design of  $q$ -ary LT codes for  $q$ -ary symmetric channels is not yet well-understood.
- (2) The  $q$ -ary BP decoder has complexity per source symbol  $O(q \log q)$ , i.e., superlinear in the alphabet cardinality.

Both problems are eliminated by the approach based on multilevel coding and multistage decoding introduced in [2]. Without loss of generality, we can assume that the source alphabet has cardinality  $q = 2^L$  for some integer  $L$  (the source alphabet can be always extended to a power of two by adding symbols with zero probability).

For the sake of simplicity, we outline the proposed multilevel approach to data compression for an i.i.d. source in the non-universal setting (i.e., assuming that the source statistics are perfectly known to both the encoder and decoder). This technique can be readily generalized to handle tree sources in a universal scheme by the concatenation with the BWT and by using the source modeler of Section 6.

Let  $X$  be an i.i.d. source with probability  $P_X$ , defined over the alphabet  $\mathcal{X}$  of cardinality  $|\mathcal{X}| = 2^L$ . For any one-to-one mapping  $B : \mathcal{X} \rightarrow \mathbb{F}_2^L$ , such that  $B(x) = (b_1(x), \dots, b_L(x))$ , we define the conditional marginal probability at level  $\ell = 1, \dots, L$  as

$$\begin{aligned} p_\ell(b_1, \dots, b_{\ell-1}) &= P(b_\ell(X) = 1 | b_1(X) = b_1, \dots, b_{\ell-1}(X) = b_{\ell-1}) \\ (5.1) \qquad \qquad \qquad &= \frac{\sum_{x \in \mathcal{X} : b_1(x) = b_1, \dots, b_{\ell-1}(x) = b_{\ell-1}, b_\ell(x) = 1} P_X(x)}{\sum_{x \in \mathcal{X} : b_1(x) = b_1, \dots, b_{\ell-1}(x) = b_{\ell-1}} P_X(x)} \end{aligned}$$

By the chain rule of entropy, we have that

$$\begin{aligned} H(X) &= \sum_{\ell=1}^L H(b_\ell(X) | b_1(X), \dots, b_{\ell-1}(X)) \\ (5.2) \qquad \qquad \qquad &= \sum_{\ell=1}^L \sum_{\substack{x \in \mathcal{X} : \\ b_1(x) = b_1, \\ \vdots \\ b_{\ell-1}(x) = b_{\ell-1}}} P_X(x) h(p_\ell(b_1, \dots, b_{\ell-1})) \end{aligned}$$

The source sequence  $(x_1, \dots, x_k)$  can be represented under the mapping  $B$  by the  $L$  binary sequences  $\{(b_\ell(x_1), \dots, b_\ell(x_k))\}_{\ell=1}^L$ . Then, we can apply the binary compressor outlined in Section 4.4 to each binary sequence successively. At level  $\ell = 1$ , we treat  $b_1(x_1, \dots, x_k)$  as an i.i.d. Bernoulli sequence with parameter  $p_1$  computed by (5.1). At level  $\ell = 2$ , we treat  $b_2(x_1, \dots, x_k)$  as an independent but non-stationary Bernoulli sequence with parameter  $p_2(0)$  for all positions  $i$  such that  $b_1(x_i) = 0$  and parameter  $p_2(1)$  for all positions  $i$  such that  $b_1(x_i) = 1$ . The encoding proceeds in this way up to level  $L$ .

At the decoder, decompression is achieved level by level in the same order as the encoder. Therefore, when the decompressor is applied to the level  $\ell$  sequence all the sequences at lower levels 1 through  $\ell - 1$  have been already recovered and the conditional prior probabilities needed by the BP decoder at level  $\ell$  are known.

Remarkably, we noticed that even for very non-uniform  $q$ -ary distributions  $P_X$ , it is often the case that for some levels  $\ell$  and bit values  $b_1, \dots, b_{\ell-1}$  the conditional probability  $p_\ell(b_1, \dots, b_{\ell-1})$  is quite close to  $1/2$ . In this case, and as explained in [6], it is convenient to pass these binary symbols directly to the output and not include them in the compression scheme of Section 4.4. This allows both complexity reduction and increase in efficiency since it shortens the input fed to the binary compression scheme and prevents the slow convergence in the BP decoder that a large fraction of very small reliability symbols can create.

## 6. Modelling and Universality

In this section we consider a universal version of our data compression scheme based on modelling the source statistical dependencies as a tree.

Suppose that the source sequence  $x = (x_1, \dots, x_k)$  takes values on a given  $q$ -ary alphabet  $\mathcal{A}$  and is generated by a stationary ergodic tree source with  $S$  states. Then, the output of the BWT, denoted in the following by  $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_k)$  converges, as  $k \rightarrow \infty$ , to a piecewise i.i.d. sequence [19, 20]. This means that there exist  $S' + 1$  *transition points*, denoted by  $1 = t_1 < t_2 < \dots < t_{S'} < t_{S'+1} = k + 1$  that partition  $\tilde{x}$  into segments  $(\tilde{x}_{t_j}, \dots, \tilde{x}_{t_{j+1}-1})$  for  $j = 1, \dots, S'$ . The empirical

*marginal* probability distribution of symbols in segment  $j$  is given by  $\widehat{P}_j(a) = N_j(a)/(t_{j+1} - t_j)$ , where we define the segment symbol counts

$$(6.1) \quad N_j(a) = \sum_{i=t_j}^{t_{j+1}-1} 1\{x_i = a\}$$

The number of segments  $S'$  may differ from the number of states  $S$  of the tree source but, for sufficiently large  $k$ , we have that  $S' = S$  and that  $\widehat{P}_j(a)$  is close to the true tree source probability distribution conditioned on state  $j$ .

Our approach consists of modelling the source tree structure and estimating its state probability by processing the BWT output sequence, along the ideas presented in [21, 22]. Since the compressor acts on  $\tilde{x}$  by treating it as a piecewise i.i.d. sequence with given marginal probabilities, our goal is to find the most efficient piecewise i.i.d. description of  $\tilde{x}$ . Generally speaking, a source statistics model  $\mathcal{M}$  is given by the number of states (or segments)  $S'$ , by the distinct transition points  $(t_2, \dots, t_{S'})$  and by the model segment distributions  $\{Q_j(a) : j = 1, \dots, S'\}$ . The cost of a model  $\mathcal{M}$  to represent  $\tilde{x}$  is measured by the total number of bits needed to describe  $\tilde{x}$ , given by

$$(6.2) \quad c(\tilde{x}, \mathcal{M}) = S'(\log_2 k + (q-1)b) + \sum_{j=1}^{S'} \sum_{a \in \mathcal{A}} N_j(a) \log_2 \frac{1}{Q_j(a)}$$

where we spend  $\log_2 k$  bits to encode the transition points<sup>4</sup>,  $(q-1)b$  bits to encode each nominal segment distribution (see later for details), and

$$\sum_{a \in \mathcal{A}} N_j(a) \log_2 \frac{1}{Q_j(a)}$$

bits to encode the  $j$ th segment symbols. Notice that this length is an *estimate* of the output length of a Shannon code for encoding the symbols in segment  $j$  using the model distribution  $Q_j(a)$ . We take this as an estimate of the cost incurred by the fountain code compressor.

In order to find the most efficient piecewise i.i.d. source model, we follow the segment merging procedure explained in [22] with a different segment cost function. The BWT output and the original source sequence are related by a data-dependent permutation  $\pi$ , such that  $x_{\pi(i)} = \tilde{x}_i$ . Hence, the depth- $d$  context of each symbol  $\tilde{x}_i$  is obtained as  $(x_{\pi(i)-d}, \dots, x_{\pi(i)-1})$ . By exploiting this fact, it is possible to partition  $\tilde{x}$  into segments of symbols with common context for a certain maximum depth  $d_{\max}$ , that is a design parameter of the algorithm. We identify segments by their context. Hence, the depth- $d_{\max}$  segments are arranged as leaves of a  $q$ -ary tree where the root is the whole sequence  $\tilde{x}$  (segment of depth 0) and where, for  $0 < d < d_{\max}$ , the segment with context  $\mathbf{s}_1^d = (s_d, \dots, s_1)$  has at most  $q$  children segments with contexts  $(a, \mathbf{s}_1^d)$ , for  $a \in \mathcal{A}$ .

Let the empirical *marginal* probability distribution of symbols in segment  $\mathbf{s}_1^d$  be denoted by  $\widehat{P}(a|\mathbf{s}_1^d) = N(a|\mathbf{s}_1^d)/L(\mathbf{s}_1^d)$ , where we define the segment symbol counts

$$(6.3) \quad N(a|\mathbf{s}_1^d) = \sum_{i \in \text{segment } \mathbf{s}_1^d} 1\{x_i = a\}$$

---

<sup>4</sup>We have  $S' - 1$  transition points plus  $\log_2 k$  bits to encode the *BWT index*, necessary to perform inverse BWT at the decompressor.

and where  $L(\mathbf{s}_1^d)$  is the segment length. The cost of directly encoding segment  $\mathbf{s}_1^d$  is given by

$$(6.4) \quad c(\mathbf{s}_1^d) = \log_2 k + (q-1)b + \sum_{a \in \mathcal{A}} N(a|\mathbf{s}_1^d) \log_2 \frac{1}{Q(a|\mathbf{s}_1^d)}$$

where  $Q(a|\mathbf{s}_1^d)$  is a quantized version of  $\hat{P}(a|\mathbf{s}_1^d)$  using  $(q-1)b$  bits. The segment merging algorithm is initialized by associating to each depth- $d_{\max}$  segment its cost  $c(\mathbf{s}_1^{d_{\max}})$ . Then, for depth  $d = d_{\max} - 1$  to  $d = 0$ , the cost associated to segment  $\mathbf{s}_1^d$  is given by the minimum between the sum of the costs of its children segments and the cost of representing it directly. The children segments are merged and their corresponding branches in the segment tree are pruned if

$$(6.5) \quad \sum_{a \in \mathcal{A}} c(a, \mathbf{s}_1^d) > c(\mathbf{s}_1^d).$$

Otherwise, the branches are kept in the tree.

The algorithm terminates when it is not possible to prune the tree further. The leaves of the pruned tree correspond to the segments of the optimal piecewise i.i.d. model  $\mathcal{M}$  for  $\hat{x}$  (subject to the tree source assumption of the original source).

The empirical probabilities on each segment of the model are quantized as follows. First, the multilevel conditional marginals defined in (5.1) are computed for each segment from the true empirical probabilities  $\hat{P}(a|\mathbf{s}_1^d)$ . Then, each value of the conditional marginal probabilities  $p = p_\ell(b_1, \dots, b_{\ell-1})$  in (5.1) is represented by  $b$  bits, by quantizing  $\log((1-p)/p)$ . This requires a total of  $(q-1)b$  bits per segment.

It is worth noticing that Rissanen's MDL bound [23] on the redundancy for sources with  $S$  states and  $q-1$  real parameters per state is given by  $\frac{S(q-1)}{2k} \log_2 k$  bits per source symbol. On the other hand, for piecewise i.i.d. sources with  $S$  segments, arbitrary unknown transition points, and  $q-1$  real parameters per segment, Merhav's bound [24] yields redundancy  $\frac{S(q+1)}{2k} \log_2 k$  bits per source symbol. By choosing  $b = \frac{1}{2} \log_2 k$  bits to quantize the empirical probabilities, the proposed source modeler for sources with  $S$  fixed and sufficiently large  $k$  achieves Merhav's bound. Remarkably, the relative overhead incurred by approximating the BWT output as a piecewise i.i.d. sequence with arbitrary transition points decreases as the source alphabet increases.

As envisioned in [6] (and implemented in state-of-the-art BWT-based compressors) it is convenient to pre-process the BWT output with the move-to-front algorithm [25] when the source has a large number of states  $S$  relatively to the blocklength  $k$ , as is often the case in practice. The move-to-front transformation replaces each symbol  $\tilde{x}_i = a$  with the number of distinct symbols appeared in  $\tilde{x}$  since the last appearance of  $a$ . Since the symbols at the input of the move-to-front algorithm are independent (or weakly dependent) the most probable symbol in the transformed sequence is 0, the next most probable is equal to 1 and so on. The beneficial effect of move-to-front on the segment merging algorithm stems from the fact that after move-to-front the empirical distributions of the segments tend to be more similar since move-to-front implicitly implements a symbol permutation that arranges probabilities in decreasing order. Therefore, the merging algorithm is more likely to merge segments after the move-to-front operation. If the number of states is small relative to the blocklength, then the segments are long enough and it

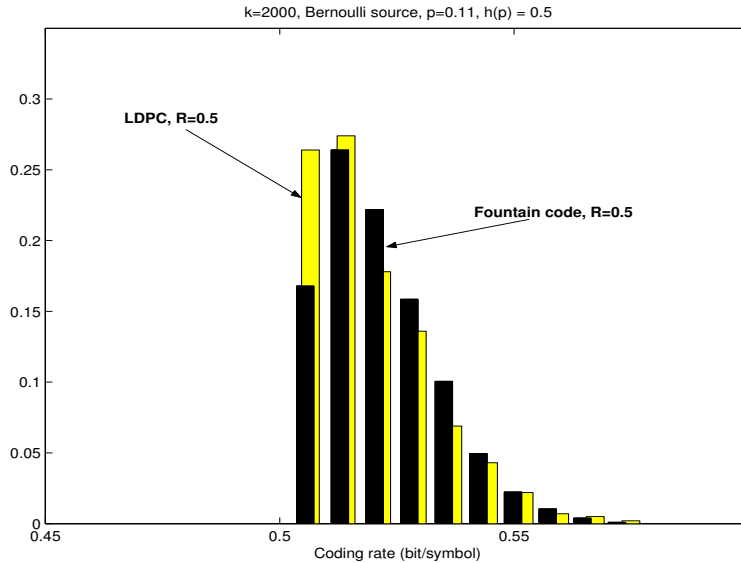


FIGURE 3. Biased coin  $p = 0.11$ ; blocklength 2000; LDPC-based compression vs. Fountain code-based compression.

is preferable not to merge segments even if their distributions are similar. Thus in those cases, move-to-front may actually incur in a small performance degradation.

## 7. Experiments

**7.1. Biased Coin.** In this subsection we test in the simplest possible nonuniversal setting the algorithm of Section 4.4. We consider an independent identically distributed binary source with bias probability  $p$ . Figure 3 shows the histogram encoding rates (compressed output length)/(source length) with  $p = 0.11$  and source length equal to 2000 bits for both the LDPC-CLID based approach of [6] and the fountain-based approach of Section 4.4. In both cases, the codes assume a nominal rate of 0.5 (equal to the entropy of the source). This translates into a choice of  $k = 2000$  and  $m = 1000$  in the encoder of Section 4.4. Note that a more efficient approach than the ones shown in Figure 3 is possible by a more aggressive choice of the rate as a function of the number of ones in the sequence. In this way, it is possible to closely track the information spectrum of the source (distribution of the log-probabilities) or, equivalently, the performance of arithmetic coding. From Figure 3 it can be seen that the performance of the compressor of [6] and the compressor in this paper are very similar.

The comparable performances of the fountain-based schemes in Sections 4.4 and 4.1 is shown in Figure 4, which plots the absolute redundancy (compression length - empirical entropy)/(source length) for a biased coin with bias  $p = 0.105$  and length equal to 5000 bits. The beneficial effect of the CLID based approach Section 4.4 is apparent from this figure. However, as we pointed out in Section 4.1, the approach therein is a strong contender in noisy-channel applications.

**7.2. Binary Sources with Memory.** In order to compare the universal version of our scheme with the leading data compression methods such as `gzip`, `bzip`

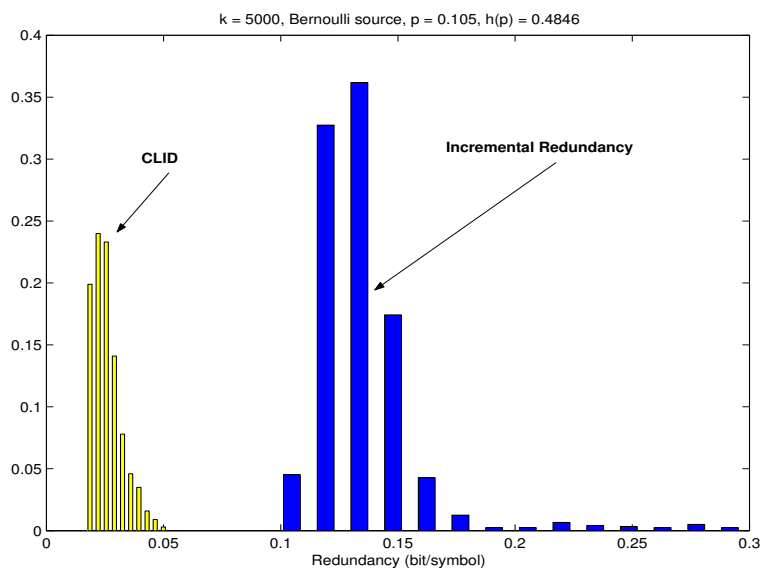


FIGURE 4. Biased coin  $p = 0.105$ ; blocklength 5000; Effect of Closed-Loop Iterative Doping.




FIGURE 5. Nonergodic binary source; 10,000 bits

and PPM we will use synthetic Markov sources whose entropy is easily computable so that we can gauge how far the various methods are from the Shannon limit. We are particularly interested in the regime of moderate length. Not only the gaps from the Shannon limit would vanish (in terms of difference between rate and entropy) for long lengths, but as we mentioned above the new proposed methods are

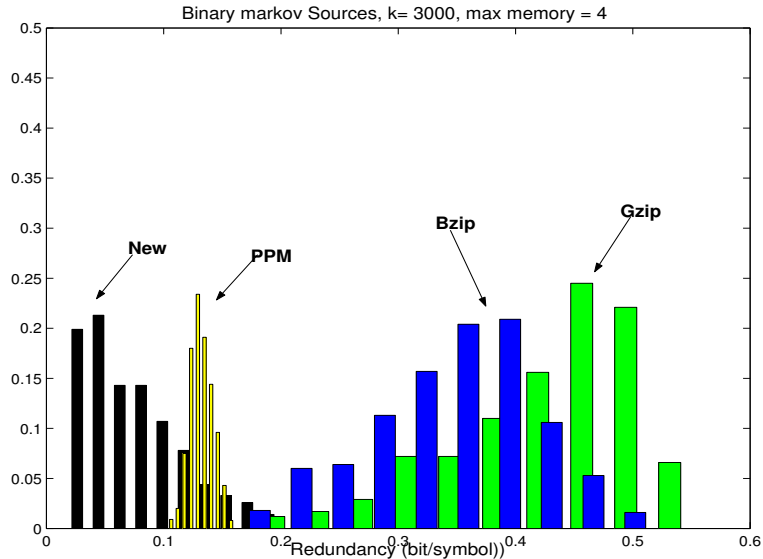


FIGURE 6. Nonergodic binary source; 3,000 bits

particularly useful for adaptation to source-channel schemes in data transmission applications where relatively short data packets are of interest.

Instead of experimenting with a given Markov source, we generate an ensemble of binary Markov sources whose transition probabilities are chosen at random. The number of states is equally likely to be 1, 2, 4, 8, 16, 32, and 64. A nonergodic source ensemble is obtained by generating independently the memory length (binary logarithm of the number of states), and then conditional distributions are also generated randomly producing sources with entropy ranging from 0.05 to 0.75 bit/symbol. The same ensemble was used to test the LDPC-based scheme in [8]. Figure 5 shows a histogram of the absolute redundancy achieved by the proposed fountain-code based method in Section 4.4 and by PPM, *gzip* (Lempel-Ziv) and *bzip* (BWT-based) in the case that the source realization has 10,000 bits. None of the four methods have any prior knowledge about the source. We can see a clear advantage with respect to *gzip* and *bzip* in both variability and average redundancy, while the comparison with respect to PPM is rather competitive.

In Figure 6 a random Markov ensemble of memory length up to 4 is tested with source realizations containing 3,000 bits. All four methods suffer degradation due to the shorter blocklength (notice the  $x$ -axis scale is wider in Figure 6) but the competitive advantage of the new method is enhanced. For the sake of clarity in the figures, we do not include the results of the LDPC-based codes in [8] using the same universal modeler as in the fountain-code based algorithms. While the new algorithm offers a slight performance degradation in terms of compression efficiency with respect to the LDPC-based method, the fountain-code method is much easier to implement in a universal setting as it avoids having to store a library of block codes with different rates. Thus, it can be viewed as providing a natural way to puncture a single code.

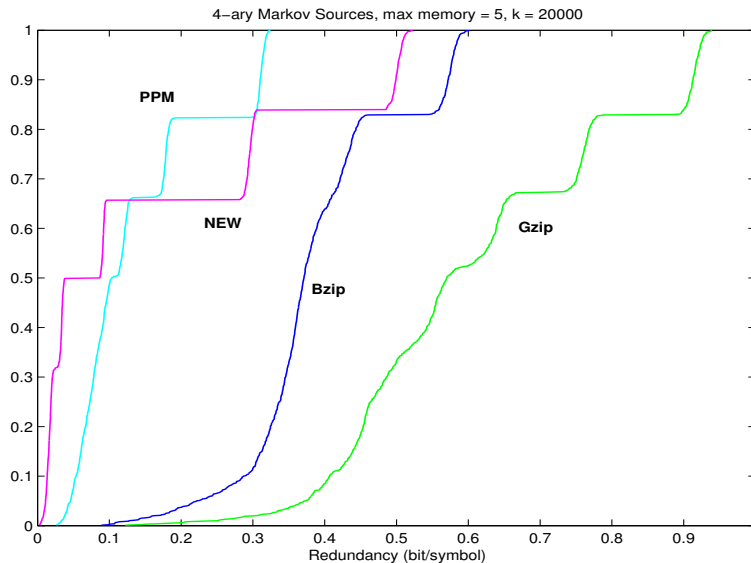


FIGURE 7. Nonergodic 4-ary source; 20,000 symbols

**7.3. Nonbinary Sources with Memory.** The universal multilevel approach of Section 5 is tested for two nonergodic sources with 4-letter and 8-letter alphabets in Figures 7 and 8 respectively. The sources are constructed as in Section 7.2 by generating Markov sources with random memory and random transition probabilities. In the experiment of the 4-letter source the maximum memory length is 5, and the conditional entropy given any state ranges from 0 to  $\log_2 3$ . In the experiment of the 8-letter source the maximum memory length is 4, and the conditional entropy given any state ranges from 0 to  $\log_2 5$ .

For clarity, we plot the cdfs of the absolute redundancy in Figures 7 and 8. The horizontal steps in these figures correspond to the different behavior observed depending on the memory length. In these examples with longer blocklengths, the comparison is more favorable to PPM than before. However, the quadratic computational complexity of PPM is considerably higher than that of the proposed algorithm in these cases.

## 8. Conclusion

Following the recent approach to universal data compression with linear error correcting codes and belief-propagation decoding [6, 7, 8], we have proposed the replacement of the LDPCs therein by fountain codes.

For applications involving data compression without protection against possible channel errors/erasures, we have seen the effectiveness of a two-stage approach that generates intermediate bits, which are then encoded by a fountain code used in conjunction with the Closed-Loop Iterative Doping algorithm.

The gains in compression efficiency over the scheme in [8] are due to a more efficient universal modeler of the source, which can also be adopted in the LDPC approach. It is interesting to note that although the experiments shown in [6, 7] used LDPCs ensembles that were optimized for binary symmetric channels using

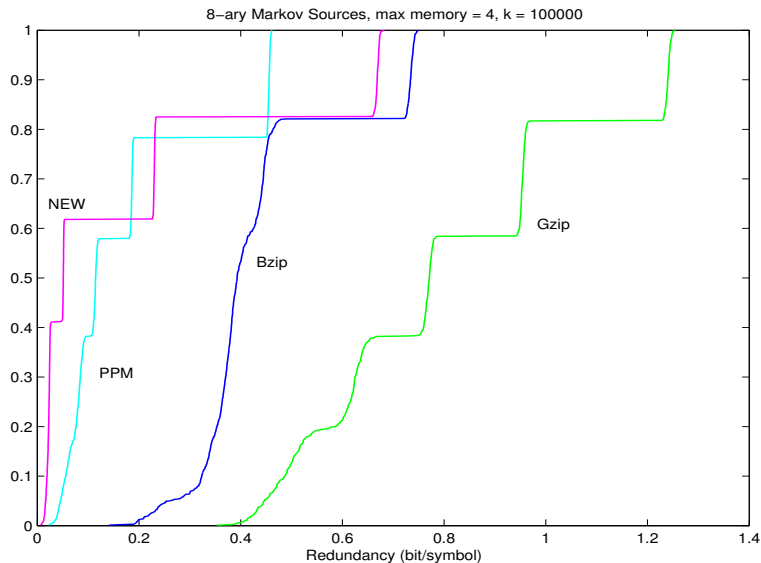


FIGURE 8. Nonergodic 8-ary source; 100,000 symbols

density evolution, in general, the LDPC approach does not result in appreciably better efficiency than the method used in this paper which uses a degree distribution originally chosen for the binary erasure channel in [10].

Compared to the LDPC approach, the method of this paper based on fountain codes offers a streamlined design where a library of codes of different rates need not be maintained. The resulting codes have linear encoding and decoding complexity, while their performance is competitive with PPM which requires quadratic complexity.

This paper has focused on the pure data compression problem. However, the major gains of linear channel codes in source coding are to be had in competing against the standard paradigm of the source-channel separation principle, through the design of effective joint source-channel coding techniques.

### References

- [1] I. Csiszar and J. Korner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*, Academic, New York, 1981.
- [2] G. Caire, S. Shamai, and S. Verdú, "Noiseless data compression with low density parity check codes," in *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, P. Gupta and G. Kramer, Eds., pp. vol. 66, pp. 263–284. American Mathematical Society, 2004.
- [3] P. E. Allard and A. W. Bridgewater, "A source encoding technique using algebraic codes," *Proc. 1972 Canadian Computer Conference*, pp. 201–213, June 1972.
- [4] H. Ohnsorge, "Data compression system for the transmission of digitalized signals," *Conf Rec IEEE Int. Conf. on Communications*, vol. II, pp. 485–488, June 1973.
- [5] T. Ancheta, "Syndrome source coding and its universal generalization," *IEEE Trans. Information Theory*, vol. 22, no. 4, pp. 432 – 436, July 1976.
- [6] G. Caire, S. Shamai, and S. Verdú, "A new data compression algorithm for sources with memory based on error correcting codes," *2003 IEEE Workshop on Information Theory*, pp. 291–295, Mar. 30- Apr. 4, 2003.

- [7] G. Caire, S. Shamai, and S. Verdú, “Lossless data compression with error correction codes,” *2003 IEEE Int. Symp. on Information Theory*, p. 22, June 29– July 4, 2003.
- [8] G. Caire, S. Shamai, and S. Verdú, “Universal data compression with LDPC codes,” *Third International Symposium On Turbo Codes and Related Topics*, pp. 55–58, Brest, France, September 1-5, 2003.
- [9] M. Burrows and D. J. Wheeler, “A block-sorting lossless data compression algorithm,” Tech. Rep. SRC 124, May 1994.
- [10] M. G. Luby, “LT codes,” *Proc. 43rd IEEE Symp. Foundations of Computer Science*, pp. 271–280, 2002.
- [11] D. MacKay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, Cambridge, UK, 2003.
- [12] A. Shokrollahi, “Raptor codes,” Preprint, 2003.
- [13] M. G. Luby, M. Mitzenmacher, M. A. Shokrollahi, and D. A. Spielman, “Efficient erasure correcting codes,” *IEEE Trans. Information Theory*, vol. 47, pp. 569–584, Feb. 2001.
- [14] E. Ordentlich, G. Seroussi, S. Verdú, K. Viswanathan and M. Weinberger, and T. Weissman, “Channel decoding of systematically encoded unknown redundant sources,” *2004 Proc. IEEE Symposium on Information Theory*, p. 165, 2004.
- [15] J. Hagenauer, A. Barros, and A. Schaefer, “Lossless Turbo source coding with decremental redundancy,” *Proc. Fifth Int. ITG Conference on Source and Channel Coding*, pp. 333–339, Jan. 2004.
- [16] G. Caire, S. Shamai, and S. Verdú, “Almost-noiseless joint source-channel coding-decoding of sources with memory,” *Proc. Fifth International ITG Conference on Source and Channel Coding (SCC)*, pp. 295–304, Jan 14-16, 2004.
- [17] O. Etesami, M. Molkaiaie, and A. Shokrollahi, “Raptor codes on symmetric channels,” *Proc. 2004 IEEE Int. Symp. Information Theory*, p. 38, June 2004.
- [18] R. Palanki and J. Yedidia, “Rateless codes on noiseless channels,” *Proc. 2004 IEEE International Symposium on Information Theory*, p. 37, June 2004.
- [19] K. Visweswariah, S. Kulkarni, and S. Verdú, “Output distribution of the Burrows-Wheeler transform,” *Proc. 2000 IEEE International Symposium on Information Theory*, p. 53, June 27– July 1, 2000.
- [20] M. Effros, K. Visweswariah, S. Kulkarni, and S. Verdú, “Universal lossless source coding with the Burrows-Wheeler transform,” *IEEE Trans. on Information Theory*, vol. 48, no. 5, pp. 1061–1081, May 2002.
- [21] N. J. Larsson, “The context trees of block sorting compression,” *Proc. 1998 Data Compression Conference*, pp. 189–198, Mar. 1998.
- [22] D. Baron and Y. Bresler, “An  $O(N)$  semi-predictive universal encoder via the BWT,” *IEEE Trans. Information Theory*, pp. 928–937, May 2004.
- [23] A. Barron, J. Rissanen, and B. Yu, “The minimum description length principle in coding and modeling,” *IEEE Trans. on Information Theory*, vol. 44, pp. 2743–2760, Oct. 1998.
- [24] N. Merhav, “On the minimum description length principle for sources with piecewise constant parameters,” *IEEE Trans. Information Theory*, vol. 39, no. 6, pp. 1962–1967, Nov. 1993.
- [25] P. Elias, “Interval and recency rank source coding: two on-line adaptive variable-length schemes,” *IEEE Trans. Information Theory*, vol. 33, pp. 3–10, Jan. 1987.

INSTITUT EURECOM, SOPHIA ANTIPOLIS, FRANCE  
*E-mail address:* `giuseppe.caire@eurecom.fr`

TECHNION, HAIFA, ISRAEL  
*E-mail address:* `sshlobo@ee.technion.ac.il`

EPFL, LAUSANNE, SWITZERLAND  
*E-mail address:* `amin.shokrollahi@epfl.ch`

PRINCETON UNIVERSITY, NEW JERSEY, USA  
*E-mail address:* `verdu@princeton.edu`