

Dv2v: A Dynamic Variable-to-Variable Compressor

Nieves R. Brisaboa*, Antonio Fariña*, Adrián Gómez-Brandón*,
Gonzalo Navarro[†], and Tirso V. Rodeiro*

*Universidade da Coruña
Facultade de Informática, CITIC
A Coruña, Spain
brisaboa@udc.es, fari@udc.es,
adrian.gbrandon@udc.es,
tirso.varela.rodeiro@udc.es

[†]University of Chile
Millennium Institute for Foundational
Research on Data (IMFD),
Department of Computer Science
Santiago, Chile
gnavarro@dcc.uchile.cl

Abstract

We present *D-v2v*, a new dynamic (one-pass) variable-to-variable compressor. Variable-to-variable compression aims at using a modeler that gathers variable-length input symbols and a variable-length statistical coder that assigns shorter codewords to the more frequent symbols. In *D-v2v*, we process the input text word-wise to gather variable-length symbols that can be either terminals (new words) or non-terminals, subsequences of words seen before in the input text. Those input symbols are set in a vocabulary that is kept sorted by frequency. Therefore, those symbols can be easily encoded with dense codes. Our *D-v2v* permits real-time transmission of data, i.e. compression/transmission can begin as soon as data become available. Our experiments show that *D-v2v* is able to overcome the compression ratios of the v2vDC, the state-of-the-art semi-static variable-to-variable compressor, and to almost reach *p7zip* values. It also draws a competitive performance at both compression and decompression.

1 Introduction

Text compression has gained relevance in the last decades along with the growth of text databases. It permitted not only to drastically reduce the storage needs of those data and the time needed to transmit them through a network, but also to handle them efficiently in compressed form.

The first compressors based on Huffman coding [1] using character-oriented modeling obtained rather poor compression ratios on text collections (compression around 60%). However, when Huffman coding was coupled with a word-based modeler during the 80s [2] the compression ratio obtained by those semi-static compressors was close to 25% when applied to English texts, and they set the basis to build modern text retrieval systems over them [3]. This boosted the interest of new compressors not only yielding fast decoding/retrieval but also allowing queries to be performed

[§] Funded in part by European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 690941 (project BIRDS). G.N. is funded by the Millennium Institute for Foundational Research on Data (IMFD), Chile. The Spanish group is funded by Xunta de Galicia/FEDER-UE [CSI: ED431G/01 and GRC: ED431C 2017/58]; by MINECO-AEI/FEDER-UE [ETOME-RDFD3: TIN2015-69951-R, Datos 4.0: TIN2016-78011-C4-1-R]; by MINECO-CDTI/FEDER-UE [INNTERCONECTA: uForest ITC-20161074]; by FPI Program [BES-2017-081390] (T.V.R); and by FPU Program [FPU16/02914] (A.G.B.).

in compressed form. At the end of the 90s, *Plain Huffman (PH)* and *Tagged Huffman (TH)* [4, 5] replaced the bit-oriented Huffman by byte-oriented Huffman to speed up decoding at the cost of losing compression effectiveness (now around 30%). In addition, *TH* reserved the first bit of each byte to gain synchronization capabilities. Compression ratios worsened to around 34% but random decompression and fast Boyer-Moore type searches became possible. In the same line, the use of *dense codes* [6] allowed *End-Tagged Dense Code (etdc)* and *(s,c) Dense Code (scdc)* to not only retain the same capabilities of *TH* but also improve its compression ratios, which became very close to those of *PH*, and a simpler coding scheme that does not depend on the Huffman tree. Indeed, assuming we have n source symbols s_i ($0 \leq i < n$) with decreasing probabilities, the codeword c_i corresponding to the i -th symbol can be obtained as $c_i \leftarrow \text{encode}(i)$, and the rank i corresponding to c_i can be obtained as $i \leftarrow \text{decode}(c_i)$. Both *encode* and *decode* algorithms perform in $O(|c_i|)$ time [7].

Unfortunately, since *PH* is the optimal word-based 256-ary zero-order compressor, all those efficient and searchable word-based compression techniques, could never reach the compression of the strongest compressors (e.g. *p7zip*). This motivated the creation of *v2vdc*, the first word-based *variable-to-variable*¹ compressor [8]. *v2vdc* obtained similar compression ratios to those of *p7zip* in English texts by parsing the text into both words and phrases (sequences of words) and then assigning codewords to them by using dense codes. Finally, the original words/phrases were replaced to create the compressed file. *v2vdc* improved the compression ratio of *etdc* by around 8–10 percentage points and besides it produced a still searchable compressed text.

Data transmission is another scenario where compression is of special interest. In some cases, the whole data is available and can be compressed with the most powerful compressors. However, there are scenarios where dynamic (one-pass) real time compression becomes necessary. That is, the compressor/sender must be capable of compressing the symbols on the fly as they arrive without the need of having the whole data before starting its compression/transmission. This could be the case of sensor data transmitted to a server, a digital library streaming a book to a electronic reader one page at a time, HTTP pages sent by a server during a HTTP session, etc. Even there exists powerful one-pass adaptive compressors such as those coupling arithmetic coding with k -order PPM-modeling [9], or those derived from the Lempel-Ziv family [10, 11], they do not match real-time requirements. Yet, we can find in the literature versions of dynamic character-based Huffman compressors [12, 13] and also word-based compressors such as the *Dynamic PH (dph)* [14], or the *Dynamic etdc (detdc)* [7, 14]. The later takes advantage of the simple on-the-fly *encode* and *decode* algorithms from *etdc* and permits both sender/compressor and receiver/decompressor to remain synchronized by simply keeping the same vocabulary of words sorted by frequency. Basically, assuming that, at a given moment, the vocabulary of the sender contains n words, when the sender inputs the next word w it could find it in its vocabulary at position i , so it simply sends $c_i \leftarrow \text{encode}(i)$ to the receiver. Otherwise, if w is a new word, it sends $c_n \leftarrow \text{encode}(n)$ (used as an escape codeword) followed by

¹Variable-to-variable compression aims at using a modeler that gathers variable-length symbols and a variable-length statistical coder that assigns shorter codewords to the more frequent symbols.

w in plain form. In any case, the encoder increases the frequency counter f of w to $f + 1$ and runs a simple *update* algorithm that swaps w with the first word that has frequency equal to f . This *update* algorithm keeps the vocabulary of words sorted by frequency and runs in $O(1)$ time. The receiver is also very simple. It receives a codeword c_i and runs $i \leftarrow \text{decode}(c_i)$. Then, if $i < n$ it has decoded the word w_i at the i -th entry of the vocabulary. Otherwise, if $i = n$ it receives a new word in plain form and adds it at the end of the vocabulary. Finally, a similar *update* procedure to that of the sender is run to increase the frequency of that word and to keep the vocabulary sorted. A variant based on *scdc* (*dscdc*) is also available [7]. Finally, more recent *lightweight* versions of *detdc* and *dscdc*, using asymmetric compression/decompression procedures to reduce the work done at decompression, were also created [15]. *detdc* displayed similar compression ratios to those of the semi-static *etdc* (around 33%) while yielding fast compression and decompression. Yet, as for *etdc*, its compression effectiveness is far from the stronger variable-to-variable counterpart *v2vdc*.

In this work, we create a dynamic (one-pass) variable-to-variable variant of *detdc* named *D-v2v*. We follow the same ideas from *detdc* to keep both sender and decoder synchronized. The variable-length symbols in our vocabulary will be arbitrary length sequences of words that can be either single-words (terminals) or repeating pairs of symbols (that again can be terminals or not) that have occurred previously in the input sequence. We use a sort of Patricia tree [16] to efficiently handle the subsequences seen before in the input sequence. Note that our ability to choose *good/relevant* variable-length sequence of words will determine the success of our new compressor. However, finding the smallest grammar for a text is a NP-complete problem [17]. To overcome this, several heuristics exist: *lz78* [11] looks for existing substrings in the already processed sequence; *repair* replaces pairs of repeating symbols recursively [18]; *Sequitur* [19] replaces a pair of repeating phrases in the processed sequence by a new phrase; or others such as [20] where the grammar induced ensures that non-terminals do only contain terminals. In *D-v2v*, our strategy to gather input symbols representing variable-length sequence of words is similar to that of *Sequitur*, whereas the semi-static *v2vdc* followed the approach from [20] supported by a *suffix array* [21] of the text and the corresponding *longest common prefix* structure. We will show that our simpler procedure also yields competitive compression values.

2 Our proposal: *D-v2v*

As we have explained, *D-v2v* is a dynamic (one-pass) compressor. *D-v2v* processes the input text and gathers symbols that represent sequences with a variable number of words. We use a sort of trie to help the parser to detect sequences of words that appeared before. We keep those symbols sorted by frequency. In this way, we can use the *etdc* encoder to encode them directly from their positions in the vocabulary. The decompressor/receiver is simpler because it only has to decode the received codewords and to keep the table of symbols sorted by frequency (synchronized with the sender).

In the next sections we conceptually describe the *parser* and the *encoder* procedures of the sender/compressor component, and also the *decoder* procedure that is




Figure 1: Non-terminal creation example.

the core of the receiver/decompressor component.

Parsing algorithm used by the sender

Our parser scans the text and splits it into tokens/symbols of one or more words that can be:

- *terminal symbols.* Those representing just one word. They are created when a new word is parsed.
- *non-terminal symbols.* Those composed by two different symbols, which can be terminals or non-terminals. Therefore, each non-terminal, represents at least a sequence of two words.

During the parsing, the sender reads the text one word at a time. If the next read word was not in our vocabulary, two symbols are created: *i*) a terminal symbol S_{new} which represents the new word (the sender will notify the receiver about this new word as we will show in the next section); and *ii*) a non-terminal $S_{\text{new}+1}$ which appends S_{new} to the previous sent symbol. For example, in Figure 1, after sending S_2 we read the new word w_{new} . Therefore, a terminal symbol S_{new} is created for that word and then a non-terminal symbol $S_{\text{new}+1}$ is created for $S_2||S_{\text{new}}$.

Otherwise, if the next read word is a prefix of any symbol from the vocabulary, we store such word in RS (read sequence). We keep reading the text word by word and append those words to RS until RS becomes an unknown sequence. At this moment, we send the symbol which corresponds to the longest known prefix of RS . Then, a new non-terminal symbol containing the current sent symbol and the previous one is created. In the example of Figure 1, let us assume that S_8 is a non-terminal that expands into the sequence $w_1w_2w_3$, and S_{10} produces $w_1w_2w_3w_4w_5w_6$, where S_{10} is the unique non-terminal symbol starting by $w_1w_2w_3w_4w_5$. After sending the symbol S_5 , we are at w_1 and we read the next words $w_1 \dots w_5$ one word at a time. We keep reading words until we reach w_7 . At that moment, $RS \leftarrow w_1w_2w_3w_4w_5w_7$ is not a prefix of the sequence in S_{10} and we stop processing the text. Note that, since the symbol containing the longest known prefix of RS corresponds to S_8 , we send S_8 (the way to encode S_8 will be explained in detail in the next section) and we create a new non-terminal symbol S_{new} for $S_5||S_8$. We will continue parsing from w_4 on.

In practice, we are using a set of known sequences KS which stores every previously created terminal and non-terminal symbols. If we are sending the message “the more I know about you the more I know about me”, at the beginning we have $KS = \emptyset$ and we read the first word “the”. Since KS is empty, there is no sequence which starts with “the”, thus we add it to KS , at position $i = |KS| = 0$, the symbol

$S_i = S_0 = \text{“the”}$, and we send S_0 to the receiver. Then we read “more” , which is also a new word. Now we have to add both the new terminal symbol $S_1 = \text{“more”}$ (S_1 is also sent to the receiver) and the new non-terminal symbol $S_2 = \text{“the more”}$ to KS .

After processing the word “you” , KS is composed by $\{S_0:\text{“the”}, S_1:\text{“more”}, S_2:\text{“the more”}, S_3:\text{“I”}, S_4:\text{“more I”}, S_5:\text{“know”}, S_6:\text{“I know”}, S_7:\text{“about”}, S_8:\text{“know about”}, S_9:\text{“you”}, S_{10}:\text{“about you”}\}$ and we continue reading “the” . Since the current read sequence $RS = \text{“the”}$ exists in KS , we read the next word and append it to RS . Now, $RS = \text{“the more”}$ matches the symbol S_2 stored in KS . In the next step, we update RS to “the more I” . Since that sequence is not included in KS , we send S_2 to the receiver and we create a new non-terminal that includes the previous and the current sent symbol: $S_{11} = S_9 || S_2 = \text{“you the more”}$.

We need a mechanism to check if RS is within the set of known sequences KS and to obtain its symbol identifier, i.e. its rank in KS . In order to perform those tasks efficiently we use a structure based on the Patricia-tree, where each branch represents a sequence, and all the sequences that start with the same prefix descend from the same node.² The last property is important, as it allows us to search incrementally for the longest sequence contained in RS . For example, after reading the second “the” ($RS = \text{“the”}$) we access to the trie of Figure 2 and go through the branch labeled with “the” reaching the node-①, which contains the identifier of S_0 . Then, we read “more” ($RS = \text{“the more”}$), hence we descend from node-① to node-②, which contains the symbol S_2 . Finally, we read “I” ($RS = \text{“the more I”}$). Since we cannot descend from node-②, the longest known sequence is “the more” and its symbol is S_2 .

Encoding procedure

Every parsed symbol must be encoded and sent to the receiver. We encode them using *etdc* dense codes. We need to keep track of the number of times each symbol was sent (frequency) because, following *etdc* procedure, the codeword of a symbol depends only on its rank within the vocabulary sorted by frequency. Recall *etdc* assigns the shortest codewords to the most frequent symbols. Note that, each time a symbol is sent, its frequency is increased, and the codewords assigned to the symbols may change. For each parsed symbol S_i we send one codeword. In addition, when we send a terminal symbol for the first time ($S_i = S_{|KS|}$), we send that codeword, which acts as an escape codeword, followed by the word in plain format.

In order to encode the symbols, we use a *codebook* where we store all the information required to compute the codeword of each symbol. Each entry in *codebook* corresponds to a symbol S_i and stores a tuple $\langle l, r, freq, voc \rangle$ as shown in Figure 2(left). l and r represent the sequence of each symbol. If the symbol is a non-terminal, l and r are pointers to the entries of *codebook* where the left and right symbols of the non-terminal are stored. Otherwise, if the symbol is a terminal, l stores the word itself and r is set to -1. $freq$ stores the frequency of the symbol³ and voc holds the position of the symbol within the vocabulary sorted by frequency. The codeword c_i

²We implemented a bit-oriented trie where unary paths are stored in their parent node.

³Every non-terminal symbol is created with frequency 0.

corresponding to the symbol S_i stored in the i -th entry of the codebook is obtained as $c_i \leftarrow \text{etdc.encode}(\text{voc}[i])$.

To keep the vocabulary sorted by frequency we use two arrays: pos and top . Array pos keeps the symbols sorted by frequency in decreasing order. Actually, $\text{pos}[i] = j$ indicates that the i -th most frequent symbol is stored in the j -th entry of the *codebook*. Consequently, note that all the symbols with the same frequency are pointed to from consecutive entries in pos . Array top contains a slot for each frequency value. For every possible value of a frequency f , $\text{top}[f] = x$ means that the first symbol with frequency f is at position x in pos . For example, in Figure 2(left) the array top indicates that the codewords of frequency 1 start at position 0 within pos . We can observe that the gap between $\text{top}[0]$ and $\text{top}[1]$ is 6, thus $\text{pos}[0..5]$ point to the 6 entries within *codebook* that hold all the symbols with frequency 1.

With the help of the arrays pos and top , we can easily add the new symbols at the end of *codebook*. Those arrays are also necessary to update the frequencies and positions in the vocabulary in $O(1)$ time without reordering the *codebook*. In our example, after inserting “*you*”, the table remains in the state of Figure 2(left). As we explained before, in the next step we send $S_2 = \text{“the more”}$ which is the symbol at position 2. Therefore, we increase the frequency at $\text{freq}[2]$ to 1. We look for the position of the first symbol with frequency 0 by using $\text{top}[0] = 6$. After that, we swap $\text{voc}[2]$ and $\text{voc}[6]$, so now $\text{voc}[2] = 6$ and $\text{voc}[6] = 9$. As we changed voc , we also have to update pos accordingly. Therefore, we modify $\text{pos}[6] = 2$ and $\text{pos}[9] = 6$. Finally, as now the list of symbols with $f = 1$ has been increased by one, the list of words with $f = 0$ starts one position further, so we update $\text{top}[0] = 7$.

Receiver procedure

The receiver works symmetrically to the sender. It decodes either a codeword corresponding to a known symbol or an escape codeword followed by a new word (terminal) in plain form. After decoding a symbol, we also add a new non-terminal composed of the last two decoded symbols to keep the codebook synchronized with the sender. This allows the receiver to rebuild the same model handled by the sender and to recover the original text. To carry this out the receiver also has a *codebook* and an auxiliary top array. The *codebook* is composed of columns *offset*, *length*, and *freq*. Each time we create a new symbol (i.e. we either received a new word or we created a new non-terminal), we set in *offset* a pointer to the position of the first occurrence of that symbol within the decompressed text. The length (in chars) of the text represented by such symbol is kept in *length*. *freq* stores the frequency of the symbol.

In Figure 2(right) we can observe the state of the receiver after decompressing “*the more I know about you the more*”. Now the sender transmits the symbol $S_8 = \text{“I know”}$ encoded with *etdc*. The receiver decodes the codeword into 8. It accesses to the *codebook* at position 8 and retrieves $\text{offset}[8] = 9$ and $\text{length}[8] = 6$, thus the decoder recovers the sequence “*I know*” from the decompressed text from position 9 to 14. Afterwards, we update the decompressed text to “*the more I know about you the more*”||“*I know*”. Then, we increase $\text{freq}[8]$, and we swap the rows in the *codebook* at positions 8 and $\text{top}[0] = 7$ (recall $\text{top}[0]$ is the first row with frequency

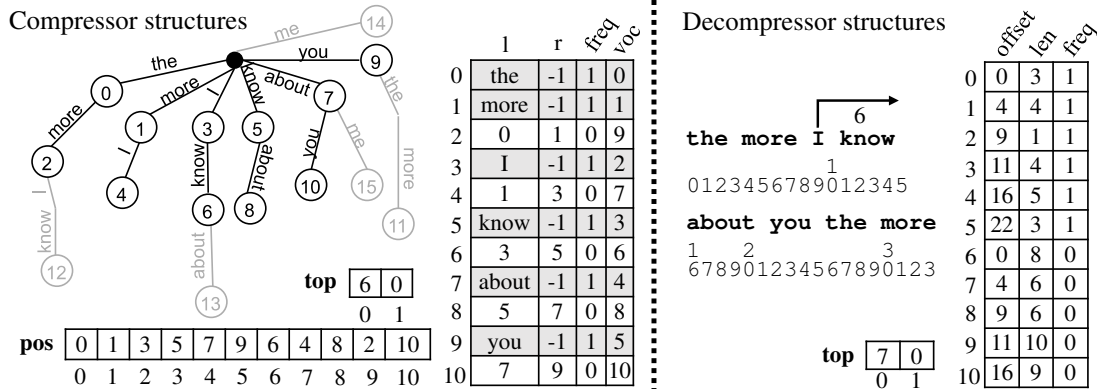


Figure 2: Structures used in both compressor (left) and decompressor (right) when processing the sentence: “the more I know about you the more I know about me”. The black branches in the trie represent its stage after processing the word “you”.

equals to 0). Finally, since the first row with frequency equals to 0 is moved to the next position, we update $top[0] = 8$.

3 Experimental evaluation

We performed experiments to compare the compression effectiveness as well as the performance at compression and decompression of $D-v2v$ with those of $detdc^4$ and $v2vdc$, which are respectively the previous dynamic word-based technique that makes up the basis of $D-v2v$, and the state-of-the-art when considering semi-static variable-to-variable compression based on dense codes. In the case of $v2vdc$, we considered the two variants proposed in [8], i.e. $v2vdc$ and $v2vdc_H$. The former one uses a simpler heuristic to gather phrases, whereas the latter uses a more complex heuristic that yields better compression at the cost of increased compression time. Given that in $D-v2v$ new words are sent in plain form, we included two variants of $v2vdc$ and $v2vdc_H$ that, as in [8], respectively represent the words in the vocabulary in plain form or compressed with $lzma$. In addition, we have included some of the most well-known representatives from different families of compressors: $p7zip$ and $lzma$,⁵ $bzip2$,⁶ and an implementation of $re-pair$,⁷ coupled with a bit-oriented Huffman.⁸

We used three text datasets from TREC-2 and TREC-4 named Ziff Data 1989-1990 (ZIFF), Congressional Record 1993 (CR) and Financial Times 1991 (FT91). In addition, we created a large dataset (ALL) including ZIFF and AP-newswire from TREC-2, as well as Financial Times 1991 to 1994 (FT91, FT92, FT93, and FT94) and ZIFF from TREC-4. We also included three highly repetitive text datasets: *world.leaders.txt* (WL), *english.001.2.txt* (ENG) and *einstein.en.txt* (EINS) from PIZZACHILI.⁹

⁴<http://vios.dc.fi.udc.es/codes>

⁵<http://www.7-zip.org>

⁶<http://www.bzip.org>

⁷<http://raymondwan.people.ust.hk/en/restore.html>

⁸https://people.eng.unimelb.edu.au/ammoffat/mr_coder/

⁹<http://pizzachili.dcc.uchile.cl>

Our test machine is an Intel(R) Core(TM) i7-3820@3.60GHz CPU (4cores-8siblings) with 64GB of DDR3-1600Mhz. It runs Ubuntu 12.04.5 LTS (kernel 3.2.0-126-generic). We compiled with gcc 4.6.4 and optimizations `-O9`. Our time results measure CPU user time.

Table 1: Compression ratio (%) with respect to the size of the plain text dataset.

	Detdc	v2vdc	v2vdc _H	v2vdc	v2vdc _H	D-v2v	Repair	lzma		p7zip	bzip2	Size Plain (KB)
		lzma words	plain words	+sHuff	def		-9 -e	def				
FT91	35,64	27,15	26,65	30,11	29,61	28,60	24,00	25,50	25,25	25,52	27,06	14.404
CR	31,99	23,55	23,13	24,73	24,31	22,86	20,16	22,05	20,83	21,63	24,14	49.888
ZIFF	33,79	24,01	23,60	24,66	24,25	23,14	20,33	23,40	21,64	22,98	25,10	180.879
ALL	33,66	22,81	–	23,39	–	22,67	–	23,23	21,34	22,80	25,98	1.055.391
WL	15,06	4,13	–	4,44	–	2,90	1,43	1,30	1,11	1,39	6,94	45.867
ENG	35,21	–	–	–	–	5,52	2,17	0,55	0,55	0,55	3,73	102.400
EINS	30,14	0,97	–	0,98	–	0,27	0,07	0,07	0,07	0,07	5,17	456.667

In Table 1, we compare the compression ratios obtained. We can see that *D-v2v* is able to improve the results of *v2vdc* (and *v2vdc_H*) in all datasets (results with ‘–’ indicate failed runs). This is remarkable since we are sending new words in plain form, while the best values of *v2vdc* are drawn when it encodes the vocabulary of words with *lzma*. As expected, by using not only words in the vocabulary of symbols allows *D-v2v* to overcome the original *detdc* by more than 10 percentage points in regular English datasets, and completely blows *detdc* out in repetitive collections. On regular texts, *D-v2v* and *p7zip* obtain similar values on the largest dataset, yet in the other datasets the fact of exploiting char- rather than word-based regularities benefit *p7zip*, *lzma*, and *re-pair*. In repetitive text collections, char-level repetitiveness is higher than at word-level, and in addition, the fact of sending words in plain form harms *D-v2v* compression. In practice, even though compression is good in *D-v2v*, it is typically far from *re-pair*, *p7zip*, and *lzma*.

In Table 2, we include both compression and decompression times. *D-v2v* is faster at compression than *p7zip*, *lzma*, and *re-pair*. It is on a par with *v2vdc_H*, and it is slower than *bzip2*. Of course *detdc*, which has not to deal with the detection of seen subsequences, is much simpler and faster than *D-v2v*.

At decompression, we can see that again *D-v2v* is the fastest technique in all cases, with the exception of *detdc* and *v2vdc* when dealing with non-repetitive English texts. Note that, in this case, *v2vdc* compression is similar to that of *D-v2v* and consequently both decode approximately the same number of codewords. However, *v2vdc* has not to perform an *update* procedure after decoding each symbol nor to generate a new non-terminal. *detdc* has to decode more symbols than *D-v2v* due to its worse compression. Yet, again it is simpler because it does not have to deal with non-terminals, only with words. In the repetitive collections *D-v2v* compresses much more than *v2vdc* and *detdc*, which leads to a compressed file with much less codewords than those of *detdc* and *v2vdc*, and this amortizes the cost of the *update* procedure required after decoding each codeword.

In Table 3, we can see memory usage at compression time. In this case, our current implementation of the trie in *D-v2v* requires lots of memory. At decompression time, we only have to deal with the *codebook* (the size of *top* is negligible), and the

memory usage becomes much more reasonable. Yet, the number of entries in the *codebook* is still very high in most datasets: {1,6M@FT91}; {4,3M@CR}; {15,4@ZIFF}; {81,2M@ALL}; {0,44M@WL}; {1,8M@ENG}; {0,3M@EINS}.

Table 2: Compression and decompression times (in seconds)

	Text	Detdc	v2vdc	v2vdcH	v2vdc	v2vdcH	D-v2v	Repair	lzma		p7zip	bzip2
			<i>lzma words</i>		<i>plain words</i>			<i>+sHuff</i>	<i>def</i>	<i>-9 -e</i>		<i>def</i>
Compr. time	FT91	0.15	1.31	2.26	1.28	2.27	5.81	8.37	9.17	10.66	9.06	1.19
	CR	0.53	5.77	19.01	5.72	19.10	19.03	39.84	32.61	44.09	33.67	4.07
	ZIFF	2.12	32.08	257.42	31.94	258.03	86.13	271.02	120.92	177.86	128.99	14.52
	ALL	13.25	292.39	–	289.05	–	573.09	–	711.23	1167.86	768.22	86.71
	WL	0.48	17.93	–	18.05	–	2.96	14.99	8.88	23.60	6.23	2.45
	ENG	1.71	–	–	–	–	13.39	58.97	28.37	57.34	28.31	8.45
	EINS	6.62	33205.00	–	33197.00	–	30.48	205.33	60.98	115.12	57.13	54.95
Decompr. time	FT91	0.09	0.08	0.09	0.06	0.06	0.09	0.15	0.18	0.17	0.19	0.47
	CR	0.28	0.22	0.23	0.20	0.21	0.31	0.65	0.54	0.54	0.54	1.54
	ZIFF	1.24	1.01	0.90	0.94	0.86	1.50	2.69	2.13	2.14	2.14	5.83
	ALL	7.68	9.13	–	8.99	–	11.63	–	12.13	12.25	12.15	33.54
	WL	0.16	0.06	–	0.06	–	0.02	0.28	0.05	0.05	0.09	1.00
	ENG	0.85	–	–	–	–	0.14	2.50	0.06	0.04	0.18	4.08
	EINS	3.22	0.33	–	0.24	–	0.05	1.71	0.15	0.15	0.71	9.40

Table 3: Memory usage (in MiB) at compression and decompression

	Text	Detdc	v2vdc	v2vdcH	v2vdc	v2vdcH	D-v2v	Repair	lzma		p7zip	bzip2
			<i>lzma words</i>		<i>plain words</i>			<i>+sHuff</i>	<i>def</i>	<i>-9 -e</i>		<i>def</i>
Compressor	FT91	24	52	52	52	52	1,194	380	94	192	165	7
	CR	53	157	157	157	157	2,635	1,286	94	504	193	7
	ZIFF	126	625	625	625	625	10,509	4,585	94	674	193	7
	ALL	207	3,509	–	3,509	–	46,160	–	94	673	193	8
	WL	49	255	–	255	–	478	1,268	94	469	193	7
	ENG	85	–	–	–	–	1,953	2,512	94	674	193	7
	EINS	152	44,821	–	9,851	–	6,521	10,859	94	674	193	8
Decompressor	FT91	4	10	20	20	10	20	13	9	15	17	4
	CR	6	65	65	65	65	51	30	9	50	19	4
	ZIFF	14	121	119	121	119	177	79	9	65	19	4
	ALL	57	378	–	378	–	931	–	9	65	19	4
	WL	5	52	–	51	–	6	11	9	46	18	4
	ENG	9	–	–	–	–	23	31	9	65	18	4
	EINS	14	73.81	–	73.81	–	5	5	9	65	18	4

4 Conclusions and future work

We have described *D-v2v*, the first word-based dynamic variable-to-variable text compressor. We showed that *D-v2v* obtains competitive compression ratios (similar to *p7zip*) in English texts and that it is fast at both compression and (mainly) decompression. Even not included in the paper, note that looking for the occurrences of a given word *P* would be possible by counting the number of escape codewords until the first occurrence of *P* (that counter indicates the initial entry of the *codebook* where we add *P*). From there on, by simulating the decompressing process we need to track the occurrences of the codeword corresponding to the terminal *P* and those codewords corresponding to all the non-terminals which include *P*.

The main drawback of *D-v2v* is that it needs lots of memory at compression to handle the subsequences (non-terminals) in the trie. As future work, we will improve

the current implementation of the trie to reduce memory requirements. We also want to apply the ideas in [15] to create an asymmetric lightweight version of $D-v2v$. This should reduce the work done by the receiver and its memory usage. In addition, the codeword associated to a given symbol S_i would not vary so often, which would allow us to implement efficient direct searches for a pattern within the compressed text.

References

- [1] D. A. Huffman, "A method for the construction of minimum-redundancy codes," *Proc. I.R.E.*, vol. 40, no. 9, pp. 1098–1101, 1952.
- [2] A. Moffat, "Word-based text compression," *Softw. Pract. Exp.*, vol. 19, no. 2, pp. 185–198, 1989.
- [3] I. Witten, A. Moffat, and T. Bell, *Managing Gigabytes: Compressing and Indexing Documents and Images*, Morgan Kaufmann Publishers, 1999.
- [4] E. Moura, G. Navarro, N. Ziviani, and R. Baeza-Yates, "Fast searching on compressed text allowing errors," in *Proc. ACM SIGIR*, 1998, pp. 298–306.
- [5] G. Navarro, E. Moura, M. Neubert, N. Ziviani, and R. Baeza-Yates, "Adding compression to block addr. inverted indexes," *Inf. Retr.*, vol. 3, no. 1, pp. 49–77, 2000.
- [6] N. Brisaboa, A. Fariña, G. Navarro, and J. Paramá, "Lightweight natural language text compression," *Inf. Retrieval*, vol. 10, no. 1, pp. 1–33, 2007.
- [7] N. Brisaboa, A. Fariña, G. Navarro, and J. Paramá, "New adaptive compressors for natural language text," *Softw. Pract. Exp.*, vol. 38, no. 13, pp. 1429–1450, 2008.
- [8] N.R. Brisaboa, A. Fariña, JR. López, G. Navarro, and E. Rodríguez, "A new searchable variable-to-variable compressor," in *Proc. DCC*, 2010, pp. 199–208.
- [9] A. Moffat and A. Turpin, *Compression and Coding Algorithms*, Kluwer Academic Publishers, 2002.
- [10] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Trans. Inf. Theor.*, vol. 23, no. 3, pp. 337–343, Sept. 2006.
- [11] J. Ziv and A. Lempel, "Compression of individual sequences via variable-rate coding," *IEEE Trans. Inf. Theor.*, vol. 24, no. 5, pp. 530–536, Sept. 2006.
- [12] N. Faller, "An adaptive system for data compression," in *Record of the 7th Asilomar Conference on Circuits, Systems, and Computers*, 1973, pp. 593–597.
- [13] R. Gallager, "Variations on a theme by huffman," *IEEE Trans. Inf. Theor.*, vol. 24, no. 6, pp. 668–674, 1978.
- [14] A. Fariña, *New compression codes for text databases*, Ph.D. thesis.
- [15] N. R. Brisaboa, A. Fariña, G. Navarro, and J. R. Paramá, "Dynamic lightweight text compression," *ACM Trans. Inf. Sys. (TOIS)*, pp. 1–32, 2010.
- [16] Donald R. Morrison, "PATRICIA - practical algorithm to retrieve information coded in alphanumeric," *J. ACM*, vol. 15, no. 4, pp. 514–534, 1968.
- [17] M. Charikar, E. Lehman, D. Liu, R. Panigrahy, M. Prabhakaran, A. Sahai, and A. Shelat, "The smallest grammar problem," *IEEE Trans. Inf. Theo.*, vol. 51, no. 7, pp. 2554–2576, 2005.
- [18] N. J. Larsson and A. Moffat, "Offline dictionary-based compression," in *Proc. DCC*, Washington, DC, USA, 1999, p. 296, IEEE Computer Society.
- [19] C. Nevill-Manning, I. Witten, and D. Maulsby, "Compression by induction of hierarchical grammars," in *Proc. DCC*, 1994, pp. 244–253.
- [20] A. Apostolico and S. Lonardi, "Off-line compression by greedy textual substitution," *Proc. IEEE*, vol. 88, no. 11, pp. 1733–1744, 2000.
- [21] U. Manber and G. Myers, "Suffix arrays: a new method for on-line string searches," *SIAM J. Comp.*, vol. 22, no. 5, pp. 935–948, 1993.