

Lempel-Ziv Decoding in External Memory

Djamal Belazzougui¹, Juha Kärkkäinen²(✉), Dominik Kempa²,
and Simon J. Puglisi²

¹ CERIST, Algiers, Algeria
dbelazzougui@cerist.dz

² Department of Computer Science, Helsinki Institute for Information
Technology HIIT, University of Helsinki, Helsinki, Finland
{juha.karkkainen,dominik.kempa,simon.puglisi}@cs.helsinki.fi

Abstract. Simple and fast decoding is one of the main advantages of LZ77-type text encoding used in many popular file compressors such as *gzip* and *7zip*. With the recent introduction of external memory algorithms for Lempel–Ziv factorization there is a need for external memory LZ77 decoding but the standard algorithm makes random accesses to the text and cannot be trivially modified for external memory computation. We describe the first external memory algorithms for LZ77 decoding, prove that their I/O complexity is optimal, and demonstrate that they are very fast in practice, only about three times slower than in-memory decoding (when reading input and writing output is included in the time).

1 Introduction

The Lempel–Ziv (LZ) factorization [18] is a partitioning of a text string into a minimal number of phrases consisting of substrings with an earlier occurrence in the string and of single characters. In LZ77 encoding [20] the repeated phrases are replaced by a pointer to an earlier occurrence (called the source of the phrase). It is a fundamental tool for data compression [6, 7, 15, 17] and today it lies at the heart of popular file compressors (e.g. *gzip* and *7zip*), and information retrieval systems (see, e.g., [6, 10]). Recently the factorization has become the basis for several compressed full-text self-indexes [5, 8, 9, 16]. Outside of compression, LZ factorization is a widely used algorithmic tool in string processing: the factorization lays bare the repetitive structure of a string, and this can be used to design efficient algorithms [2, 12–14].

One of the main advantages of LZ77 encoding as a compression technique is a fast and simple decoding: simply replace each pointer to a source by a copy of the source. However, this requires a random access to the earlier part of the text. Thus the recent introduction of external memory algorithms for LZ77 factorization [11] raises the question: Is fast LZ77 decoding possible when the text length exceeds the RAM size? In this paper we answer the question positively by describing the first external memory algorithms for LZ77 decoding.

This research is partially supported by Academy of Finland through grant 258308 and grant 250345 (CoECGR).

In LZ77 compression, the need for external memory algorithms can be avoided by using an encoding window of limited size. However, a longer encoding window can improve the compression ratio [6]. Even with a limited window size, decompression on a machine with a small RAM may require an external memory algorithm if the compression was done on a machine with a large RAM. Furthermore, in applications such as text indexing and string processing limiting the window size is not allowed. While most of these applications do not require decoding, a fast decoding algorithm is still useful for checking the correctness of the factorization.

Our Contribution. We show that in the standard external memory model of computation [19] the I/O complexity of decoding an LZ77-like encoding of a string of length n over an alphabet of size σ is $\Theta\left(\frac{n}{B \log_\sigma n} \log_{M/B} \frac{n}{B \log_\sigma n}\right)$, where M is the RAM size and B is the disk block size in units of $\Theta(\log n)$ bits. The lower bound is shown by a reduction from permuting and the upper bound by describing two algorithms with this I/O complexity.

The first algorithm uses the powerful tools of external memory sorting and priority queues while the second one relies on plain disk I/O only. Both algorithms are relatively simple and easy to implement. Our implementation uses the STXXL library [4] for sorting and priority queues.

Our experiments show that both algorithms scale well for large data but the second algorithm is much faster in all cases. This shows that, while external memory sorting and priority queues are extremely useful tools, they do have a significant overhead when their full power is not needed. The faster algorithm (using a very modest amount of RAM) is only 3–4 times slower than an in-memory algorithm that has enough RAM to perform the decoding in RAM (but has to read the input from disk and write the output to disk).

Our algorithms do not need a huge amount of disk space in addition to the input (factorization) and output (text), but we also describe and implement a version, which can reduce the additional disk space to less than 3% of total disk space usage essentially with no effect on runtime.

2 Basic Definitions

Strings. Throughout we consider a string $X = X[1..n] = X[1]X[2] \dots X[n]$ of $|X| = n$ symbols drawn from the alphabet $[0..\sigma - 1]$ for $\sigma = n^{\mathcal{O}(1)}$. For $1 \leq i \leq j \leq n$ we write $X[i..j]$ to denote the *substring* $X[i]X[i+1] \dots X[j]$ of X . By $X[i..j]$ we denote $X[i..j - 1]$.

LZ77. The *longest previous factor* (LPF) at position i in string X is a pair $\text{LPF}[i] = (p_i, \ell_i)$ such that $p_i < i$, $X[p_i..p_i + \ell_i] = X[i..i + \ell_i]$, and ℓ_i is maximized. In other words, $X[i..i + \ell_i]$ is the longest prefix of $X[i..n]$ which also occurs at some position $p_i < i$ in X . There may be more than one potential value of p_i , and we do not care which one is used.

The LZ77 factorization (or LZ77 parsing) of a string X is a greedy, left-to-right parsing of X into longest previous factors. More precisely, if the j th LZ factor (or *phrase*) in the parsing is to start at position i , then $\text{LZ}[j] = \text{LPF}[i] = (p_i, \ell_i)$ (to represent the j th phrase), and then the $(j + 1)$ th phrase starts at position $i + \ell_i$. The exception is the case $\ell_i = 0$, which happens iff $X[i]$ is the leftmost occurrence of a symbol in X . In this case $\text{LZ}[j] = (X[i], 0)$ (to represent $X[i..i]$) and the next phrase starts at position $i + 1$. This is called a *literal phrase* and the other phrases are called *repeat phrases*. For a repeat phrases, the substring $X[p_i..p_i + \ell_i]$ is called the *source* of the phrase $X[i..i + \ell_i]$. We denote the number of phrases in the LZ77 parsing of X by z .

LZ77-type Factorization. There are many variations of LZ77 parsing. For example, the original LZ77 encoding [20] had only one type of phrase, a (potentially empty) repeat phrase always followed by a literal character. Many compressors use parsing strategies that differ from the greedy strategy described above to optimize compression ratio after entropy compression or to speed up compression or decompression. The algorithms described in this paper can be easily adapted for most of them. For purposes of presentation and analysis we make two assumptions about the parsing:

- All phrases are either literal or repeat phrases as described above.
- The total number of repeat phrases, denoted by z_{rep} , is $\mathcal{O}(n/\log_{\sigma} n)$.

We call this an *LZ77-type factorization*. The second assumption holds for the greedy factorization [18] and can always be achieved by replacing too short repeat phrases with literal phrases. We also assume that the repeat phrases are encoded using $\mathcal{O}(\log n)$ bits and the literal phrases using $\mathcal{O}(\log \sigma)$ bits. Then the size of the whole encoding is never more than $\mathcal{O}(n \log \sigma)$ bits.

3 On I/O Complexity of LZ Decoding

Given an LZ77-type factorization of a string encoded as described above, the task of LZ77 decoding is to recover the original string. In this section, we obtain a lower bound on the I/O complexity of LZ decoding by a reduction from permuting.

We do the analysis using the standard external memory model [19] with RAM size M and disk block size B , both measured in units of $\Theta(\log n)$ bits. We are primarily interested in the I/O complexity, i.e., the number of disk blocks moved between RAM and disk.

Given a sequence $\bar{x} = x_1, x_2, \dots, x_n$ of n objects of size $\Theta(\log n)$ bits each and a permutation $\pi[1..n]$ of $[1..n]$, the task of permuting is to obtain the permuted sequence $\bar{y} = y_1, y_2, \dots, y_n = x_{\pi[1]}, x_{\pi[2]}, \dots, x_{\pi[n]}$. Under the mild assumption that $B \log(M/B) = \Omega(\log(n/B))$, the I/O complexity of permuting is $\Theta\left(\frac{n}{B} \log_{M/B} \frac{n}{B}\right)$, the same as the I/O complexity of sorting [1].

We show now that permuting can be reduced to LZ decoding. Let X be the string obtained from the sequence \bar{x} by encoding each x_i as a string of length

$h = \Theta(\log_\sigma n)$ over the alphabet $[0..\sigma)$. Let Y be the string obtained in the same way from the sequence \bar{y} . Form an LZ77-type factorization of XY by encoding the first half using literal phrases and the second half using repeat phrases so that the substring representing y_i is encoded by the phrase $(h\pi[i] + 1 - h, h)$. This LZ factorization is easy to construct in $\mathcal{O}(n/B)$ I/Os given \bar{x} and π . By decoding the factorization we obtain XY and thus \bar{y} .

Theorem 1. *The I/O complexity of decoding an LZ77-type factorization of a string of length n over an alphabet of size σ is*

$$\Omega\left(\frac{n}{B \log_\sigma n} \log_{M/B} \frac{n}{B \log_\sigma n}\right).$$

Proof. The result follows by the above reduction from permuting a sequence of $\Theta(n/\log_\sigma n)$ objects. \square

For comparison, the worst case I/O complexity of naive LZ decoding is $\mathcal{O}(n/\log_\sigma n)$.

4 LZ Decoding Using EM Sorting and Priority Queue

Our first algorithm for LZ decoding relies on the powerful tools of external memory sorting and external memory priority queues.

We divide the string X into $\lceil n/b \rceil$ segments of size exactly b (except the last segment can be smaller). The segments must be small enough to fit in RAM and big enough to fill at least one disk block. If a phrase or its source overlaps a segment boundary, the phrase is split so that all phrases and their sources are completely inside one segment. The number of phrases increases by at most $\mathcal{O}(z_{\text{rep}} + n/b)$ because of the splitting.

After splitting, the phrases are divided into three sequences. The sequence R_{far} contains repeat phrases with the source more than b positions before the phrase (called far repeat phrases) and the sequence R_{near} the other repeat phrases (called near repeat phrases). The sequence L contains all the literal phrases. The repeat phrases are represented by triples (p, q, ℓ) , where p is the starting position of the source, q is the starting position of the phrase and ℓ is the length. The literal phrases are represented by pairs (q, c) , where q is the phrase position and c is the character. The sequence R_{far} of far repeat phrases is sorted by the source position. The other two sequences are not sorted, i.e., they remain ordered by the phrase position.

During the computation, we maintain an external memory priority queue Q that stores already recovered far repeat phrases. Each such phrase is represented by a triple (q, ℓ, s) , where q and ℓ are as above and s is the phrase as a literal string. The triples are extracted from the queue in the ascending order of q . The maximum length of phrases stored in the queue is bounded by a parameter ℓ_{max} . Longer phrases are split into multiple phrases before inserting them into the queue.

The string X is recovered one segment at a time in left-to-right order and each segment is recovered one phrase at a time in left-to-right order. A segment recovery is done in a (RAM) array $Y[0..b]$ of size b . At any moment in time, for some $i \in [0..b]$, $Y[0..i]$ contains the already recovered prefix of the current segment and $Y[i..b]$ contains the last $b - i$ characters of the preceding segment. The next phrase starting at $Y[i]$ is recovered in one of three ways depending on its type:

- A literal phrase is obtained as the next phrase in the sequence L .
- A near repeat phrase is obtained as the next phrase in the sequence R_{near} . The source of the phrase either starts in $Y[0..i]$ or is contained in $Y[i..b]$, and is easily recovered in both cases.
- A far repeat phrase is obtained from the priority queue with the full literal representation.

Once a segment has been fully recovered, we read all the phrases in the sequence R_{far} having the source within the current segment. Since R_{far} is ordered by the source position, this involves a single sequential scan of R_{far} over the whole algorithm. Each such phrase is inserted into the priority queue Q with its literal representation (splitting the phrase into multiple phrases if necessary).

Theorem 2. *A string of length n over an alphabet of size σ can be recovered from its LZ77 factorization in $\mathcal{O}\left(\frac{n}{B \log_\sigma n} \log_{M/B} \frac{n}{B \log_\sigma n}\right)$ I/Os.*

Proof. We set $\ell_{\text{max}} = \Theta(\log_\sigma n)$ and $b = \Theta(B \log_\sigma n)$. Then the objects stored in the priority queue need $\mathcal{O}(\log n + \ell_{\text{max}} \log \sigma) = \mathcal{O}(\log n)$ bits each and the total number of repeat phrases after all splitting is $\mathcal{O}(z_{\text{rep}} + n/\log_\sigma n) = \mathcal{O}(n/\log_\sigma n)$. Thus sorting the phrases needs $\mathcal{O}\left(\frac{n}{B \log_\sigma n} \log_{M/B} \frac{n}{B \log_\sigma n}\right)$ I/Os. This is also the I/O complexity of all the external memory priority queue operations [3]. All other processing is sequential and needs $\mathcal{O}\left(\frac{n}{B \log_\sigma n}\right)$ I/Os. \square

We have implemented the algorithm using the STXXL library [4] for external memory sorting and priority queues.

5 LZ Decoding Without Sorting or Priority Queue

The practical performance of the algorithm in the previous section is often bounded by in-memory computation rather than I/O, at least on a machine with relatively fast disks. In this section, we describe an algorithm that reduces computation based on the observation that we do not really need the full power of external memory sorting and priority queues.

To get rid of sorting, we replace the sorted sequence R_{far} with $\lceil n/b \rceil$ unsorted sequences R_1, R_2, \dots , where R_i contains all phrases with the source in the i th segment. In other words, sorting R_{far} is replaced with distributing the phrases

into R_1, R_2, \dots . If n/b is less than M/B , the distribution can be done in one pass, since we only need one RAM buffer of size B for each segment. Otherwise, we group M/B consecutive segments into a supersegment, distribute the phrases first into supersegments, and then into segments by scanning the supersegment sequences. If necessary, further layers can be added to the segment hierarchy. This operation generates the same amount of I/O as sorting but requires less computation because the segment sequences do not need to be sorted.

In the same way, the priority queue is replaced with $\lceil n/b \rceil$ simple queues. The queue Q_i contains a triple (q, ℓ, s) for each far repeat phrase whose phrase position is within the i th segment. The order of the phrases in the queue is arbitrary. Instead of inserting a recovered far repeat phrase into the priority queue Q it is appended into the appropriate queue Q_i . This requires a RAM buffer of size B for each queue but as above a multi-round distribution can be used if the number of segments is too large. This approach might not reduce the I/O compared to the use of a priority queue but it does reduce computation. Moreover, the simple queue allows the strings s to be of variable sizes and of unlimited length; thus there is no need to split the phrases except at segment boundaries.

Since the queues Q_i are not ordered by the phrase position, we can no more recover a segment in a strict left-to-right order, which requires a modification of the segment recovery procedure. The sequence R_{near} of near repeat phrases is divided into two: R_{prev} contains the phrases with the source in the preceding segment and R_{same} the ones with the source in the same segment.

As before, the recovery of a segment X_j starts with the previous segment in the array $Y[0..b)$ and consists of the following steps:

1. Recover the phrases in R_{prev} (that are in this segment). Note that each source is in the part of the previous segment that is still untouched.
2. Recover the literal phrases by reading them from L .
3. Recover the far repeat phrases by reading them from Q_j (with the full literal representation).
4. Recover the phrases in R_{same} . Note that each source is in the part of the current segment that has been fully recovered.

After the recovery of the segment, we read all the phrases in R_j and insert them into the queues Q_k with their full literal representations.

We want to minimize the number of segments. Thus we choose the segment size to occupy at least half of the available RAM and more if the RAM buffers for the queues Q_k do not require all of the other half. It is easy to see that this algorithm does not generate asymptotically more I/Os than the algorithm of the previous section. Thus the I/O complexity is $\mathcal{O}\left(\frac{n}{B \log_{\sigma} n} \log_{M/B} \frac{n}{B \log_{\sigma} n}\right)$. We have implemented the algorithm using standard file I/O (without the help of STXXL).

6 Reducing Disk Space Usage

The algorithm described in the previous section can adapt to a small RAM by using short segments, and if necessary, multiple rounds of distribution. However, reducing the segment size does not affect the disk space usage and the algorithm will fail if it does not have enough disk space to store all the external memory data. In this section, we describe how the disk space usage can be reduced.

The idea is to divide the LZ factorization into parts and to process one part at a time recovering the corresponding part of the text. The first part is processed with the algorithm of the previous section as if it was the full string. To process the later parts, a slightly modified algorithm is needed because, although all the phrases are in the current part, the sources can be in the earlier parts. Thus we will have the R_j queues for all the segments in the current and earlier parts but the Q_j queues only for the current part. The algorithm processes first all segments in the previous parts performing the following steps for each segment X_j :

- Read X_j from disk to RAM.
- Read R_j and for each phrase in R_j create the triple (q, ℓ, s) and write it to the appropriate queue Q_k .

Then the segments of the current part are processed as described in the previous section.

For each part, the algorithm reads all segments in the preceding parts. The number of additional I/Os needed for this is $\mathcal{O}(np/(B \log_\sigma n))$, where p is the number of parts. In other respects, the performance of the algorithm remains essentially the same.

We have implemented this partwise processing algorithm using greedy on-line partitioning. That is, we make each part as large as possible so that the peak disk usage does not exceed a given disk space budget. An estimated peak disk usage is maintained while reading the input. The implementation needs at least enough disk space to store the input (the factorization) and the output (the recovered string) but the disk space needed in addition to that can usually be reduced to a small fraction of the total with just a few parts.

7 Experimental Results

Setup. We performed experiments on a machine equipped with two six-core 1.9 GHz Intel Xeon E5-2420 CPUs with 15 MiB L3 cache and 120 GiB of DDR3 RAM. The machine had 7.2 TiB of disk space striped with RAID0 across four identical local disks achieving a (combined) transfer rate of about 480 MiB/s. The STXXL block size as well as the size of buffers in the algorithm based on plain disk I/O was set to 1 MiB.

The OS was Linux (Ubuntu 12.04, 64bit) running kernel 3.13.0. All programs were compiled using g++ version 4.7.3 with `-O3 -DNDEBUG` options. The machine had no other significant CPU tasks running and only a single thread of execution

Table 1. Statistics of data used in the experiments. All files are of size 256 GiB. The value of n/z (the average length of a phrase in the LZ77 factorization) is included as a measure of repetitiveness.

| Name | σ | n/z |
|-----------|----------|---------|
| hg.reads | 6 | 52.81 |
| wiki | 213 | 84.26 |
| kernel | 229 | 7767.05 |
| random255 | 255 | 4.10 |

was used for computation. All reported runtimes are wallclock (real) times. In the experiments with a limited amount of RAM, the machine was rebooted with a kernel boot flag so that the unused RAM is unavailable even for the OS.

Datasets. For the experiments we used the following files varying in the number of repetitions and alphabet size (see Table 1 for some statistics):

- hg.reads: a collection of DNA reads (short fragments produced by a sequencing machine) from 40 human genomes¹ filtered from symbols other than $\{A, C, G, T, N\}$ and newline;
- wiki: a concatenation of three different English Wikipedia dumps² in XML format dated: 2014-07-07, 2014-12-08, and 2015-07-02;
- kernel: a concatenation of ~ 16.8 million source files from 510 versions of Linux kernel³;
- random255: a randomly generated sequence of bytes.

Experiments. In the first experiment we compare the implementation of the new LZ77 decoding algorithm not using external-memory sorting or priority queue to a straightforward internal-memory LZ77 decoding algorithm that scans the input parsing from disk and decodes the text from left to right. All copying of text from sources to phrases happens in RAM.

We use the latter algorithm as a baseline since it represents a realistic upper bound on the speed of LZ77 decoding. It needs enough RAM to accommodate the output text as a whole, and thus we were only able to process prefixes of test files up to size of about 120 GiB. In the runtime we include the time it takes to read the parsing from disk (we stream the parsing using a small buffer) and to write the output text to disk. The new algorithm, being fully external-memory algorithm, can handle full test instances. The RAM usage of the new algorithm was limited to 3.5 GiB.

The results are presented in Fig. 1. In nearly all cases the new algorithm is about three times slower than the baseline. This is due to the fact that in

¹ <http://www.1000genomes.org/>.

² <http://dumps.wikimedia.org/>.

³ <http://www.kernel.org/>.

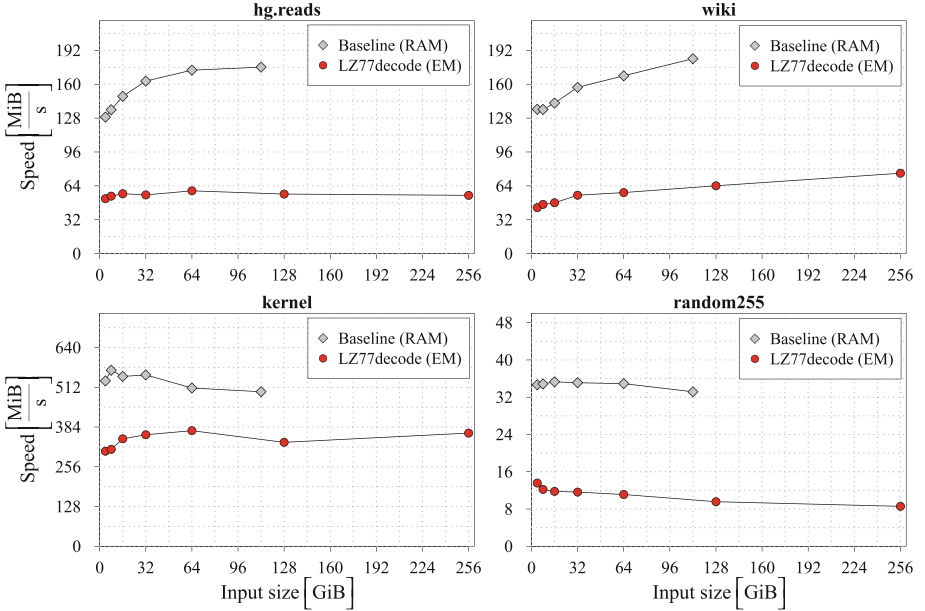


Fig. 1. Comparison of the new external memory LZ77 decoding algorithm based on plain disk I/O (“LZ77decode”) with the purely in-RAM decoding algorithm (“Baseline”). The latter represents an upper bound on the speed of LZ77 decoding. The unit of decoding speed is MiB of output text decoded per second.

the external memory algorithm each text symbol in a far repeat phrase is read or written to disk three times: first, when written to a queue Q_j as a part of a recovered phrase, second, when read from Q_j , and third, when we write the decoded text to disk. In comparison, the baseline algorithm transfers each text symbol between RAM and disk once: when the decoded text is written to disk. Similarly, while the baseline algorithm usually needs one cache miss to copy the phrase from the source, the external memory algorithm performs about three cache misses per phrase: when adding the source of a phrase to R_j , when adding a literal representation of a phrase into Q_j , and when copying the symbols from Q_j into their correct position in the text. The exception of the above behavior is the highly repetitive kernel testfile that contains many near repeat phrases, which are processed as efficiently as phrases in the RAM decoding algorithm.

In the second experiment we compare our two algorithms described in Sects. 4 and 5 to each other. For the algorithm based on priority queue we set $\ell_{\max} = 16$. The segment size in both algorithms was set to at least half of the available RAM (and even more if it did not lead to multiple rounds of EM sorting/distribution), except in the algorithm based on sorting we also need to allocate some RAM for the internal operations of STXXL priority queue. In all instances we allocate 1 GiB for the priority queue (we did not observe a notable effect on performance from using more space).

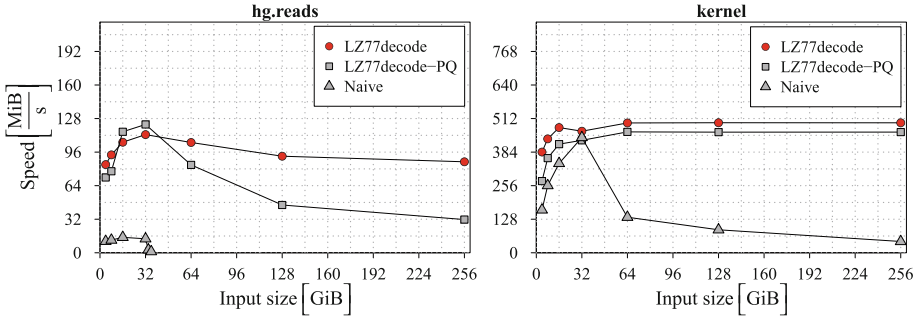


Fig. 2. Comparison of the new external memory LZ77 decoding algorithm based on plain disk I/O (“LZ77decode”) to the algorithm implemented using external memory sorting and priority queue (“LZ77decode-PQ”). The comparison also includes the algorithm implementing naive approach to LZ77 decoding in external memory. The speed is given in MiB of output text decoded per second.

In the comparison we also include a naive external-memory decoding algorithm that works essentially the same as baseline RAM algorithm except we do not require that RAM is big enough to hold the text. Whenever the algorithm requests a symbol outside a window, that symbol is accessed from disk. We do not explicitly maintain a window of recently decoded text in RAM, and instead achieve a very similar effect by letting the operating system cache the recently accessed disk pages. To better visualize the differences in performance, all algorithms were allowed to use 32 GiB of RAM.

The results are given in Fig. 2. For highly repetitive input (kernel) there is little difference between the new algorithms, as they both copy nearly all symbols from the window of recently decoded text. The naive algorithm performs much worse, but still finishes in reasonable time due to large average length of phrases (see Table 1).

On the non-repetitive data (hg.reads), the algorithm using external-memory sorting and priority queue clearly gets slower than the algorithm using plain disk I/O as the size of input grows. The difference in constant factors is nearly three for the largest test instance. The naive algorithm maintains acceptable speed only up to a point where the decoded text is larger than available RAM. At this point random accesses to disk dramatically slow down the algorithm.

Note also that the speed of our algorithm in Fig. 2 is significantly higher than in Fig. 1. This is because the larger RAM (32 GiB vs. 3.5 GiB) allows larger segments, and larger segments mean that more of the repeat phrases are near repeat phrases which are faster to process than far repeat phrases.

In the third experiment we explore the effect of the technique described in Sect. 6 aiming at reducing the peak disk space usage of the new algorithm. We executed the algorithm on 32 GiB prefixes of two testfiles using 3.5 GiB of RAM and with varying disk space budgets. As shown in Fig. 3, this technique allows reducing the peak disk space usage to very little over what is necessary to store

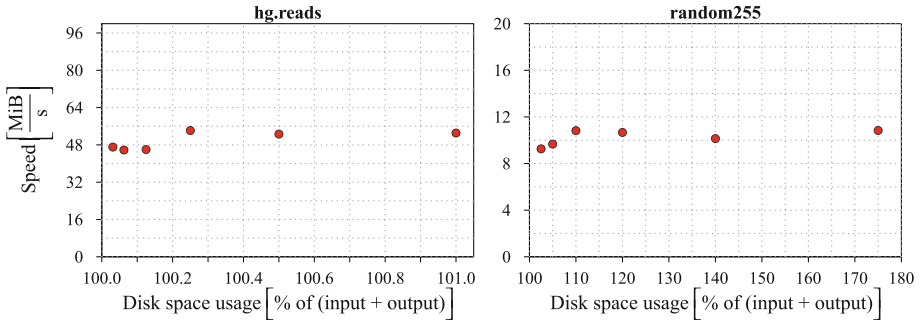


Fig. 3. The effect of disk space budget (see Sect. 6) on the speed of the new external-memory LZ77 decoding algorithm using plain disk I/O. Both testfiles were limited to 32 GiB prefixes and the algorithm was allowed to use 3.5 GiB of RAM. The rightmost data-point on each of the graphs represents a disk space budget sufficient to perform the decoding in one part.

the input parsing and output text and does not have a significant effect on the runtime of the algorithm, even on the incompressible random data.

8 Concluding Remarks

We have described the first algorithms for external memory LZ77 decoding. Our experimental results show that LZ77 decoding is fast in external memory setting too. The state-of-the-art external memory LZ factorization algorithms are more than a magnitude slower than our fastest decoding algorithm, see [11].

References

1. Aggarwal, A., Vitter, J.S.: The input/output complexity of sorting and related problems. *Commun. ACM* **31**(9), 1116–1127 (1988). doi:[10.1145/48529.48535](https://doi.org/10.1145/48529.48535)
2. Badkobeh, G., Crochemore, M., Toopsuwan, C.: Computing the maximal-exponent repeats of an overlap-free string in linear time. In: Calderón-Benavides, L., González-Caro, C., Chávez, E., Ziviani, N. (eds.) *SPIRE 2012*. LNCS, vol. 7608, pp. 61–72. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-34109-0_8](https://doi.org/10.1007/978-3-642-34109-0_8)
3. Brodal, G.S., Katajainen, J.: Worst-case efficient external-memory priority queues. In: Arnborg, S. (ed.) *SWAT 1998*. LNCS, vol. 1432, pp. 107–118. Springer, Heidelberg (1998). doi:[10.1007/BFb0054359](https://doi.org/10.1007/BFb0054359)
4. Dementiev, R., Kettner, L., Sanders, P.: STXXL: standard template library for XXL data sets. *Softw. Pract. Exper.* **38**(6), 589–637 (2008). doi:[10.1002/spe.844](https://doi.org/10.1002/spe.844)
5. Ferrada, H., Gagie, T., Hirvola, T., Puglisi, S.J.: Hybrid indexes for repetitive datasets. *Phil. Trans. R. Soc. A* **372** (2014). doi:[10.1098/rsta.2013.0137](https://doi.org/10.1098/rsta.2013.0137)
6. Ferragina, P., Manzini, G.: On compressing the textual web. In: *Proceedings of 3rd International Conference on Web Search and Web Data Mining (WSDM)*, pp. 391–400. ACM (2010). doi:[10.1145/1718487.1718536](https://doi.org/10.1145/1718487.1718536)

7. Gagie, T., Gawrychowski, P., Kärkkäinen, J., Nekrich, Y., Puglisi, S.J.: A faster grammar-based self-index. In: Dediu, A.-H., Martín-Vide, C. (eds.) LATA 2012. LNCS, vol. 7183, pp. 240–251. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-13089-2_23](https://doi.org/10.1007/978-3-642-13089-2_23)
8. Gagie, T., Gawrychowski, P., Puglisi, S.J.: Faster approximate pattern matching in compressed repetitive texts. In: Asano, T., Nakano, S., Okamoto, Y., Watanabe, O. (eds.) ISAAC 2011. LNCS, vol. 7074, pp. 653–662. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25591-5_67](https://doi.org/10.1007/978-3-642-25591-5_67)
9. Gagie, T., Gawrychowski, P., Kärkkäinen, J., Nekrich, Y., Puglisi, S.J.: A faster grammar-based self-index. In: Dediu, A.-H., Martín-Vide, C. (eds.) LATA 2012. LNCS, vol. 7183, pp. 240–251. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-28332-1_21](https://doi.org/10.1007/978-3-642-28332-1_21)
10. Hoobin, C., Puglisi, S.J., Zobel, J.: Relative Lempel-Ziv factorization for efficient storage and retrieval of web collections. *Proc. VLDB* **5**(3), 265–273 (2011)
11. Kärkkäinen, J., Kempa, D., Puglisi, S.J.: Lempel-Ziv parsing in external memory. In: Proceedings of 2014 Data Compression Conference (DCC), pp. 153–162. IEEE (2014). doi:[10.1109/DCC.2014.78](https://doi.org/10.1109/DCC.2014.78)
12. Kolpakov, R., Bana, G., Kucherov, G.: MREPS: efficient and flexible detection of tandem repeats in DNA. *Nucleic Acids Res.* **31**(13), 3672–3678 (2003). doi:[10.1093/nar/gkg617](https://doi.org/10.1093/nar/gkg617)
13. Kolpakov, R., Kucherov, G.: Finding maximal repetitions in a word in linear time. In: Proceedings of 40th Annual Symposium on Foundations of Computer Science (FOCS), pp. 596–604. IEEE Computer Society (1999). doi:[10.1109/SFFCS.1999.814634](https://doi.org/10.1109/SFFCS.1999.814634)
14. Kolpakov, R., Kucherov, G.: Finding approximate repetitions under hamming distance. *Theor. Comput. Sci.* **303**(1), 135–156 (2003). doi:[10.1016/S0304-3975\(02\)00448-6](https://doi.org/10.1016/S0304-3975(02)00448-6)
15. Kreft, S., Navarro, G.: LZ77-like compression with fast random access. In: Proceedings of 2010 Data Compression Conference (DCC), pp. 239–248 (2010). doi:[10.1109/DCC.2010.29](https://doi.org/10.1109/DCC.2010.29)
16. Kreft, S., Navarro, G.: Self-indexing based on LZ77. In: Giancarlo, R., Manzini, G. (eds.) CPM 2011. LNCS, vol. 6661, pp. 41–54. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-21458-5_6](https://doi.org/10.1007/978-3-642-21458-5_6)
17. Kuruppu, S., Puglisi, S.J., Zobel, J.: Relative Lempel-Ziv compression of genomes for large-scale storage and retrieval. In: Chavez, E., Lonardi, S. (eds.) SPIRE 2010. LNCS, vol. 6393, pp. 201–206. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-16321-0_20](https://doi.org/10.1007/978-3-642-16321-0_20)
18. Lempel, A., Ziv, J.: On the complexity of finite sequences. *IEEE Trans. Inf. Theor.* **22**(1), 75–81 (1976). doi:[10.1109/TIT.1976.1055501](https://doi.org/10.1109/TIT.1976.1055501)
19. Vitter, J.S.: Algorithms and data structures for external memory. *Found. Trends Theoret. Comput. Sci.* **2**(4), 305–474 (2006). doi:[10.1561/0400000014](https://doi.org/10.1561/0400000014)
20. Ziv, J., Lempel, A.: A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theor.* **23**(3), 337–343 (1977). doi:[10.1109/TIT.1977.1055714](https://doi.org/10.1109/TIT.1977.1055714)