

Massively-Parallel Lossless Data Decompression

Evangelia Sitaridi*, Rene Mueller†, Tim Kaldewey‡, Guy Lohman† and Kenneth A. Ross*

*Columbia University: {eva, kar}@cs.columbia.edu

†IBM Almaden Research: {muellerr, lohman}@us.ibm.com

‡IBM Watson: tkaldew@us.ibm.com

Abstract—Today’s exponentially increasing data volumes and the high cost of storage make compression essential for the Big Data industry. Although research has concentrated on efficient compression, fast decompression is critical for analytics queries that repeatedly read compressed data. While decompression can be parallelized somewhat by assigning each data block to a different process, break-through speed-ups require exploiting the massive parallelism of modern multi-core processors and GPUs for data decompression *within a block*. We propose two new techniques to increase the degree of parallelism during decompression. The first technique exploits the massive parallelism of GPU and SIMD architectures. The second sacrifices some compression efficiency to eliminate data dependencies that limit parallelism during decompression. We evaluate these techniques on the decompressor of the DEFLATE scheme, called Inflate, which is based on LZ77 compression and Huffman encoding. We achieve a $2\times$ speed-up in a head-to-head comparison with several multi-core CPU-based libraries, while achieving a 17% energy saving with comparable compression ratios.

I. INTRODUCTION

With exponentially increasing data volumes and the high cost of enterprise data storage, data compression has become essential for reducing storage costs in the Big Data era. There exists a plethora of compression techniques, each having a different trade-off between its compression ratio (compression efficiency) and its speed of execution (bandwidth). Most research so far has focused on the speed of compressing data as data is loaded into an information system, but the speed of decompressing that data can be even more important for Big Data workloads – usually data is compressed only once at load time but repeatedly decompressed as it is read when executing analytics or machine learning jobs. Decompression speed is, therefore, crucial to minimizing the response time of these applications, which are typically I/O-bound.

In an era of flattening processor speeds, parallelism provides our best hope of speeding up any process. In this work, we leverage the massive parallelism provided by Graphics Processing Units (GPUs) to accelerate decompression. GPUs have already been successfully used to accelerate several other data processing problems, while providing a better Performance/Watt ratio than conventional CPUs, as well. However, accelerating decompression on massively parallel processors like GPUs presents new challenges. Straightforward parallelization methods, in which the input block is simply split into many, much smaller data blocks that are then processed independently by each processor, result in poorer compression efficiency, due to the reduced redundancy in the smaller blocks, as well as diminishing performance returns caused by

per-block overheads. In order to exploit the high degree of parallelism of GPUs, with potentially thousands of concurrent threads, our implementation needs to take advantage of both intra-block parallelism and inter-block parallelism. For intra-block parallelism, a group of GPU threads decompresses the same data block concurrently. Achieving this parallelism is challenging due to the inherent data dependencies among the threads that collaborate on decompressing that block.

In this paper, we propose and evaluate two approaches to address this intra-block decompression challenge. The first technique exploits the SIMD-like execution model of GPUs to coordinate the threads that are concurrently decompressing a data block. The second approach avoids data dependencies encountered during decompression by proactively eliminating performance-limiting back-references during the compression phase. The resulting speed gain comes at the price of a marginal loss of compression efficiency. We present **Gompresso/Bit**, a parallel implementation of an Inflate-like scheme [1] that aims at high decompression speed and is suitable for massively-parallel processors such as GPUs. We also implement **Gompresso/Byte**, based on LZ77 with byte-level encoding. It trades off slightly lower compression ratios for an average $3\times$ higher decompression speed.

In summary, the contributions of this paper are:

- A technique to achieve massive intra-block parallelism during decompression by exploiting the SIMD-like architecture of GPUs.
- Improved intra-block parallelism by eliminating data dependencies during compression.
- An evaluation of the impact of both techniques on compression ratio and speed.
- Comparisons of Gompresso’s decompression speed and energy efficiency on the Tesla K40 GPU against several state-of-the-art multi-core CPU libraries, showing that **Gompresso/Bit** is $2\times$ faster while achieving a 17% energy saving.

Section II discusses related work. In Section III we analyze how Gompresso parallelizes decompression to harvest the massive parallelization of GPUs. Section IV focuses on the alternative dependency resolution strategies we designed for LZ77. Section V presents the experimental results of tuning and comparing Gompresso against state-of-the-art parallel CPU libraries. Finally, in Section VI we conclude and suggest some interesting directions for future work.

II. RELATED WORK

In this section, we discuss related work in parallelizing decompression. Although there are numerous compression schemes, we focus in this section on just the parallelization attempts of the best-known compression schemes. In the extended version of the paper [2], we provide a brief description of the relevant background on data compression and a short introduction to relevant aspects of modern GPU architectures.

a) *Parallel CPU Implementations*: A parallel implementation for CPUs of gzip compression in the pigz library [3] achieves a linear speed-up of compression with the number of CPU cores. Decompression in pigz, however, has to be single-threaded because of its variable-length blocks. Another CPU compression library, pbzip2 [4], parallelizes the set of algorithms implemented by the bzip2 scheme. The input is split into data blocks that can be compressed and decompressed in parallel. As already described in the Introduction, this inter-block parallelism alone is insufficient and results in poor performance on GPUs.

b) *Hardware-Accelerated Implementations*: Parallelizing compression schemes within a block is a bigger challenge for massively-parallel processors. For example, the GPU implementation of bzip2 did not improve performance against the single-core CPU bzip2 [5]. The major bottleneck was the string sort required for the Burrow-Wheeler-Transform (BWT) compression layer. Future compressor implementations could be accelerated by replacing string sort with suffix array construction [6], [7], [8].

Most research has focused on accelerating compression, rather than decompression [9]. Here, we address the thread dependencies that limit the parallelism of the LZ77 decompression. In our implementation each thread writes multiple back-reference characters at a time, avoiding the high per character cost. A parallel algorithm for LZ decompression, depending on the type of data dependencies, does not guarantee efficient GPU memory access[10]. Huffman encoding is typically added to improve the compression ratio[11]. However, decoding is hard to parallelize because it has to identify codeword boundaries for variable-length coding schemes. Our parallel decoding method splits data blocks into smaller sub-blocks to increase the available parallelism. We trade-off a little of compression efficiency but only make only one pass over the encoded data. Alternative parallel decoding algorithms do not affect the compression ratio but they require multiple passes to decode the data for BWT decompression: A first pass to determine the codeword boundaries and a second for the actual decoding [7].

Simpler compression schemes have been implemented on GPUs in the context of a database system [12], but while these algorithms achieve good compression ratios for database columns, they are not efficient for Big Data workloads that might be unstructured. FPGAs and custom hardware have also been used to accelerate compression, resulting in high speed-ups [13], [14]. However, these hardware devices have very different characteristics and constraints than GPUs, so their parallelization techniques generally aren't applicable to GPUs.

III. GOMPRESSO OVERVIEW

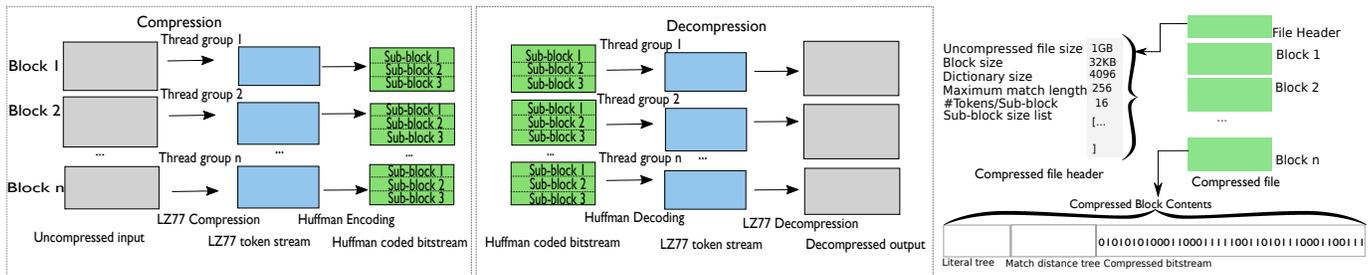
Here, we provide an overview of **Gompresso**, which exploits parallelism between and also within data blocks. The main design goal for **Gompresso** is high decompression speed, while maintaining a “reasonable” compression ratio. **Gompresso** implements both compression and decompression, and defines its own file format. Figure 1a gives an overview of the **Gompresso** compression and decompression algorithms. In this paper, we focus on the parallel decompression.

Gompresso/Byte can combine decoding and decompression in a single pass because of its fixed-length byte-level coding scheme. Figure 1b shows the structure of the compressed file format in detail. The token streams can be read directly from the compressed output. **Gompresso/Bit** uses a variable-length coding scheme for a higher compression ratio, and therefore needs to first decode the bitstream into a stream of tokens before proceeding with the LZ77 decompression. **Gompresso** assigns a group of GPU threads to collaborate on the Huffman decoding and LZ77 decompression on the independently compressed data blocks. This permits an additional degree of parallelism within data blocks.

1) *Huffman Decoding*: Each thread of a group decodes a different sub-block of the compressed data block. The starting offset of each sub-block in the bitstream is computed from the sub-block sizes in the file header. All sub-blocks of a given data block decode their bitstreams using look-up tables created from the same two Huffman trees for that block and stored in the software-controlled, on-chip GPU memories. We can retrieve the original token symbol with a single lookup in each table, which is much faster than searching through the (more compact) Huffman trees, which would introduce branches and hence divergence of the threads' execution paths. The output of the decoder is the stream of literal and back-references, and is written back to the device memory.

2) *LZ77 Decompression*: Each data block is assigned to a single GPU warp (32 threads operating in lock-step) for decompression. We chose to limit the group size to one warp to take advantage of the efficient voting and shuffling instructions within a warp. Larger thread groups would require explicit synchronization and data exchange via on-chip memory. We found that the potential performance gain by the increased degree of parallelism is canceled out by this additional coordination overhead.

We first group consecutive literals into a single literal string. We further require that a literal string is followed by a back-reference and vice versa, similar to the LZ4 [15] compression scheme. A literal string may have zero length if there is no literal token between two consecutive back-references. A pair consisting of a literal string and a back-reference is called a *sequence*. We assign each sequence to a different thread (see Figure 2). In our experiments, we found that this grouping results in better decompression speed since it not only assigns each thread a larger unit of work but its uniformity suits the lock-step execution model of the GPU. All threads in the warp concurrently alternate between executing instructions for



(a) Compression and decompression algorithm parallelization in Gompesso.

(b) Gompesso file format.

Fig. 1: Gompesso overview.

string literals and for back references. For each sequence, each thread performs: (a) read its sequence from device memory and compute the start position of its string literal, (b) determine the output position of its literal, and copy its string literal to the output buffer, and (c) resolve and write its back-reference. We now describe each step in more detail:

a) *Reading sequences*: Each warp uses the block offset to determine the location of the first decoded token in the device memory. Each thread in the warp will read a different sequence (see Figure 2). The threads then determine the start location of their literal strings in the token stream by computing an intra-warp exclusive prefix sum from the literal lengths of their sequences. We use NVIDIA’s shuffle instructions to efficiently compute this prefix sum without memory accesses, a common GPU technique.

b) *Copying literal strings*: Next, the threads compute write positions in the decompressed output buffer. Since all blocks, except potentially the last, have the same uncompressed size, the threads can also easily determine the start position of their block in the uncompressed output stream. The start position of each thread’s literal string is determined by a second exclusive prefix sum, which is then added to the start position of the block. This prefix sum is computed from the total number of bytes that each thread will write for its sequence, i.e., the length of its literal string positions are

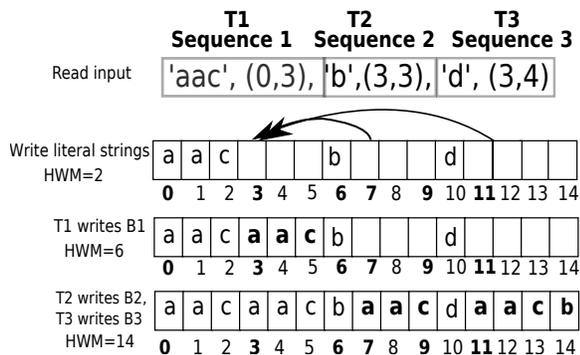


Fig. 2: Multi-Round Resolution (MRR) execution

determined from the two prefix sums, the threads can copy the literal strings from the token stream into the output buffer.

c) *Copying back-references*: This is the most challenging step for parallel decompression because of the data dependencies between threads in a warp. These dependencies arise when a back-reference points to another back-reference, and thus cannot be resolved before the former has been resolved. We address these nested back-references in Section IV. After all the back-references have been resolved, the warp continues with the next 32 sequences.

IV. DATA DEPENDENCIES IN NESTED BACK-REFERENCES

Before processing a back-reference, the data pointed to by this reference needs to be available in the output. This introduces a data dependency and stalls threads with dependent references until the referenced data becomes available. The problem is illustrated in Figure 2. Threads T2 and T3 will have to wait for T1 to finish processing its sequence, because they both have back-references that point into the range that is written by T1. Resolving back-references sequentially would produce the correct output, but would also under-utilize the available thread resources. To maximize thread utilization, we propose two strategies to handle these data dependencies. The first strategy uses warp shuffling and voting instructions to process dependencies as soon as possible, i.e., as soon as all of the referenced data becomes available. The second strategy avoids data dependencies altogether by prohibiting nested back-references during compression.

```

1: function MRR(HWM, read_pos, write_pos, length)
2:   pending  $\leftarrow$  true            $\triangleright$  thread has not written any output
3:   do
4:     if pending and read_pos+length $\leq$ HWM then
5:       copy length bytes from read_pos to write_pos
6:       pending  $\leftarrow$  false
7:     end if
8:     votes  $\leftarrow$  ballot(pending)
9:     last_writer  $\leftarrow$  count_leading_zero_bits(votes)
10:    HWM  $\leftarrow$  shfl(write_pos+length, last_writer)
11:    while votes > 0            $\triangleright$  Repeat until all threads done
12:    return HWM
13: end function

```

Fig. 3: Multi-Round Resolution (MRR) Algorithm

A. Multi-Round Resolution (MRR) of Nested back-references

Figure 3 shows the Multi-Round Resolution (MRR) algorithm for iterative resolution of nested back-references, which is executed by every thread in the warp. We follow the GPU

programming convention in which each of the variables is *thread-private* unless it is explicitly marked as locally or globally shared. The Boolean variable *pending* is initially set on 2 and is cleared once the thread has copied its back-reference to the output (line 6).

Before calling MRR, all threads have written their literal string from their sequence to the output, but no thread in the warp has written a back-reference yet. In order to determine when the referenced data becomes available, the threads keep track of the high-water mark (HWM) position of the output that has been written so far without gaps. A back-reference whose referenced interval is below the HWM can, therefore, be resolved. In each iteration, threads that have not yet written their output use the high-water mark (HWM) to determine whether their back reference can be resolved (line 4). If so, they copy the data from the referenced sequence to the output, and indicate that they completed their work (lines 5 and 6).

The HWM is updated at the end of each iteration. The algorithm determines the last sequence that was completed by the warp, and sets the HWM past the highest write position of that sequence’s back-reference. The threads can determine the last sequence without accessing shared memory by exploiting the warp-voting instruction *ballot* on the *pending* flag (line 8). This produces a 32-bit bitmap that contains the *pending* states of all threads in this warp. Each thread receives this bitmap and then counts the number of leading zeros in the bitmap to determine the ID of the *last_writer* thread that completed the last sequence. A subsequent warp-shuffle instruction broadcasts the new HWM computed by the *last_writer* thread to all other threads in the warp (line 10). The iteration completes when all threads have processed their back-references.

Figure 2 illustrates the execution of MRR. Initially, all threads write in parallel their string of literals. In the next step, T1 copies the back-reference of Sequence 1. In the last step, after Sequence 1 has been processed, the dependencies of T2 and T3 are satisfied, so both threads can proceed to copy their back-references.

At least one back-reference is resolved during each iteration which guarantees termination of the algorithm. The degree of achievable parallelism depends on nesting of back-references. As soon as the referenced ranges falls below the HWM they can be resolved simultaneously. Back-references that do not depend on data produced by other back-references from the same warp can be resolved in one round leading to maximum parallelism of the warp. In the worst-case scenario, all but one back-reference depends on another back-reference in the same warp. MRR then leads to sequential execution. The next section describes a strategy that avoids this scenario.

B. Dependency Elimination (DE)

In this strategy, we trade off some compression efficiency to avoid MRR’s run-time cost of iteratively detecting and resolving dependencies during decompression. During compression, we prohibit nested back-references that would create data dependencies within the same warp. This does not eliminate

```

1: pos ← 0
2: while pos < blocksize do
3:   warpHWM ← pos
4:   s ← 0
5:   literal_str ← ""
6:   while s < 32 do
7:     match ← find_match_below_hwm(dict, input,
8:                                   warpHWM)
9:     if match found then
10:      emit_sequence((literal_str, match))
11:      update_dictionary_with_backref(dict, match)
12:      pos ← pos + match.length
13:      s ← s + 1
14:      literal_str ← ""
15:     else
16:      b ← get next byte from input
17:      literal_str ← literal_str | b
18:      update_dictionary_with_literal_byte(dict, b)
19:      pos ← pos + 1
20:     end if
21:   end while
22: end while

```

Fig. 4: Modified LZ77 compression algorithm with DE

all nested back-references, only those that would depend on other back-references *within the same warp*. Prohibiting these same-warp back-references generally results in a slightly lower compression ratio and more effort during compression, due to the additional checking and bookkeeping. As we will show in Section V, the degradation in compression ratio and compression speed is acceptable. In return, however, we get a 2–3× gain in decompression speed.

Dependency elimination works as follows: For every group of 32 sequences that will eventually be decompressed by the same warp of threads, we only look for dictionary matches below a certain warp high-water mark (*warpHWM*). By choosing the *warpHWM* to be the cursor position in the input that has been completed previously by the warp, we avoid back-references that would otherwise lead to data dependencies. Figure 4 shows the modified LZ77 compression algorithm. The *warpHWM* is updated only after a group of 32 sequences have been completely processed (line 3). Threads that cooperate in the compression perform the string matching in parallel in

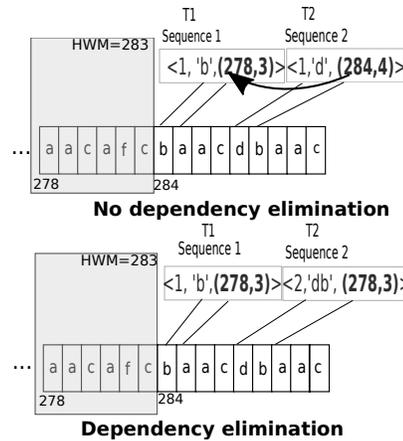


Fig. 5: Resulting token stream without and with DE

`find_match_below_hwm` (line 8). They only look for a match below the current warpHWM. If no match is found, the next input byte is added to the literal string (line 17) and to the dictionary (line 18). Otherwise, if a match is found, the thread closes and emits the output sequence comprising the current literal string and the found match as a back-reference (line 10). Then the dictionary is updated with the found match. The variable “pos” keeps track of the cursor position in the processed input. Figure 5 illustrates the algorithm with an example. The dependency of T2 on T1 is avoided by choosing a shorter match in the back-references for Sequence 2.

Since our **Gompresso** work is focused on decompression, our implementation of the compressor is not as highly optimized as the most commonly used data compression libraries. We decided to also implement the DE algorithm in the LZ4 compression library (CPU-only) [15] to measure the impact that the dependency elimination has on compression speed and the resulting compression ratio. In addition to the DE algorithm itself, we also had to implement the logic for `find_match_below_hwm()` (line 8) by modifying the match-finding component in the LZ4 library so that it only returns matches below a certain HWM.

V. EXPERIMENTAL EVALUATION

We evaluate **Gompresso** using two different datasets. The first is a 1 GB XML dump of the English Wikipedia [16]. The second dataset is the “Hollywood-2009” sparse matrix stored as a 0.77 GB Matrix Market file [17]. Both sets are highly compressible. For comparison, the `gzip` tool achieves a compression ratio of 3.09:1 for the former and 4.99:1 for the latter, using the default compression level setting (-6). The performance measurements are conducted on a dual-socket system with two Intel E5-2620 v2 CPUs, 2×6 cores running 24 hardware threads. We add an NVIDIA Tesla K40 with 2,880 CUDA cores to the system for the GPU measurements. The device is connected via a PCI Express (PCIe) 3.0 x16 link with a nominal bandwidth of 16 GB/sec in each direction. We report bandwidth numbers that include PCIe transfers. When the PCIe bandwidth becomes the bottleneck, we report the bandwidth with input and output data residing in the GPU’s device memory. ECC is turned on in our measurements. We determine the decompression bandwidth as the ratio of the size of the *uncompressed data* over the total processing time. Unless otherwise noted, we are using a data block size of 256 KB and a sliding window of 8 KB. For compression, we look at the next 64 bytes in the input for each match search in the 8 KB window. To facilitate parallel Huffman decoding in **Gompresso/Bit**, we split the sequence stream into sub-blocks that are 16 sequences long.

Performance Impact of Nested back-references: We first focus on just the LZ decompression throughput of **Gompresso/Byte**, i.e., with no entropy coding, for different resolution strategies in Figure 6. *Sequential Copying (SC)* is our baseline resolution strategy, in which threads copy their back-references in a sequential order without intra-block parallelism. The figure shows that Dependency Elimination (DE)

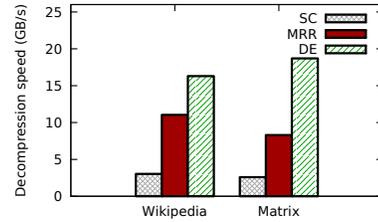


Fig. 6: Decompression speed of **Gompresso/Byte** (data transfer cost not included)

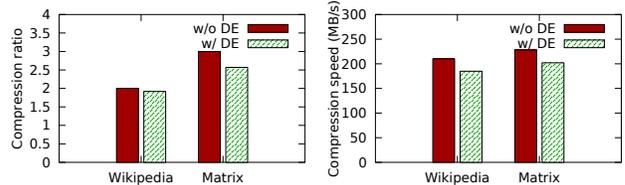


Fig. 7: Degradation in compression ratio and speed

is the fastest strategy for decompression. It is at least $5\times$ faster than SC. We place the compressed input and the decompressed output in device memory in this setup, and ignore PCIe transfers. The figure shows that the decompression throughput is higher than the theoretical maximal bandwidth of the PCIe link. As expected, Multi-Round Resolution (MRR) performs better than SC due to the higher degree of parallelism, while DE outperforms MRR because it achieves an even higher degree of parallelism.

Impact of DE on Compression Ratio and Speed: Figure 7 shows the degradation in compression ratio and compression speed when eliminating dependencies using the Dependency Elimination (DE) algorithm we implemented by modifying the LZ4 library. The maximum degradation is 13% in compression speed and 19% in compression ratio, which is acceptable when we are aiming at fast decompression. In the remaining experiments, we use DE for decompression.

GPU vs. Multi-core CPU Performance: Lastly, we compare the performance of **Gompresso** to state-of-the-art parallel CPU libraries regarding decompression speed and overall energy consumption. We used a power meter to measure energy consumption at the wall socket. For CPU-only environments, we physically removed the GPUs from our server to avoid including the GPU’s idle power. We parallelized the single-threaded implementations of the CPU-based state-of-the-art compression libraries by splitting the input data into equally-sized blocks that are then processed by the different cores in parallel. We chose a block size of 2 MB, as this size resulted in the highest decompression speeds for the parallelized libraries. Once a thread has completed decompressing a data block, it immediately processes the next block from a common queue. This balances the load across CPU threads despite input-dependent processing times for the different data blocks.

Figures 8a and 8b show the trade-offs between decompression speed and compression ratio. In addition to the measurements of our **Gompresso** system, we include the

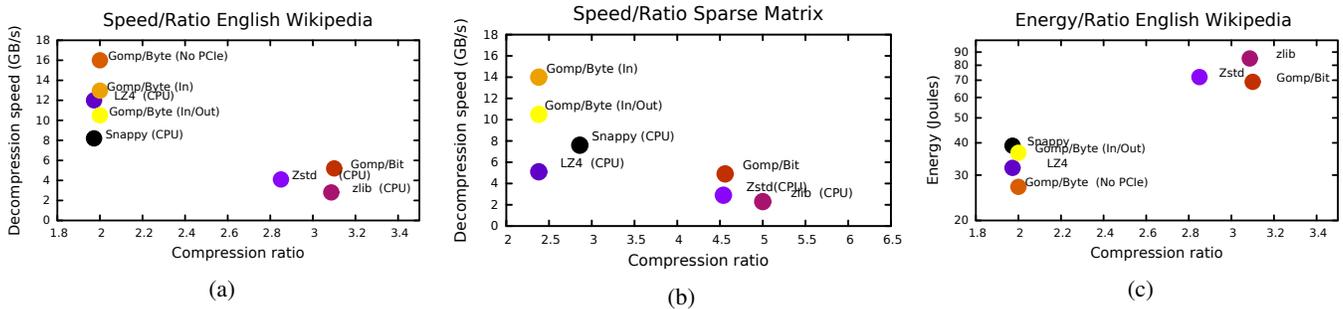


Fig. 8: Gompesso evaluation against multi-core state-of-the-art CPU libraries

performance of two byte-level compression libraries (LZ4, Snappy) and for two libraries using bit-level encoding (gzip, zlib) for comparison. Zstd implements a different coding algorithm on top of LZ-compression that is typically faster than Huffman decoding, and we include it in our measurements for completeness [18]. zlib implements the DEFLATE scheme for the CPU. For the GPU measurements, we show the end-to-end performance, including times for: (a) both compressed input and uncompressed output over PCIe, marked (*In/Out*); (b) only the input transfers, marked as (*In*); and (c) ignoring data transfers altogether, marked as *No PCIe*.

For **Gompesso/Byte**, PCIe transfers turned out to be the bottleneck. In separate bandwidth tests, we were able to achieve a PCIe peak bandwidth of 13 GB/sec. **Gompesso/Bit**, though not PCIe-bound, is still $2\times$ faster than zlib and **Gompesso/Byte** is $1.35\times$ faster than LZ4. For the matrix dataset, the decompression speed of **Gompesso/Bit** is around $2\times$ faster than zlib. There is around 9% degradation in compression ratio because we use limited-length Huffman coding. Although it lowers the compression efficiency, it enables us to fit more Huffman decoding tables into the on-chip memory.

Finally, we compare the energy consumed to decompress the Wikipedia dataset. In general, faster decompression on the same hardware platform results in improved energy efficiency. This is because the power drawn at the system level, i.e., at the wall plug, does not differ significantly for different algorithms. More interesting is the energy efficiency when comparing different implementations on different hardware platforms, e.g., a parallel CPU vs. a GPU solution. Figure 8c shows the overall energy consumption versus the compression ratio for **Gompesso** and a number of parallelized CPU-based libraries. **Gompesso/Bit** consumes around 17% less energy than the parallel zlib library. It also has similar energy consumption to Zstd, which implements a faster coding algorithm.

VI. CONCLUSIONS AND FUTURE WORK

Here, we developed techniques within our compression framework, Gompesso, for massively parallel decompression using GPUs. We presented one solution for parallelizing Huffman decoding by using parallel sub-blocks, and two techniques to resolve back-references in parallel. The first technique iteratively resolves back-references and the second eliminates data dependencies during compression that will stall parallelism among collaborating threads concurrently decom-

pressing that set of sub-blocks. Gompesso decompresses two real-world datasets $2\times$ faster than the state-of-the-art block-parallel variant of zlib running on a modern multi-core CPU, while suffering no more than a 10% penalty in compression ratio. **Gompesso** also uses 17% less energy by using GPUs. Future work includes determining the extent to which our techniques can be applied to alternative coding and context-based compression schemes, and evaluating their performance.

ACKNOWLEDGMENT

Authors Evangelia Sitaridi and Kenneth Ross were partially supported by National Science Foundation grant IIS-1218222 and by an equipment gift from NVIDIA Corporation.

REFERENCES

- [1] P. Deutsch, "DEFLATE compressed data format specification version 1.3," RFC 1951 (Informational), IETF, may 1996.
- [2] E. Sitaridi, R. Mueller, T. Kaldewey, G. Lohman, and K. Ross, "Massively-parallel lossless data decompression," *CoRR*, vol. abs/1606.00519, 2016. [Online]. Available: <http://arxiv.org/abs/1606.00519>
- [3] M. Adler, "Parallel gzip," <http://zlib.net/pigz/>, accessed: 2015-04-02.
- [4] J. Gilchrist and Y. Nikolov, "Parallel BZIP2," <http://compression.ca/pbzip2/>, accessed: 2015-04-02.
- [5] R. A. Patel, Y. Zhang, J. Mak, A. Davidson, and J. D. Owens, "Parallel lossless data compression on the GPU," in *InPar*, May 2012, p. 9.
- [6] M. Deo and S. Keely, "Parallel suffix array and least common prefix for the GPU," in *PPoPP*, 2013, pp. 197–206.
- [7] J. A. Edwards and U. Vishkin, "Parallel algorithms for Burrows–Wheeler compression and decompression," *TCS*, vol. 525, pp. 10 – 22, 2014.
- [8] L. Wang, S. Baxter, and J. D. Owens, "Fast parallel suffix array on the GPU," in *Euro-Par*, 2015, pp. 573–587.
- [9] A. Ozsoy and M. Swany, "CULZSS: LZSS lossless data compression on CUDA," in *CLUSTER*, 2011, pp. 403–411.
- [10] S. D. Agostino, "Speeding up parallel decoding of LZ compressed text on the PRAM EREW," in *SPIRE*, 2000, pp. 2–7.
- [11] A. Ozsoy, D. M. Swany, and A. Chauhan, "Optimizing LZSS compression on GPGPUs," *FGCS*, vol. 30, pp. 170–178, 2014.
- [12] W. Fang, B. He, and Q. Luo, "Database compression on graphics processors," *Proc. VLDB Endow.*, vol. 3, no. 1-2, pp. 670–680, 2010.
- [13] Xilinx, "GUNZIP/ZLIB/Inflate data decompression core," <http://www.xilinx.com/products/intellectual-property/1-79drsh.html>, accessed: 2015-11-14.
- [14] M. S. Abdelfattah, A. Hagiescu, and D. Singh, "Gzip on a chip: High performance lossless data compression on FPGAs using OpenCL," in *IWOCL*, 2014, pp. 4:1–4:9.
- [15] Y. Collet, "LZ4-extremely fast compression," <https://github.com/Cyan4973/lz4>, accessed: 2015-04-02.
- [16] M. Mahoney. Large text compression benchmark. [Online]. Available: <http://mattmahoney.net/dc/enwik9.zip>
- [17] UF sparse matrix collection. [Online]. Available: <http://cise.ufl.edu/research/sparse/MM/LAW/hollywood-2009.tar.gz>
- [18] Y. Collet, "Zstandard-fast and efficient compression algorithm," <https://github.com/Cyan4973/zstd>, accessed: 2015-04-02.