

2009

Accelerating Lossless Data Compression with Graphics Processing Units

R. L. Cloud

M. L. Curry

H. L. Ward

A. Skjellum

P. Bangalore

Follow this and additional works at: <https://digitalcommons.library.uab.edu/inquire>

 Part of the [Higher Education Commons](#)

Recommended Citation

Cloud, R. L.; Curry, M. L.; Ward, H. L.; Skjellum, A.; and Bangalore, P. (2009) "Accelerating Lossless Data Compression with Graphics Processing Units," *Inquire, the UAB undergraduate science research journal*. Vol. 2009: No. 3, Article 22.

Available at: <https://digitalcommons.library.uab.edu/inquire/vol2009/iss3/22>

This content has been accepted for inclusion by an authorized administrator of the UAB Digital Commons, and is provided as a free open access item. All inquiries regarding this item or the UAB Digital Commons should be directed to the [UAB Libraries Office of Scholarly Communication](#).

Accelerating Lossless Data Compression with Graphics Processing Units

R.L. Cloud, M.L. Curry, H.L. Ward, A. Skjellum, P. Bangalore

Abstract

Huffman compression is a statistical, lossless, data compression algorithm that compresses data by assigning variable length codes to symbols, with the more frequently appearing symbols given shorter codes than the less. The work to be presented is a modification of the Huffman algorithm which permits data to be decomposed into independently compressible and decompressible blocks, permitting concurrent compression or decompression on multiple processors. We implemented this modified algorithm on a NVIDIA GPU using the CUDA API as well as on a current Intel chip and the performance results are compared, showing higher performance compression and decompression on the GPU.

Introduction

Lossless data compression is important in application domains and usage environments where bandwidth or storage limitations may negatively impact application or system performance. Generally classifiable into statistical or dictionary methods, lossless data compression algorithms can range widely in compression speed and efficiency (compression factor). Certain algorithms, especially the more efficient, can be quite computationally expensive, and as the data processing needs of current scientific endeavor continue to scale with more rapidity than storage or bandwidth, compression becomes increasingly necessary, but questions remain as to how to accelerate it and how to do so without consuming a large amount of computational resources.

The use of graphics processing units (GPUs) for general purpose computation, i.e. problems outside the graphical domain, is a relatively recent development. First this was achieved through third party toolkits, e.g. Stanford's BrookGPU, but even more recently have GPU manufacturers themselves begun to offer general purpose tools which give the programmer a lower level communion with the chip than earlier GPGPU programming interfaces which are built upon OpenGL and DirectX. One of these, and currently the most prominent, is the Compute Unified Device Architecture (CUDA) from the NVIDIA corporation. The potential benefits of GPUs in general purpose computation are great, but potential must be emphasized, more so even than for parallel programming on the x86. To achieve anywhere near the theoretical maximums in performance on the GPU, the computation patterns underlying a solution's algorithm must be very near to the traditional usage of the GPU; a prospective algorithm's implementation on the GPU should be, in order of importance to performance, highly data parallelizable, logically simple, and have relatively many computations to memory accesses. In essence, to use the GPU to maximum effect, the abstractable computation patterns underlying a solution should

be co-linear to the GPU's original task, graphics rendering. Our problem domain, I/O, while it does not perfectly fit these criteria, has already benefited from GPUs to enhance storage redundancy [5]; we attempt now their utilization in lossless data compression

One major difficulty here in achieving good speedup with slim negative side effects is that lossless data compression algorithms can generally not be, in their unaltered form, thought of as highly parallelizable. Indeed, if one wishes to express these algorithms in parallel, one often needs to consider tradeoffs between compression efficiency and performance. Nevertheless, we hope to effectively demonstrate that it is possible to come to a reasonable middle ground with respect to coding acceleration and efficiency loss.

Huffman Encoding

Statistical methods of data compression perform analysis on the nature of the data to make intelligent decisions about how it can be represented more efficiently in compressed form. The Huffman encoding algorithm falls within this genus and operates by counting the appearance of every distinct symbol in the uncompressed data, then representing the more frequent with shorter codes than the less frequent. Every symbol in the data is replaced with its code, and if the data is non-random, i.e. a few symbols appear with greater frequency than others, compression can be achieved. The Huffman compression algorithm is old by the standards of our science [6], but is still used, and has the attractive quality of being a primitive of several more modern and common algorithms, e.g. Deflate [1] and potentially the algorithm described by Burrows and Wheeler [4].

Parallel Huffman Coding

There is literature on parallel Huffman coding and of varying goals, ranging from the actual construction of Huffman codes in parallel [3], [2], to [7] which addresses details of decomposition for parallel Huffman decoding and demonstrates some moderate decoding speedups while maintaining optimally encoded data by making use of the observation that Huffman codes can frequently synchronize. Because of limitations in our architecture, we must try to create the simplest encoding routine possible. In doing this we make a minor modification to the output of the Huffman algorithm.

A modification is necessary because of the nature of Huffman codes, i.e. they are of a variable length; an encoded data string is composed of these codes packed together in a nature where bit codes can cross byte boundaries. Simple decomposition of the encoded data stream into blocks of static size would result in the practical certainty that decoding would take an erroneous path, which is discussed in some detail in [7]. One counter to

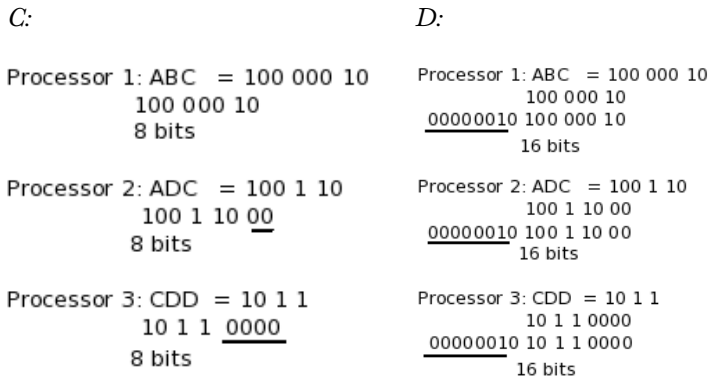
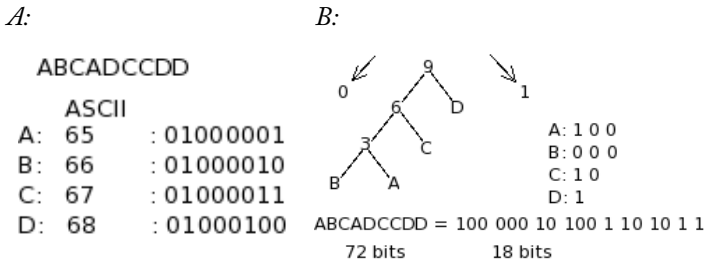


Figure 1: A: The original ASCII encoded string. B: The binary tree and encoded representation of the original string. C: Decomposing the string into three symbol blocks and adding packing bits to the nearest byte. D: The addition of a length delimiter at the start of the block. Single bytes are used for the overhead in the diagrams for simplicity. In our implementation, we pack the block to four bytes and use a four byte integer to represent the block length.

this is to pack the blocks to byte boundaries, introducing some size overhead. One more change is necessary. Because the codes are of variable lengths, even if we encode a constant number of symbols in each block, the resulting length of the encoded block will vary, sometimes dramatically. For this reason, we must encode an indication of where the block starts and ends. Our approach is again simple; at the start of the encoded block we give the length of the block which is known by making an additional pass over the unencoded block and summing the lengths of the code representation of the symbols. Our implementation stores this length as an unencoded four byte integer for simplicity, and because of this and the requirements of our architecture, we pack the blocks to four byte boundaries.

The overhead of our modifications range therefore, from between 32 bits and 63 bits per block, with the variation being because if the size of the encoded block is evenly divisible by four bytes, it is unnecessary to add packing bits to its tail. This overhead naturally becomes less significant as the length of the block is increased, which is indicated in the figure measuring block size against overhead. The time required for summing the block lengths is measurable but undramatic and most noticeable when comparing the runtimes of a sequential block encoder to a sequential traditional(non-block) encoder.

To parallelize decoding, it is sufficient to build a table of offsets into the encoded data from these block length delimiters. The computation threads on the GPU can then index into the encoded data and decode in parallel, storing decoded data in a sequential array indexable by thread and block numbers.

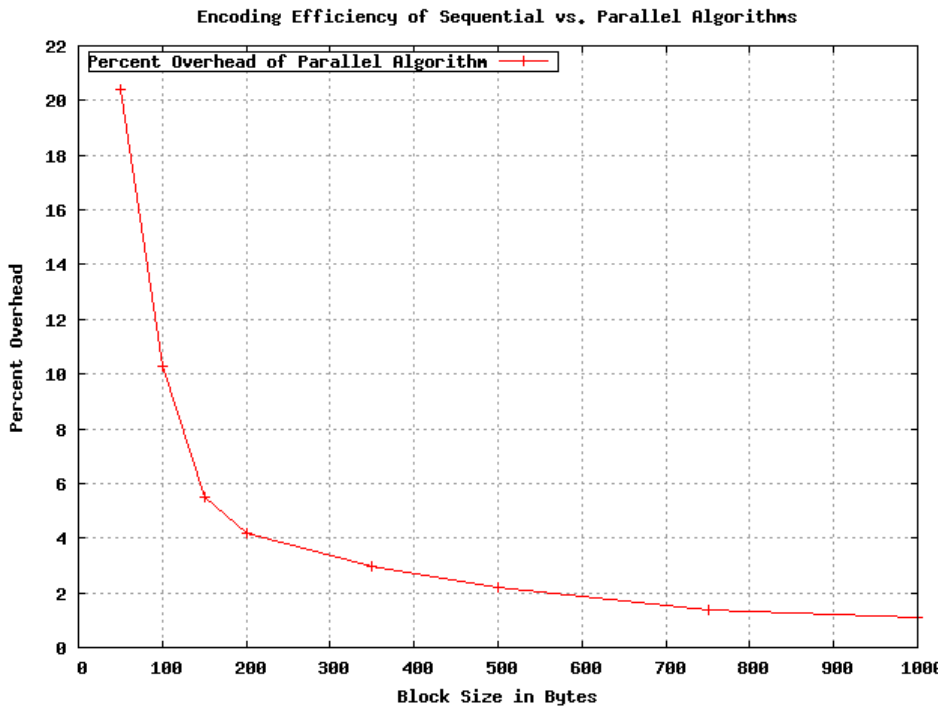


Figure 2: The size overhead of using the parallel Huffman algorithm graphed against the block size. The number of bytes overhead per block remains a constant, so as the block size increases the overhead becomes less significant. At large block sizes, the overhead per block can be less than one percent.

Performance Comparisons

Encoding

Acceleration over our sequential implementation was achieved for both encoding and decoding. This comparison is most meaningful in terms of throughputs, the amount of data which can be encoded or decoded per second. Following is the comparison of our sequential encoder to our parallel GPU encoder and a parallel CPU encoder programmed with OpenMP. The GPU used in these experiments is the NVIDIA GeForce GTX 285 with 240 cores at 1.5 GHz, and the CPU used is the Intel Core i7 Extreme Edition 965 with four cores at 3.2 GHz. Despite the GPU having 60 times the number of cores as our CPU, the differences in throughput between the GPU encoder and the OpenMP encoder are not dramatic. This paradox can be largely resolved by recalling that the architecture of the

GPU was developed for the SIMD, single instruction multiple data, programming model while our CPU was developed with MIMD, multiple instruction multiple data, in mind.

The processors in the GPU are organized into 30 groups of 8 cores. Each group of cores is known as a multiprocessor and contains a single control unit and a small amount of high speed memory shared between the cores in the multiprocessor. The control unit broadcasts an instruction to all the cores, and optimal performance can only be achieved when every core can execute it. If, for example, the instruction is a branching statement, then there is a likelihood that some cores will not follow the jump, and in this case, some cores must remain inactive until they either themselves satisfy the branching instruction or control passes beyond the branching sections of the code. Therefore, in the worst case, when only one core can satisfy the jump and the other seven are left idle, our GPU behaves more like a crippled 30 core shared memory MIMD machine with a slow clock speed and no automatic memory caching. Our encoder consists of complicated branching statements for the bit manipulation which makes

worst case behavior relatively likely. This also illustrates that in heterogeneous programming environments, one must be very aware of the strengths and weaknesses of the various architectures so that programming effort can be directed where benefits are most likely to be found.

Decoding

Our decoding routine consists of reading bits and traversing a binary tree repeatedly for each code string. This contains branching instructions, but markedly fewer than the encoding routine, and the factor of acceleration on the GPU is greater than that of the encoding routine. Also interestingly, the measured increases in throughput from using OpenMP on the CPU, compared to the sequential implementation, are even better than linear by number of cores on the CPU. By launching increasing numbers of threads, we can hide latency by issuing more memory requests. In this way, we saw continued performance improvements through increasing thread counts up to 8. Intel's Hyper-Threading technology assists significantly in this.

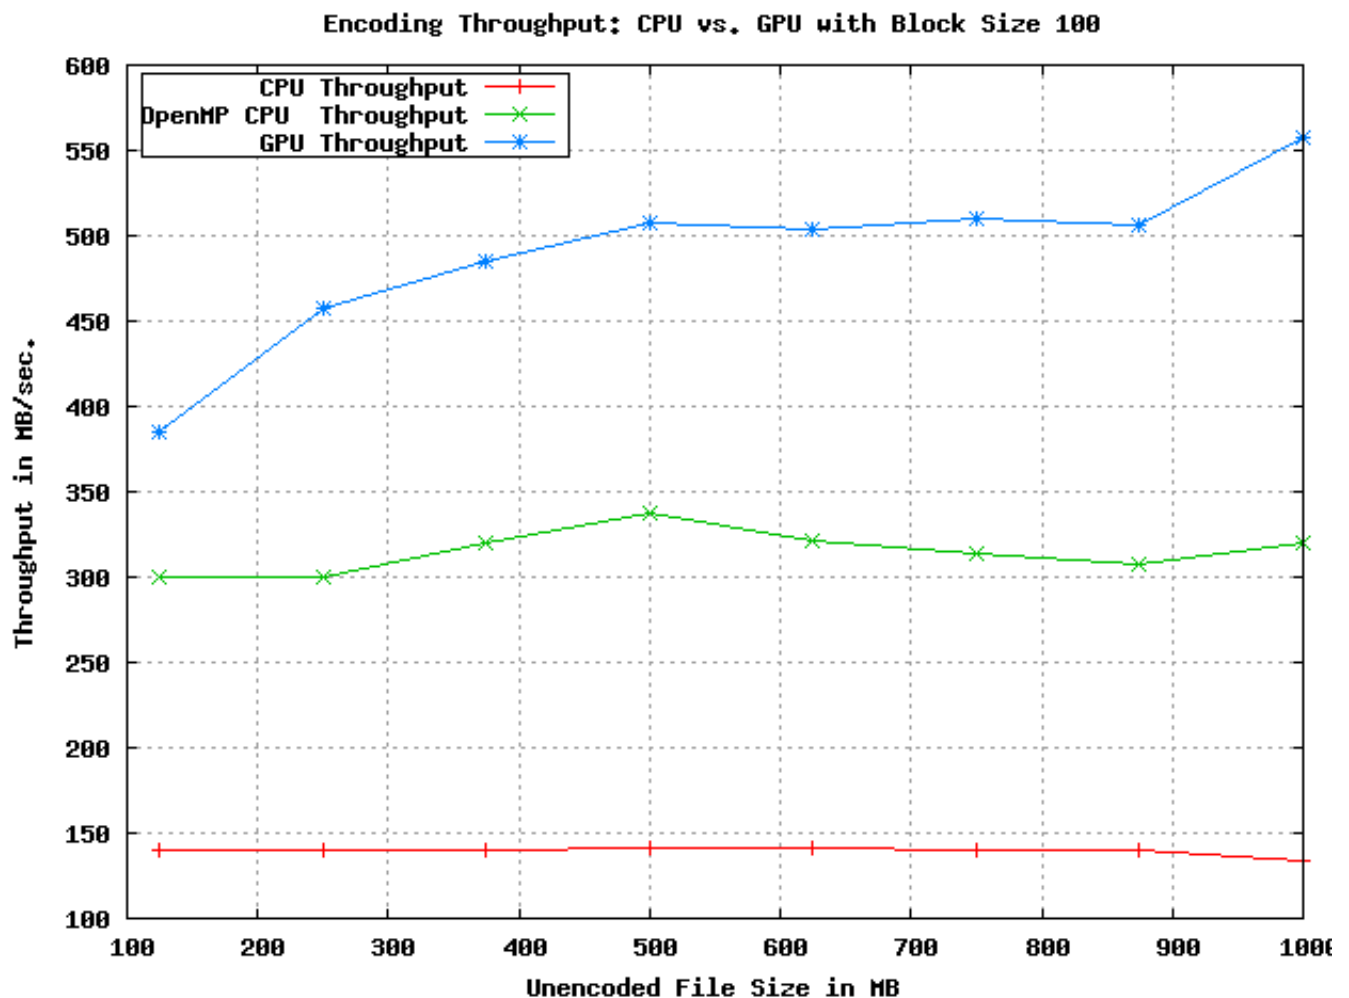


Figure 3: We saw superior performance with the GPU based encoder compared to our multi-core CPU encoder and our single threaded CPU implementation

Conclusions

The data presented here suggests that the strengths of the GPU architecture are robust enough to give performance benefits to applications which, while data parallel, still have a not insignificant level of logical complexity. Optimal use of the GPU's SIMD cores requires the complete elimination of divergence within warps, which, in practicality, requires the complete absence of if statements from the GPU sub-routine; however, sub-optimal performance, through the emulation of MIMD, can still be acceptable. Despite the large number of divergent threads in a warp, our encoder kernel is capable of throughputs, sans memory transfer times to and from the GPU, in excess of 4 GB/sec. Total encoding throughputs using the GPU are weighed down by the need to transfer data to and from the card; however, in an online system, or when encoding very large amounts of data, this could be somewhat ameliorated by using asynchronous data transfers with the GPU to fully exploit bus resources while encoding.

Realistically, current performance levels for our GPU encoder and decoder do not warrant the use of the program as a standalone encoding system. The Huffman algorithm itself is not the best choice for such purposes and even the strengths of the GPU do not make up for the algorithm's deficiencies. However, our encoding system could be used as an auxiliary process to a GPU application. Much greater coding performance than that shown in the above figures could be seen were the data to be encoded already on the GPU.

References

- [1] M. Adler, Deflate algorithm. <http://www.gzip.org/algorithm.txt>.
- [2] M. J. Atallah, S. R. Kosaraju, L. L. Larmore, G. L. Miller, and S.-H. Teng. Constructing Trees in Parallel. In SPAA '89:

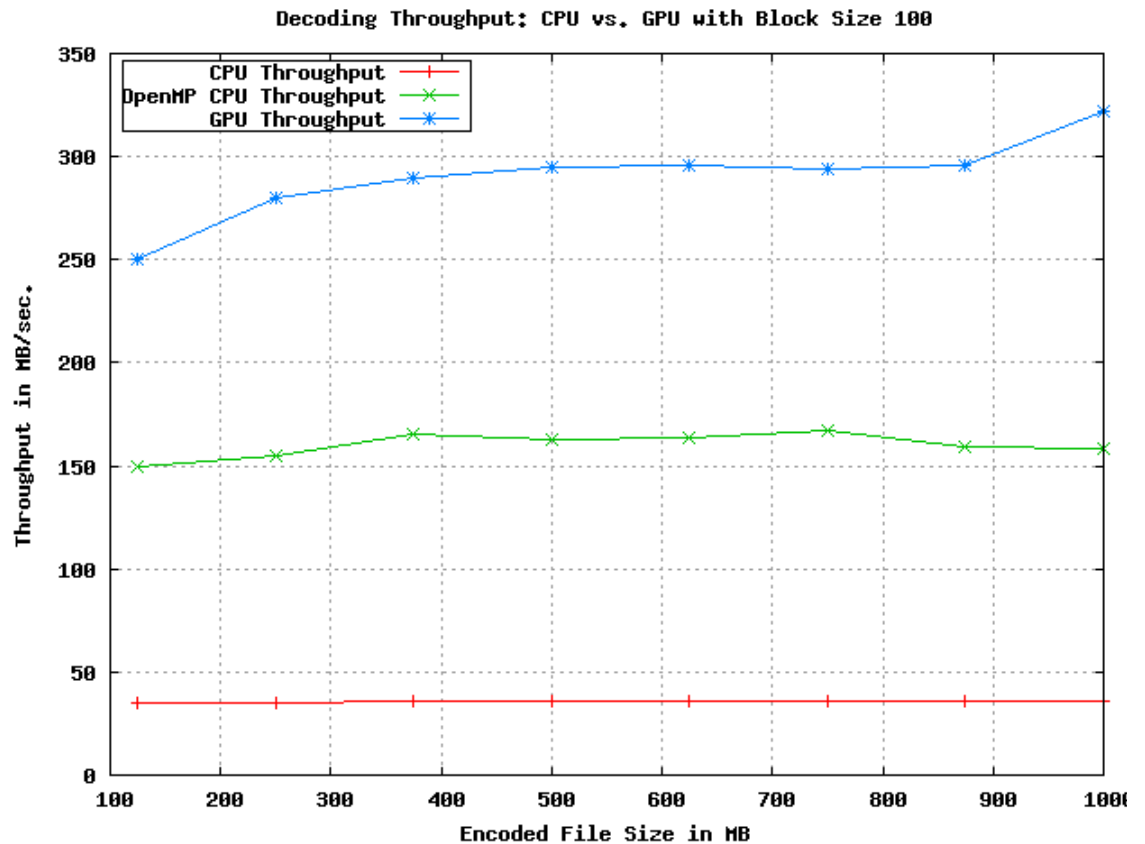


Figure 4: Again, our GPU based decoder gave better performance than both CPU decoders.

Proceedings of the first annual ACM symposium on Parallel algorithms and architectures, pages 421–431, New York, NY, USA, 1989. ACM.

- [3] P. Berman, M. Karpinski, and Y. Nekrich. Approximating Huffman Codes in Parallel. *Journal of Discrete Algorithms*, 5(3):479–490, 2007.
- [4] M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical report, Digital SRC Research Report, 1994.
- [5] M. Curry, A. Skjellum, H. Ward, and R. Brightwell. Accelerating Reed-Solomon Coding in Raid Systems with Gpus. *IEEE International Symposium on Parallel and Distributed Processing*, pages 1–6, 2008.
- [6] D. Huffman. A Method for the Construction of Minimum-Redundancy Codes. *Proceedings of the IRE*, 40:1098–1101, 1952.
- [7] S. T. Klein and Y. Wiseman. Parallel Huffman Decoding with Applications to Jpeg Files. *The Computer Journal*, 46:487–497, 2003.