

# GABAC: an arithmetic coding solution for genomic data

Tom Paridaens, Jan Voges, Mikel Hernaez, Jan Fostier, and Jörn Ostermann

In an effort to provide a response to the ever-expanding generation of genomic data, MPEG (Moving Picture Experts Group), under the auspices of ISO (International Organization for Standardization), is designing a new solution for the representation, compression and management of genomic sequencing data: the MPEG-G standard. This standard specifies an abstract representation of sequencing data and offers support for, among others, multi-dimensional random access, extensive meta-data information, data and privacy protection, data storage and streaming. Part 2 of the MPEG-G standard focuses on specifying the coding of the sequencing data. Note that the standard is a set of specifications (i.e., a book) assembled after an open competition for technologies, and thus, it does not provide, *per se*, actual encoding implementations. This paper discusses the first baseline implementation of an MPEG-G compliance entropy encoder/decoder: GABAC. GABAC combines proven coding technologies, such as context-adaptive binary arithmetic coding (CABAC) [1], binarization schemes, and transformations into one straight-forward solution for the compression of the sequencing data. GABAC will be part of the open-source full MPEG-G encoder/decoder suite Genie, currently under development by the Mitogen initiative<sup>1</sup>.

**Methods:** The MPEG-G standard represents genomic information by splitting the data into a set of descriptor streams. Each descriptor stream has been designed to contain one specific type of data (e.g., mapping positions, quality scores or mismatch information). Thanks to this approach, all data contained in a descriptor stream is homogeneous. Note that other genomic data compressors also employ a similar approach of generating streams containing statistically similar data [2, 3]. This property allows for effective data compression, as redundancies are more probable and value distributions are more predictable.

Given an input stream, the compression process specified in the MPEG-G standard consists of a five-step pipeline: input parsing, (optional) 3-step transformations, symbol binarization, context selection, and CABAC.

In the input parsing step, the descriptor stream is parsed into a stream of symbols. These symbols are then processed by the 3-step transformation step, which converts the symbol stream into one or more transformed sub-streams. Available transformations are as follows. i) run-length encoding, where repetitions of symbols are translated into a symbol sub-stream and a length sub-stream); ii) match coding, an LZ77-style transformation where blocks of symbols are replaced by a pointer and length value, indicating either the position and the length of a perfect match of this block in the previously encoded symbols, or indicating a raw symbol if no match is found; iii) equality coding, where a symbol is replaced by a flag indicating equality of the symbol and its predecessor and a correction symbol, if required; iv) a look-up table transformation; and v) differential coding.

In the first part of the transformation step, the symbols are processed using either run-length encoding, match coding, or equality coding. In the second part of the transformation step,

which is applied to each transformed sub-stream separately, a look-up table transformation can be performed. Finally, in the third part of the transformation step, a differential coding can be applied.

Each transformed sub-stream is then processed separately during the rest of the process. For each sub-stream a binarization algorithm, used for conversion of each symbol into a bit string, is chosen together with a context selection algorithm. In the last step, each bit of the binarization (called a bin) is combined with a context (selected using the context selection algorithm) and both are processed using CABAC.

**Results:** To analyze the performance of the GABAC encoder, a test set of 206 descriptor stream files (cropped to a maximum size of 200 MiB each to emulate random access capabilities) has been selected. These files contain data (such as mapping positions, pairing information, or unmapped reads) generated from items 02, 03, 05, 07, 08, 09, 10, and 20 of the MPEG-G Genomic Information Database<sup>2</sup>. The total uncompressed size is 12.97 GiB.

The performance of GABAC was compared to the codecs that are used in CRAM, i.e. gzip, bzip2, xz, rANS order-0, and rANS order-1. The rANS codecs were extracted from CRAM and are publicly available<sup>3</sup>. The tests were performed in parallel on a set of five servers, each equipped with 2 Intel Xeon E5-2650 v3 CPUs and 128 GiB of RAM, running Ubuntu 14.04.

Table 1 shows the total compressed size and the total encoding and decoding times for each of the tested codecs for all considered streams. Additionally, the table shows the compressed size for the complete test set when, for each file, the codec with the highest compression ratio across the CRAM encoders is selected (CRAM), and across the CRAM encoders and GABAC (CRAM + GABAC). GABAC offers the highest compression ratio across all coding solutions, while being 5.5 times faster than the second best compressor. Additionally, adding GABAC to the CRAM set of encoders offers an additional compression gain of 79 MiB and a speed-up with a factor of 2.4 in compression time.

	Compressed Size (MiB)	Encoding Time	Decoding Time
gzip	3,524	3h 25m 18s	<b>06m 02s</b>
bzip2	3,088	33m 55s	20m 00s
xz	2,944	4h 47m 38s	09m 25s
rANS-0	4,143	<b>06m 01s</b>	07m 08s
rANS-1	3,400	06m 54s	08m 20s
GABAC	<b>2,877</b>	45m 25s	20m 18s
CRAM	2,879	2h 25m 58s	<b>09m 32s</b>
CRAM + GABAC	<b>2,800</b>	<b>1h 01m 17s</b>	20m 08s

Table 1: Total compressed size and decoding & encoding times.

Figure 1 shows the compression ranking for each codec and each descriptor stream file. The x-axis shows the MPEG-G Ge-

<sup>2</sup><https://mpeg.chiariglione.org/standards/MPEG-G/genomic-information-representation/MPEG-G-genomic-information-database>

<sup>3</sup><https://github.com/voges/rans>

<sup>1</sup><https://github.com/mitogen>

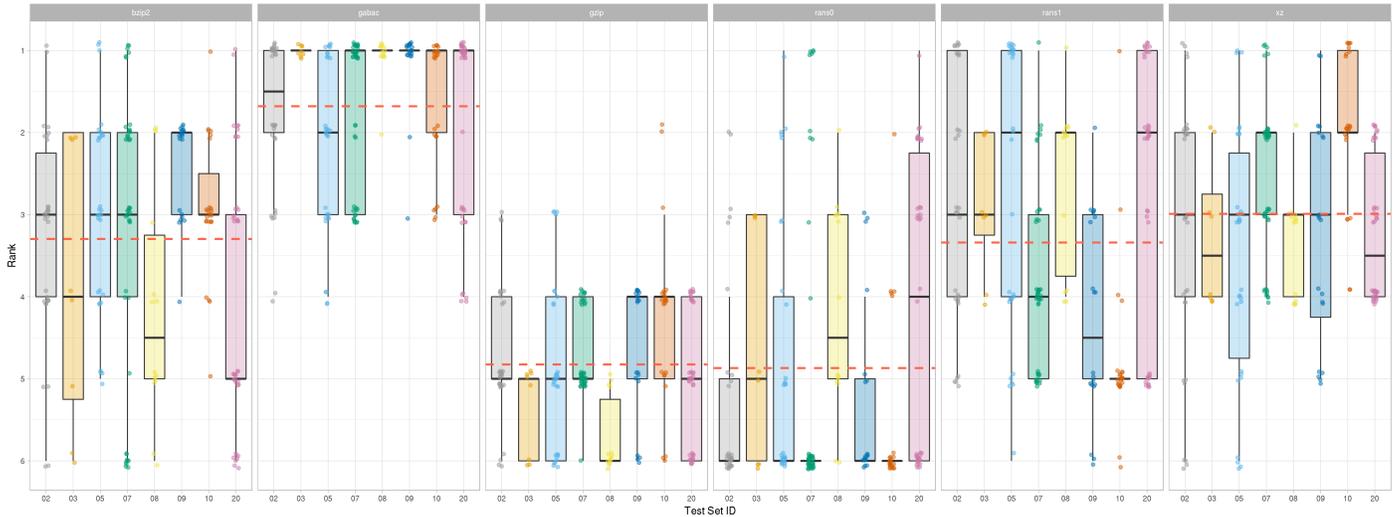


Figure 1: Rank of compression performance. Dots were jittered for clarity.

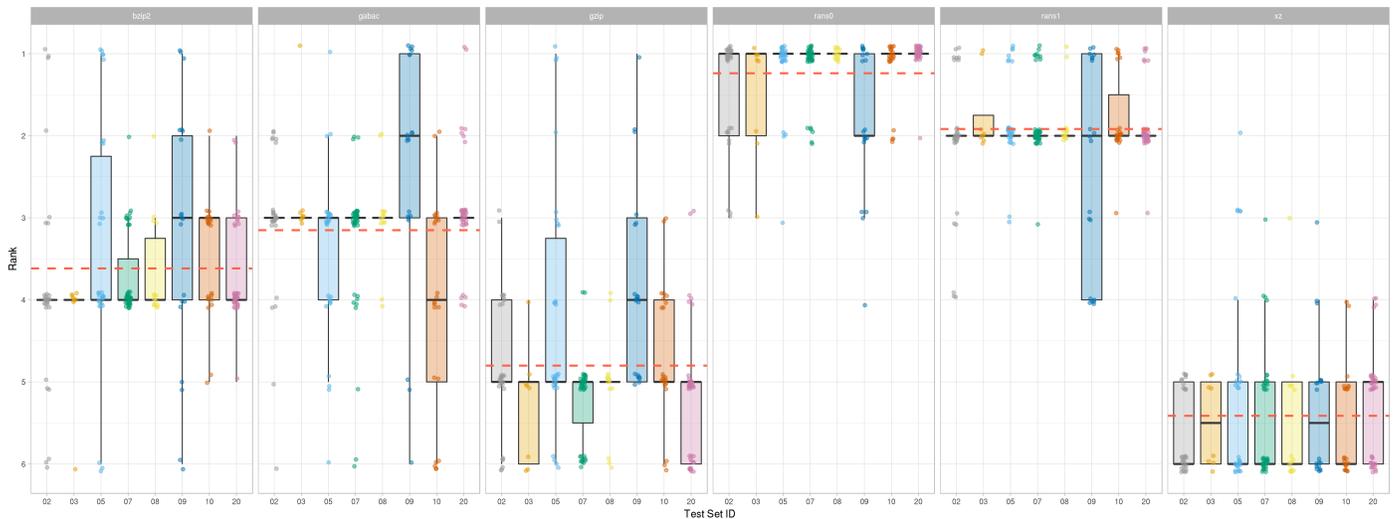


Figure 2: Rank of compression speed performance. Dots were jittered for clarity.

conomic Information Database Test Set ID from which the descriptor stream files were generated. The y-axis denotes the actual ranks. Each dot depicts the ranking of a codec achieved on one specific descriptor stream file. The boxes visualize the following statistics: the black lines denote the medians, the lower and upper hinges correspond to the first and third quartiles, and the whiskers extends to the smallest/largest values no further than 1.5 times the inter-quartile range from the hinges. The dotted red lines denote the mean compression rank for each tool, averaged over all Test Set IDs. Overall, GABAC yields the smallest compressed size among all codecs for 127 out of 206 test items, which is over 60%. As expected gzip never achieves the best compression while bzip2, xz and rANS-1 show good compression results for certain test items. Furthermore, note that GABAC is never ranked in either of the last two positions.

Figure 2 shows the compression speed ranks for each codec and each descriptor stream file. The axis and box plot semantics are the same as those used in Figure 1. GABAC ranks 3rd in terms of mean compression speed and it is only outperformed by the

two RANS variants.

**In conclusion:** This work presents the first implementation of a codec which is compliant with the entropy decoding process as specified in the MPEG-G specifications. We have shown that the proposed implementation already outperforms well established entropy codecs improving in several cases both the compression ratio and the compression speed. We hope that it will serve as a baseline of future implementations as the performance of new implementations is expected to improve over time.

**Acknowledgments:** The authors thank IDLab (Ghent University - imec) for the use of their computing cluster.

- [1] D. Marpe, H. Schwarz and T. Wiegand, "Context-based adaptive binary arithmetic coding in the H.264/AVC video compression standard", *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, no. 7, pp. 620–636, 2003.
- [2] M. Hsi-Yang Fritz, R. Leinonen, G. Cochrane and E. Birney, "Efficient storage of high throughput DNA sequencing data using reference-based compression", *Genome Research*, vol. 21, no. 5, pp. 734–740, 2011.
- [3] J.K. Bonfield, "The Scramble conversion tool", *Bioinformatics*, vol. 30, no. 19, p. 2818, 2014.