


# ROOT I/O compression algorithms and their performance impact within Run 3



Oksana Shadura, Brian Paul Bockelman

University of Nebraska-Lincoln, USA

oksana.shadura@cern.ch, bbockelm@cse.unl.edu


## Introduction

We have performed a survey of the performance of the new compression techniques. We also provide insight into solutions applied to the bottlenecks in compression algorithms for the improved ROOT performance.

## ROOT compression algorithms

- ZLIB** - a LZ77 preprocessor with Huffman coding [ROOT default]
  - other modifications - **zlib-cf** or **Intel zlib**
- LZMA** - a variant of LZ77 with huge dictionary sizes and special support for repeatedly used match distances, whose output is then encoded with a range encoder, using a complex model to make a probability prediction of each bit.
- LZ4** - a LZ77-type compressor with a fixed, byte-oriented encoding and no Huffman coding pass [new ROOT default]
- ZSTD** - a dictionary-type algorithm (LZ77) with large search window and fast implementations of entropy coding stage, using either very fast Finite State Entropy (tANS) or Huffman coding. [Facebook]
- Snappy** - a byte aligned LZ77 algorithm intended for high speed rather than good compression. [Google]
- Old ROOT compression algorithm (backward compatibility)

## ROOT I/O overview



→ We can try to arrange bits of the values to compress sequences for primitive branches efficiently [e.g. Kudu, Apache].

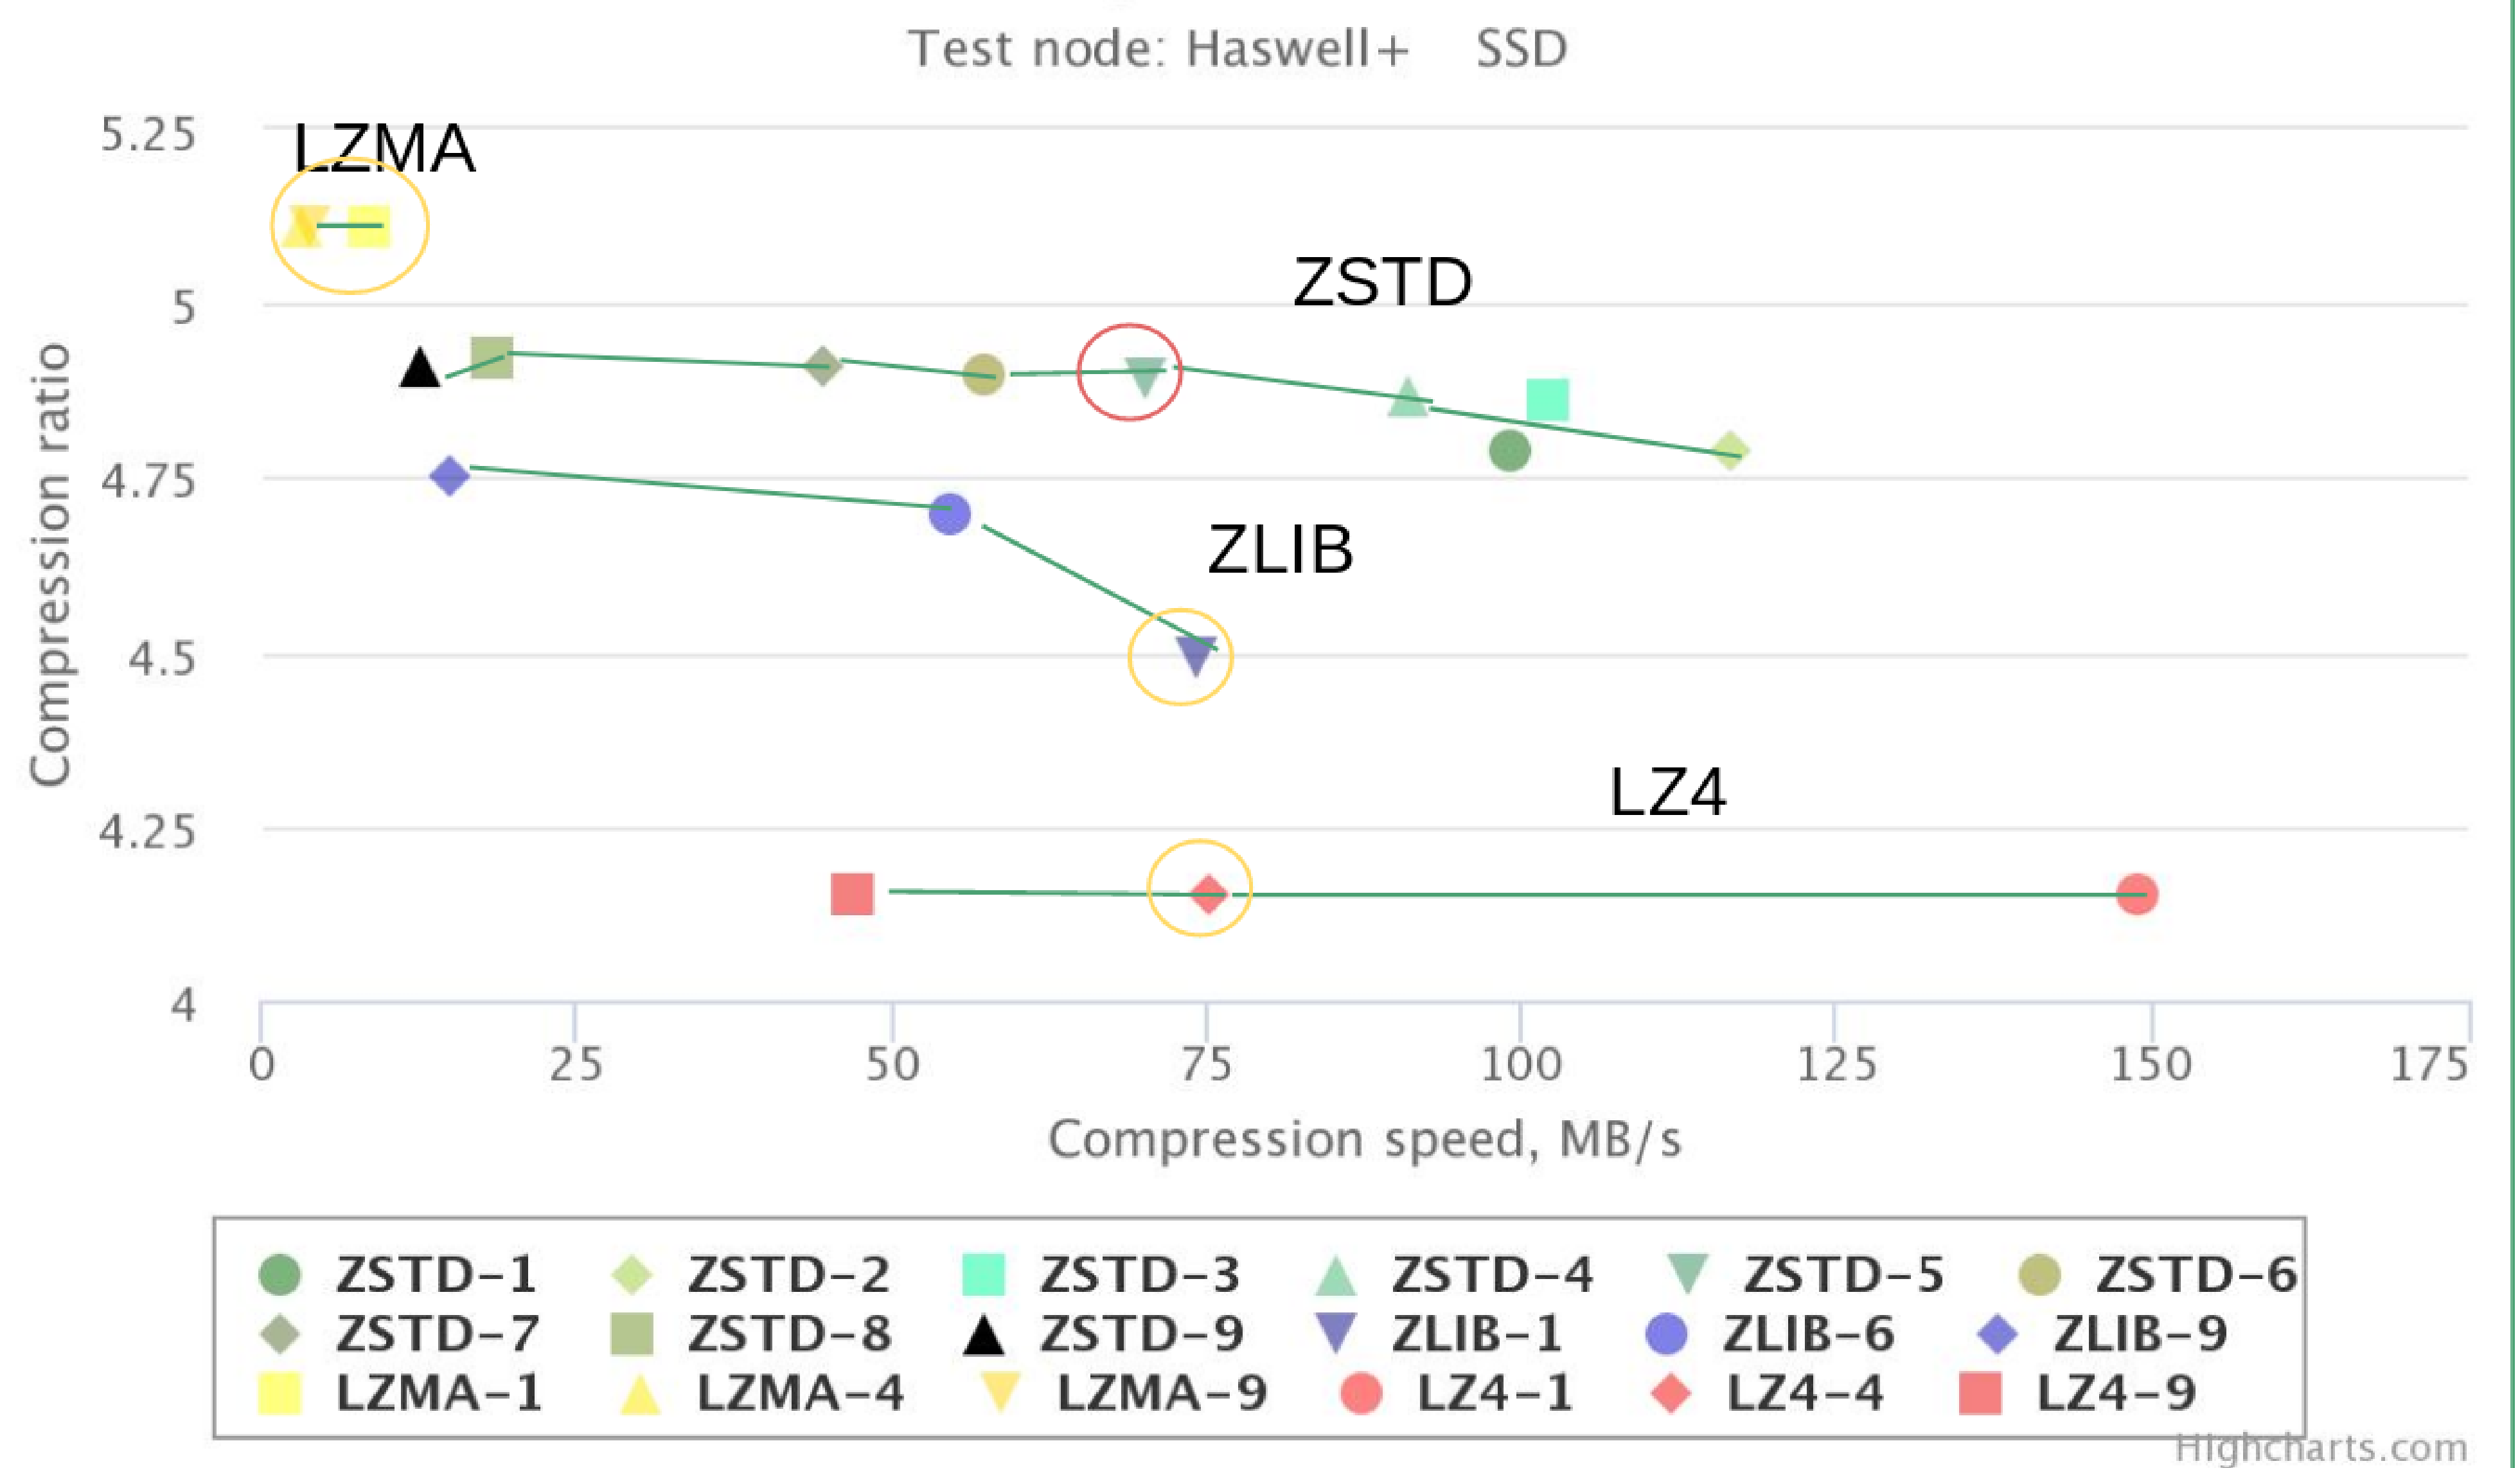
## Compression Pre-conditioners

Algorithms that rearranges typed, binary data for improving compression ratio:

- Shuffle** - integrated byte shuffle preconditioner (available in Blosc)
- BitShuffle** - integrated bit shuffle preconditioner (available in Blosc)

## Comparison of ROOT compression algorithms -compression

### Compression speed vs Compression Ratio for compression algorithms

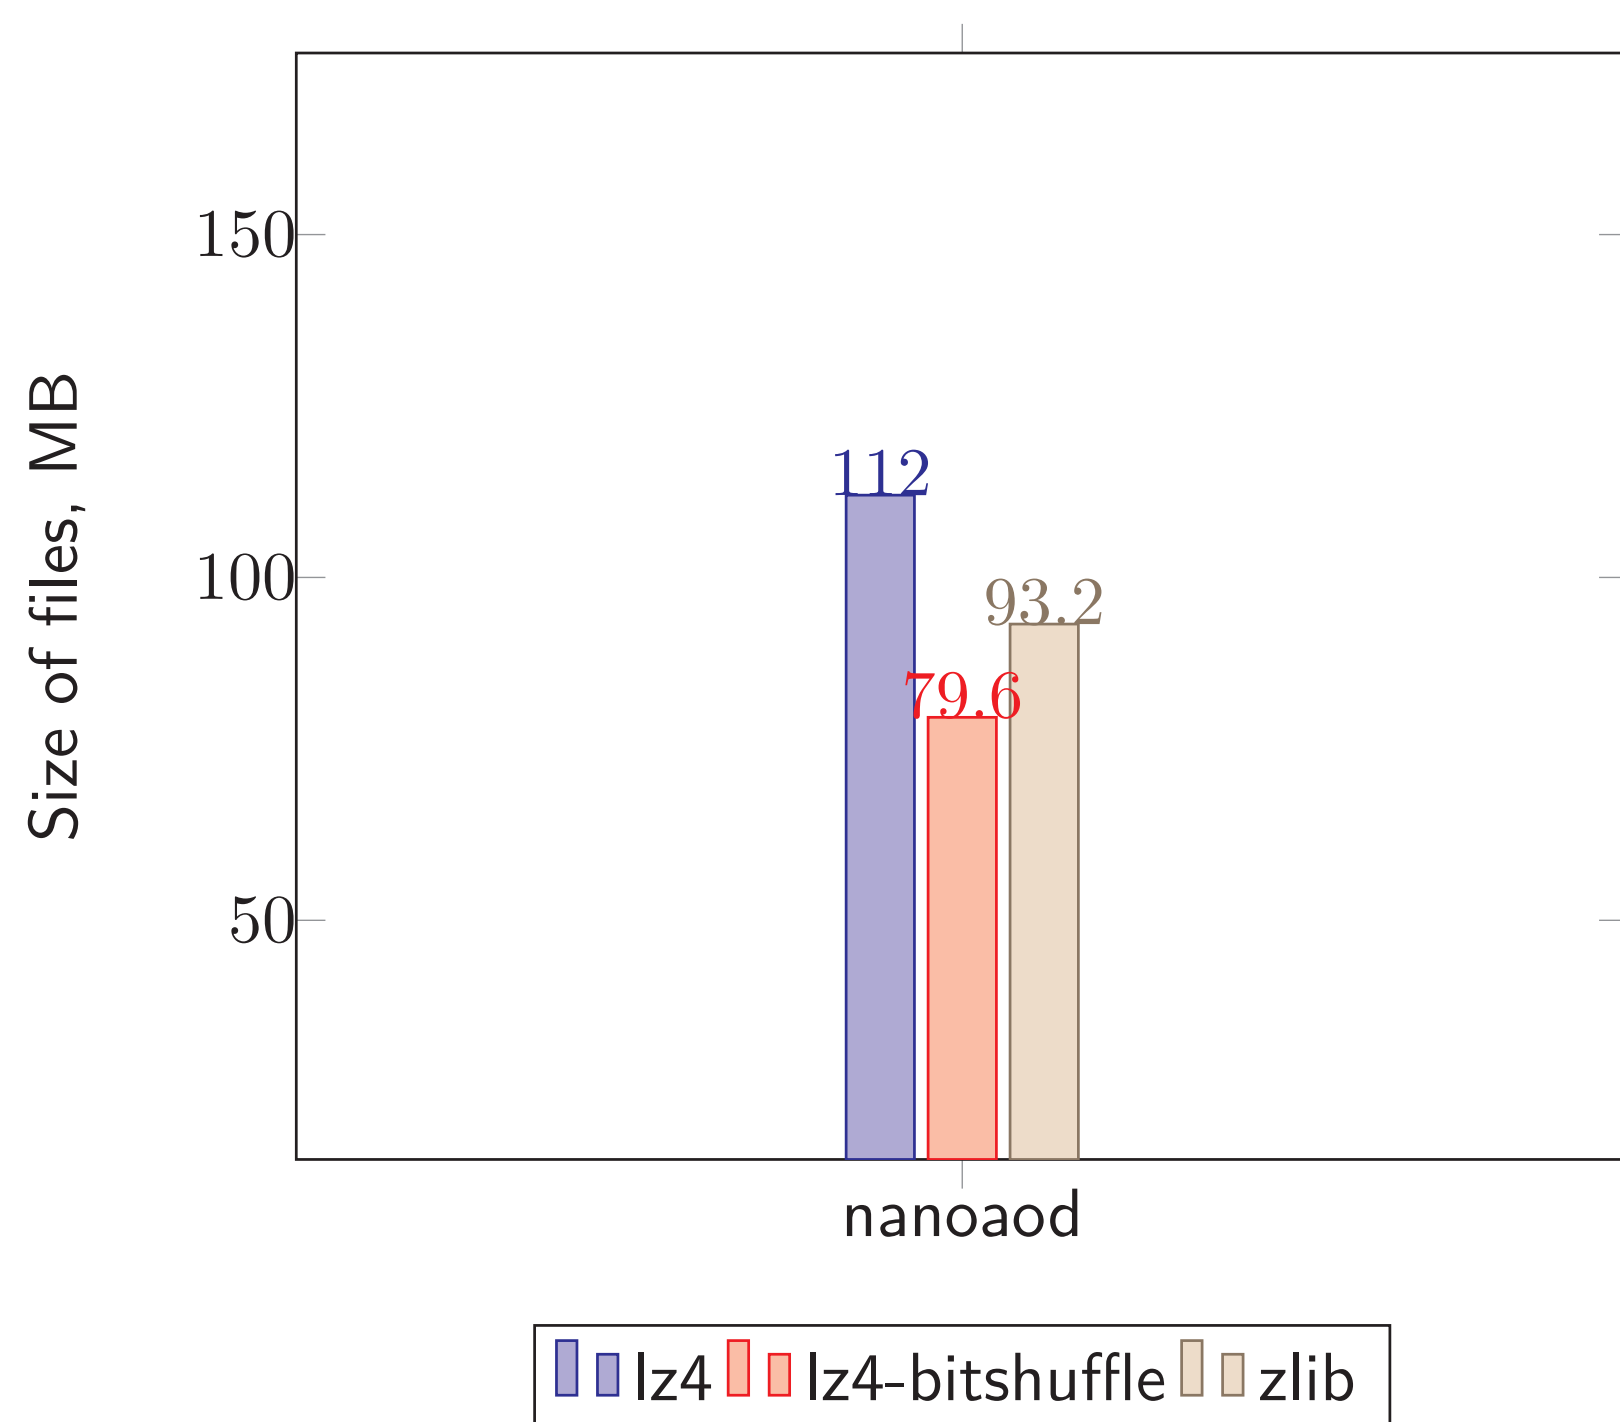


→ ZSTD is showing promising results, even though compression dictionary procedure is still not optimal.

## Recommendation to the users

- Ratio between compression ratio and compression/decompression speed:** LZ4
- Size of the file:** LZMA
- Recovering data from partial file (in case of crash):** tune AutoSave!
  - Default frequency is to save the meta-data every 10 clusters.
- Memory use or physical I/O performance:** tune AutoFlush!
  - Default is number of entries needed to reach 32 Mb of compressed data [number of entries or compressed data size]


## WIP: LZ4 + Bitshuffle



→ It is a promising result that shows that LZ4 can outperform ZLIB by better compression ratio, while compression time should be still optimized.

## Comparison of current ROOT compression algorithms - decompression

### Decompression speed, 2000 event TTree, MB/s



→ We know that decompression for LZ4 and ZSTD is faster than reading decompressed data (it depends on the medium type, e.g., SSD): significantly less data is coming from the I/O subsystem.

## Acknowledgements