Compression (in HDF5) in the projects ICOMEX and AIMES

Julian M. Kunkel Nathanel Hübbe

kunkel@dkrz.de

German Climate Computing Center (DKRZ), Research Group/Scientific Computing (WR)

18-01-2016

Outline

Overview

Lossless compression

- Evaluated existing compressors on climate data (GZ, BZ, LZMA, SZIP, ...)
- Developed pre-conditioner and adaptive compression scheme
  - Re-orders and/or re-formats (floating point) data
  - Improves compression (relative size -8 to -15%)
  - Compression ratio around 2.5:1

Lossy compression

- Evaluated APAX vs. GRIB2
- Working on user-defined compression qualities

Support analysis via reference data sets

- Program to create datasets with different properties
  Random, interpolated, procedural (landscape like), ...

Analysis of Compressability on Mistral

- Compression of LZMA, ZIP, SZIP, BZIP2 on production data
  - Selected a growing random sample of files (5 TiB, 300k files)
  - Analyzed compression ratio

Relative size of the compressed file vs. files scanned

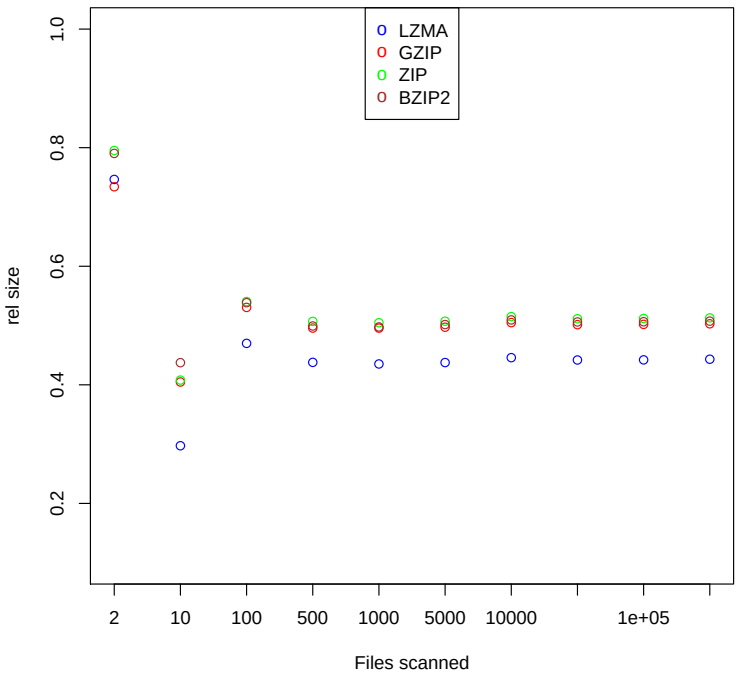

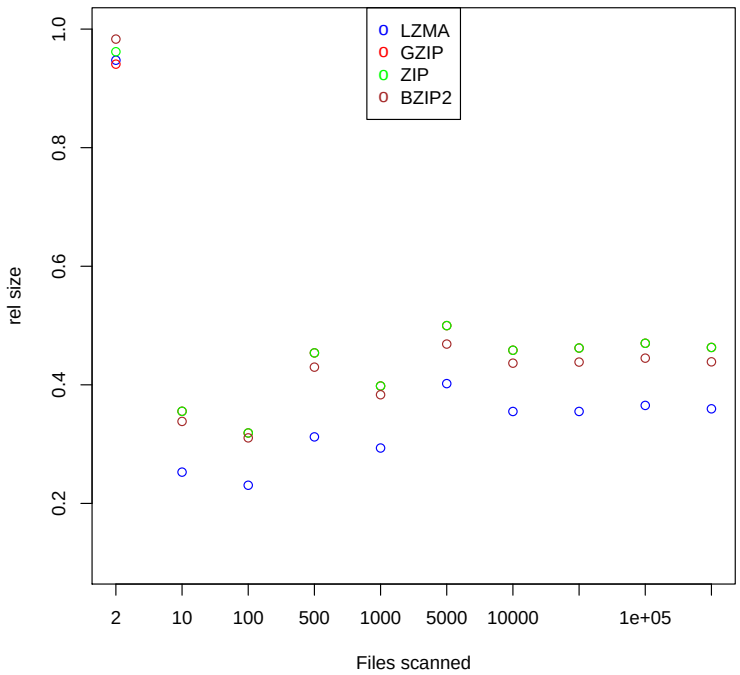

Left avg. across file number., right avg. over the data volume

Detailed Results

Summary

5 TiB, 335k files

	Avg. ratio x:1		Speed (MiB/s) per core	
Algoritm	files	size	Compression	Decompression
LZMA	2.25	2.78	3.0	28
BZIP2	1.97	2.28	4.8	20
GZIP	2.00	2.16	24.0	130
ZIP	1.95	2.16	24.9	135

Compression ratio for different projects (here LZMA)

- Rel. ratio: 3, 3.4, 1.4, 2.6, 1.3, 2.9
- Rel. avg. size ratio: 2.7, 2.8, 2.3, 3.2, 3.9

MAFISC: Lossless Compression for Multi-dim. Data

- How much lossless can be achieved on climate data?
- MAFISC: pre-conditioner (filter) and adaptive compressor
  - Example filter: compute binary deltas of FP values
    In several directions in a multi-dimensional dataset
  - Applies existing compression schemes (GZ, ...)
- One algorithm is not enough:
  MAFISC uses 180 different filter chains
- HDF5 filter has been created: `http://wr.informatik.uni-hamburg.de/research/projects/icomex/mafisc`
- 15% better compression of CMIP-5 dataset than LZMA
  - At the cost of CPU time
  - Costs per block can be bound by
    1. applying the currently best algorithm
    2. evaluating another algorithm (to see if this one is better)

Compressing a CMIP-5 Sub-set

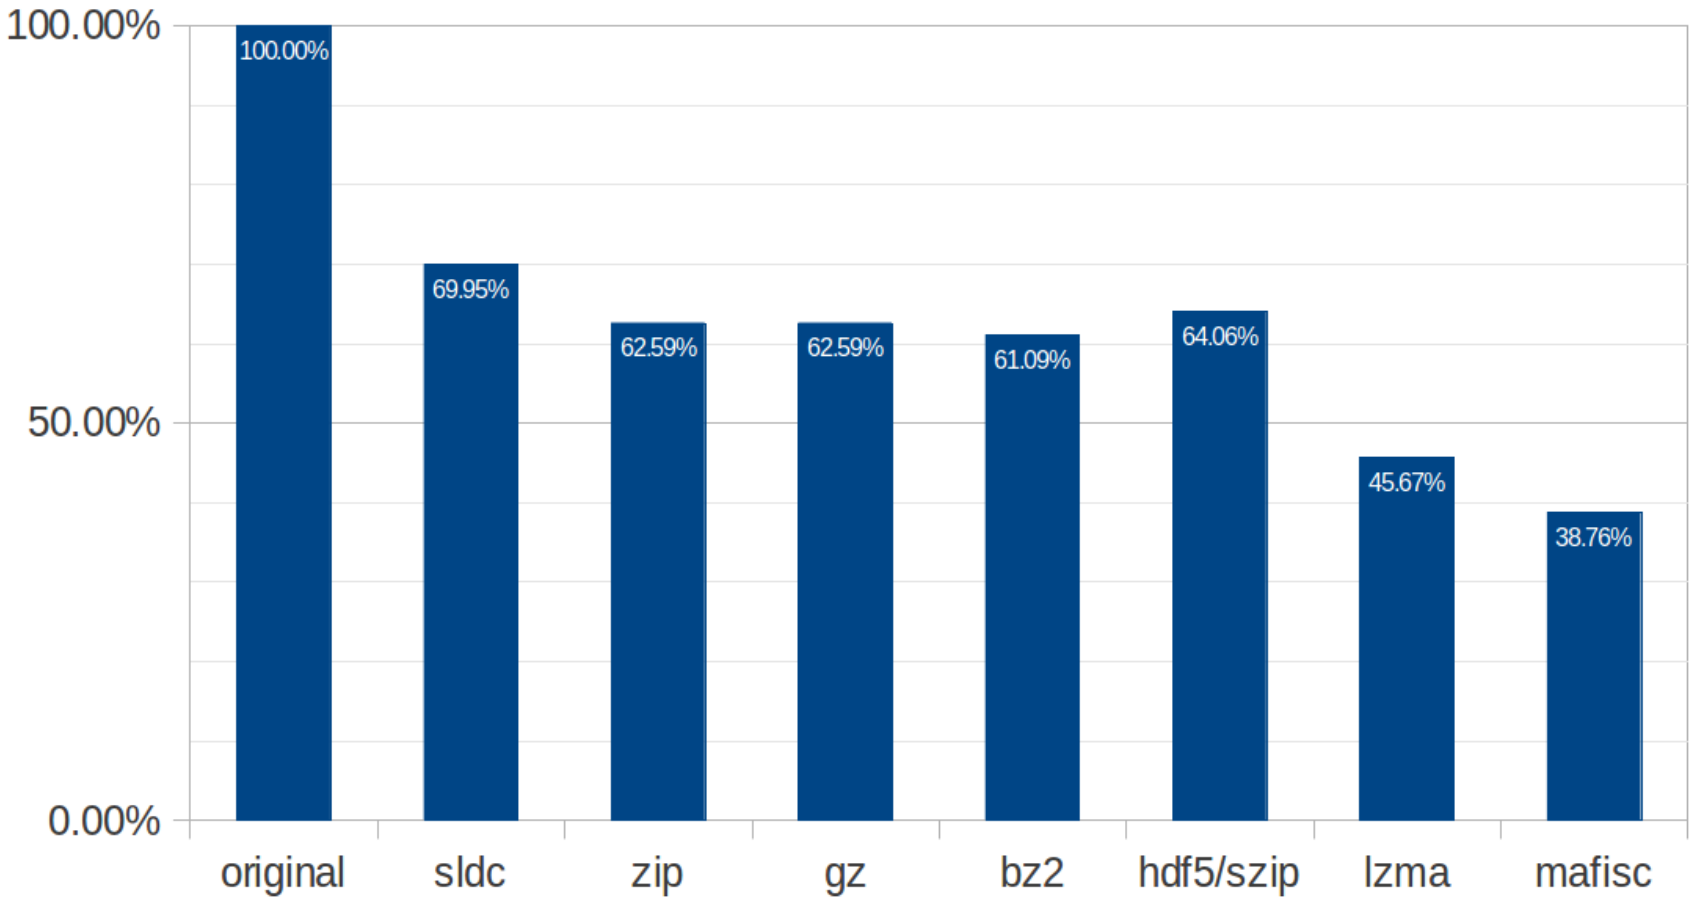

Compressed size relative to uncompressed size

Lossy Compression

Study of GRIB2, GRIB2 + JPEG2000, GRIB2 + LZMA, APAX

- APAX: Cooperation with SAMPLIFY (now out of business)

GRIB2 Compression Scheme

- Choose the number of bits (accuracy)
- Identify Min/Max of FP values
- Quantize values to $h = (max - min)/2^{bits}$
- Encode value v as integer i
  $\hat{v} = min + i \cdot h$

Analyzing 200+ Output Variables of ECHAM

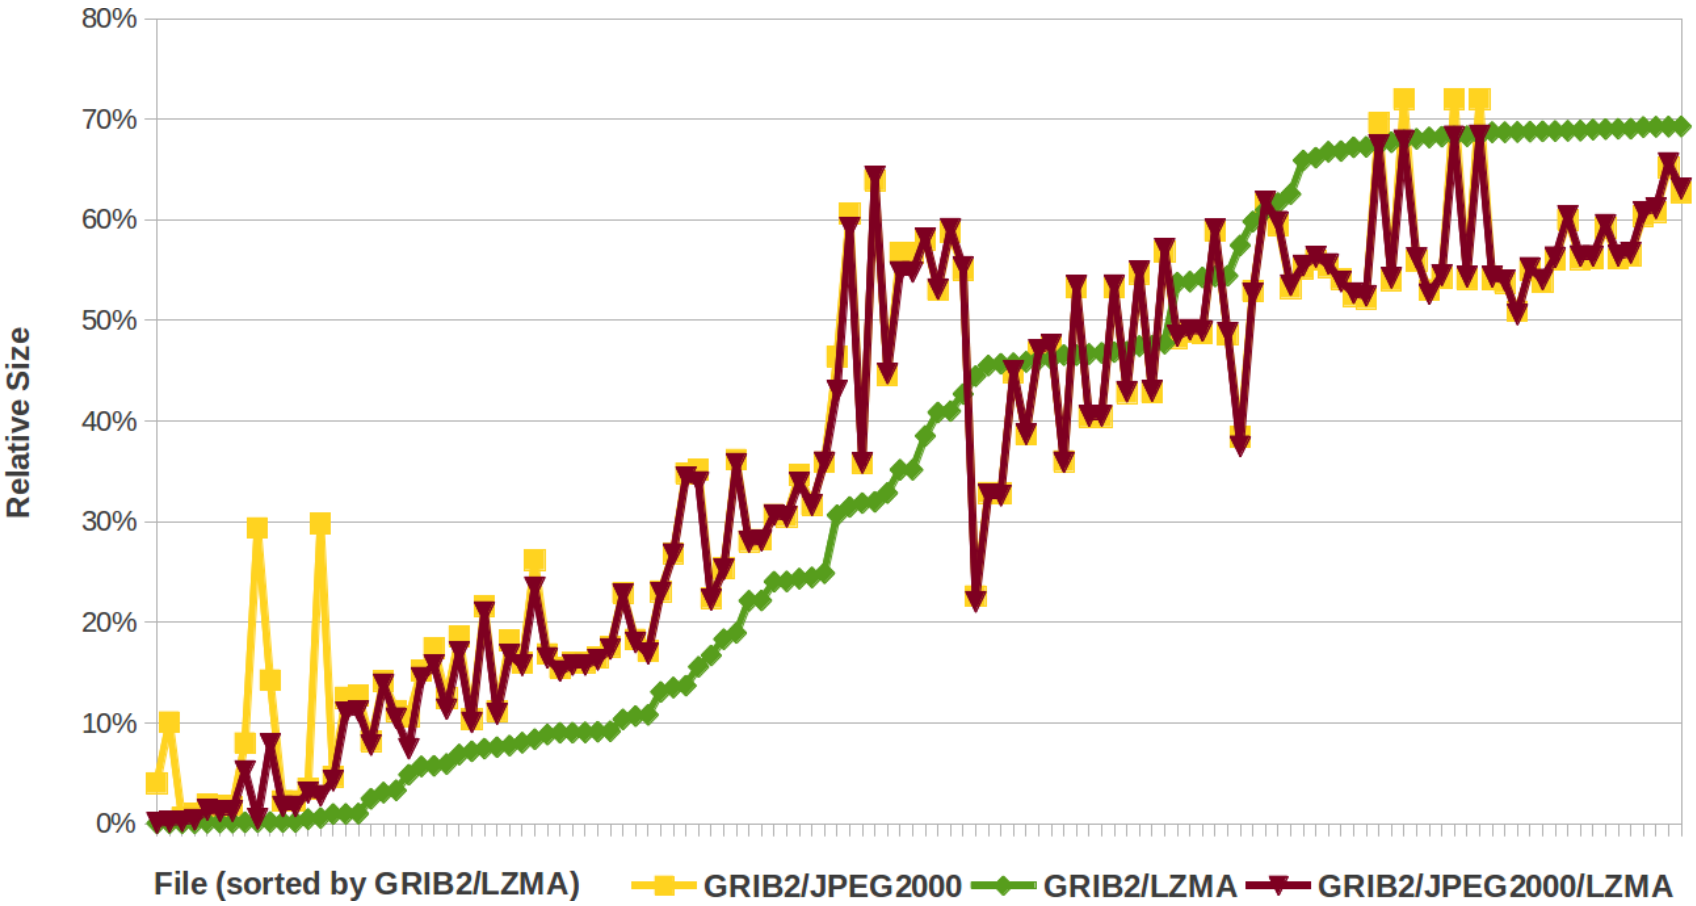

Relative compression size of variables sorted by GRIB2/LZMA

Consequence: One lossy compression scheme is not sufficient

Analyzing 200+ Output Variables of ECHAM

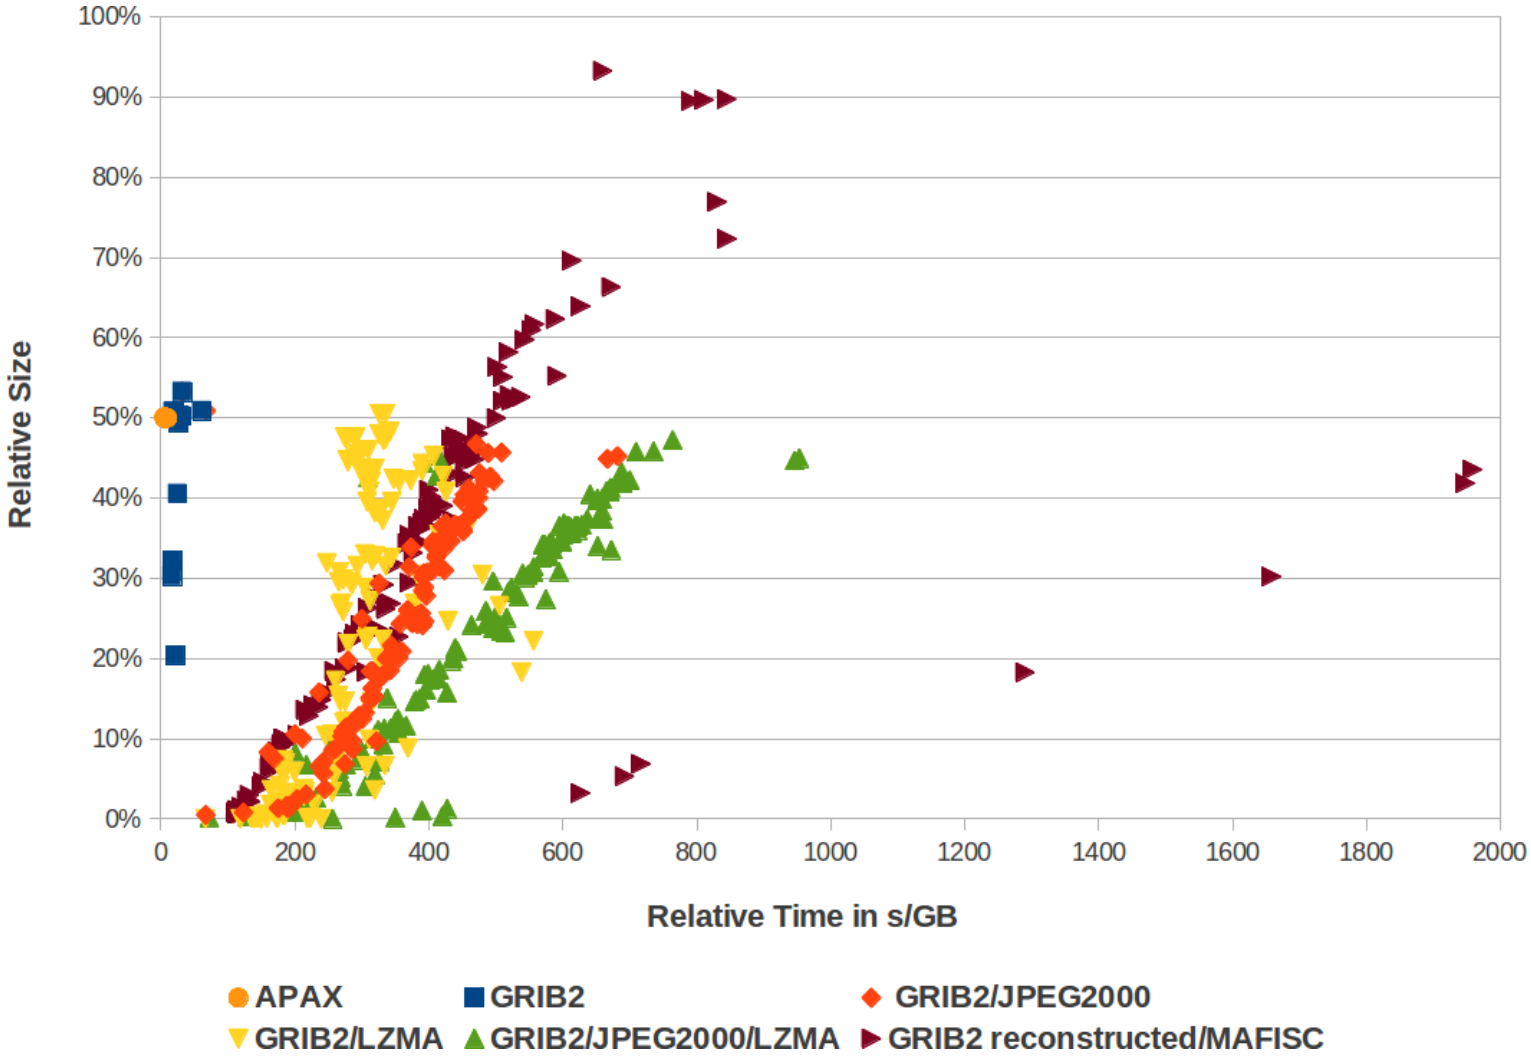

Performance vs. compression

Raw GRIB2 and APAX are very fast ($>$ 100 MiB/s)

More compressability $\rightarrow$ faster

Future Work in AIMES

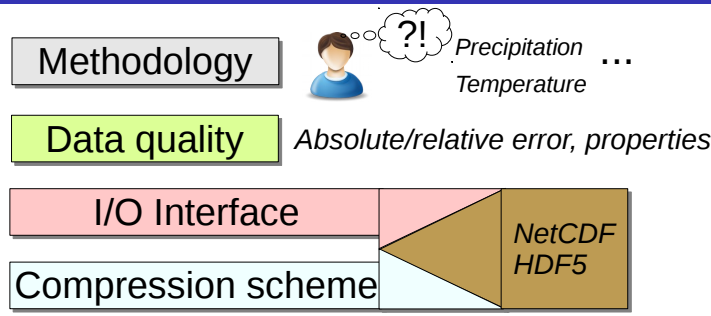

Goals

- Exploit lossy compression for scientists
- De-couple selection of compr. scheme from needed accuracy

Tasks

1. Develop user-defined metrics to tune accuracy
   - e.g. relative accuracy, significant digits, absolute accuracy
2. Develop an API to allow users to set the needed quality
3. Exploit the user-defined accuracy by extending compressors
4. Provide tools helping users to identify the appropriate accuracy

Current implementation strategy

- Develop an own API/library which performs compression [2],[3]
- Integrate it into HDF5 as filter and use HDF properties

Appendix

Scientific Computing

- Research Group of Prof. Ludwig at the University of Hamburg
- Embedded into DKRZ

Research

- Analysis of parallel I/O
- I/O & energy tracing tools
- Middleware optimization
- Alternative I/O interfaces
- Data reduction techniques
- Cost & energy efficiency