

FZ Hands-On Handout

Martin Burtscher, Franck Cappello, Sheng Di, Hanqi Guo, Samuel Li, Xin Liang,
Peter Lindstrom, Dingwen Tao, Kento Sato, Seung Woo Son, Robert Underwood, Kai Zhao *

February 14, 2024

*names listed alphabetically

Introduction

Welcome to the FZ Developer's meeting. During the hands-on portion, we hope that you will explore the compressors that are most interesting to you given the overview we presented. Each section of the tutorial is intended to be self-contained, so the order in which you visit the sections shouldn't matter. The developers of these compressors will be walking around to answer questions.

You have two options to work on the hands-on: via SSH (preferred) or on your own laptop via Docker

Connecting to the Machines

We have provided a remote system that you can log into via SSH to run the hands on. This is the recommended way to access the tutorial. The password is CUT LENT ROVE CAL LEFT CUBA GLUE OMEN EEL BEAR GAL

```
1 ssh -p 80 exouser@10.8.0.28
```

Using Your Laptop/Supercomputer

We have a pre-built container that runs on modern versions of Linux, MacOS, and Windows. This runs best on amd64/x86_64 processors (e.g. Intel/AMD processors), but it is possible to run on arm64/aarch64/m-series silicon as well. Please note that you may not be able to run all compressors in this way if you lack access to the appropriate hardware.

While the commands listed here are for podman, equivalent commands exist for Singularity, Apptainer, and Docker.

```
1 #Linux
2 sudo apt install podman #Debian/Ubuntu
3 sudo dnf install podman #RHEL/Fedora (probably installed)
4
5 #Windows
6 choco.exe install podman
7 podman machine init
8 podman machine start
9
10 #MacOS
11 brew install podman qemu
12 podman machine init
13 podman machine start
14
15 # all platforms
16 podman run -it --rm --platform linux/amd64 ghcr.io/robertu94/sz-zfp-zchecker:latest
```

Installation via Spack/Conda

We maintain spack and in many cases conda packages for this software as well. Please ask us about this if you need this for your environment.

Contents

1	SZ2/SZ3	4
2	cuSZ	6
3	ZFP	9
4	MGARD	12
5	SPERR	15
6	LC Framework	23
7	TEZip	30
8	DCTZ	32
9	LibPressio	33

1 SZ2/SZ3

contact: Kai Zhao, Sheng Di

1.1 CLI Interface (SZ2&SZ3)

```
1 # print help for the sz command
2 sz
3
4 # compress a 3D array in sz, with an absolute error bound of 1e-3
5 # -f : FP32
6 # -M ABS -A 1E-3: absolute error bound of 1e-3
7 # -i source_file : input
8 # -z compressed_file : output
9 # -3 500 500 100: array[100][500][500], reverse of C sequence
10 sz -f -M ABS -A 1E-3 -i /usecases/Hurricane/Uf48.dat -z \
11     /usecases/Hurricane/Uf48.dat.sz -3 500 500 100
12
13 # decompress the above data in sz, and print various statistics
14 # -i source_file : input, used for statistics
15 # -s compressed_file : input
16 # -x decompressed_file : output
17 # -3 500 500 100: array[100][500][500], reverse of C sequence
18 # -a : print statistics
19 sz -f -i /usecases/Hurricane/Uf48.dat -s /usecases/Hurricane/Uf48.dat.sz -x \
20     /usecases/Hurricane/Uf48.dat.sz.out -3 500 500 100 -a
21
22 # the command line of sz3 is backward compatible with sz2
23 # you can replace 'sz' with 'sz3' in the above lines to use sz3
```

1.2 C/C++ Interface (SZ3)

```
1 #include "SZ3/api/sz.hpp"
2 #compress
3 SZ3::Config conf(100, 500, 500); // array[100][500][500]
4 conf.cmprAlgo = SZ3::ALGO_INTERP_LORENZO;
5 conf.errorBoundMode = SZ3::EB_ABS; // refer to def.hpp for all supported error bound mode
6 conf.absErrorBound = 1E-3; // absolute error bound 1e-3
7 size_t cmpSize = 100 * 500 * 500; // cmpSize is overwritten by compressed data size
8 /**
9  * @tparam T source data type
10  * @param config compression configuration
11  * @param data source data
12  * @param cmpSize compressed data size in bytes
13  * @return compressed data, remember to 'delete []' when the data is no longer needed.
14  */
15 char *cmpData = SZ_compress<float>(conf, data, cmpSize);
16
17 #decompress
18 auto decompressedData = new float[100*500*500];
19 SZ3::Config conf1;
20 /**
21  * @param conf configuration placeholder. Overwritten by the compression configuration
22  * @param cmpData compressed data
```

```

23  * @param cmpSize compressed data size in bytes
24  * @param decData pre-allocated memory space for decompressed data
25  */
26  SZ_decompress<float>(conf1, cmpData, cmpSize, decData);
27

```

1.3 Python Interface (SZ3)

```

1  #python support is experimental
2  #you need to locate 1) sz3c lib file and 2) pysz.py
3  #sz3c lib can be found in your_cmake_folder/tools/sz3c/
4  #pysz.py can be found in sz3/tools/pysz/
5
6  import numpy as np
7  from pathlib import Path
8  from pysz import SZ
9  import sys
10
11  # prepare your data in numpy array format
12  data = np.fromfile('Uf48.bin.dat', dtype=np.float32)
13  data = np.reshape(data, (100, 500, 500))
14
15
16  lib_extention = {
17      "darwin": "libSZ3c.dylib",
18      "windows": "SZ3c.dll",
19  }.get(sys.platform, "libSZ3c.so")
20  # Please change the path to the SZ dynamic library file in your system
21  sz = SZ("/path/to/your/sz/lib/folder".format(lib_extention))
22
23  """
24  Compress data with SZ
25  :param data: original data, numpy array format, dtype is FP32 or FP64
26  :param eb_mode:# error bound mode, integer (0: ABS, 1:REL)
27  :param eb_abs: optional, abs error bound, double
28  :param eb_rel: optional, rel error bound, double
29  :param eb_pwr: optional, pwr error bound, double
30  :return: compressed data, numpy array format, dtype is np.uint8
31          compression ratio
32  """
33  data_cmpr, cmpr_ratio = sz.compress(data, 0, 1e-3, 0, 0)
34  print("compression ratio = {:.5G}".format(cmpr_ratio))
35
36  # decompress, both input and output data are numpy array
37  data_dec = sz.decompress(data_cmpr, data.shape, data.dtype)
38
39  # verify
40  sz.verify(data, data_dec)

```

2 cuSZ

For this section you will need an NVIDIA GPU.

contact: [Jiannan Tian](#)

cuSZ is a GPU adoption of SZ. cuSZ incorporates several practical components, including fully parallelized prediction & error quantization stage and coarsely parallelized Huffman encoding (the released version).

2.1 CLI Interface

1. The interpretation of data shape is important to utilize the dimensional information. cuSZ puts two ways to specify data shape,
 - fastest to slowest (inner to outer loop if sequential, also CUDA dim3)
`--dim3 [X]x[Y]x[Z]` or `--len [X]x[Y]x[Z]`
 - slowest to fastest (outer to inner loop if sequential)
`--size [Z]x[Y]x[X]`
2. Two predictors are supported:
 - Lorenzo predictor of higher throughput
`--predictor lorenzo`
 - cubic spline interpolation of higher quality (3D only)
`--predictor spline`
3. Two modes are supported: `abs` and `r2r` (relative to data range).

```
1 # print help for the cusz command
2 cusz
3
4 # print even longer help for the cusz command
5 cusz -h
6
7 # Currently, only 32-bit fp-type is supported.
8 # To compress, -z is needed.
9 # By default, Lorenzo predictor is used.
10 # To use 3D cubic spline interpolation,
11 # --predictor spline should be appended in the command line.
12 cuszi -t f32 -m rel -e [ErrorBound] -i [/PATH/TO/DATA] --dim3 [X]x[Y]x[Z] -z --report time
13
14 # To decompress, -x is needed
15 # The optional --compare [/PATH/TO/DATA] can be specified to
16 # compared the original and reconstructed data.
17 cuszi -i [/PATH/TO/DATA].cusza -x --report time --compare [/PATH/TO/DATA]
```

2.2 C/C++ Interface

cuSZ uses namespace `psz`, which stands for GPU-parallelized **SZ**.

```

1 // header zone
2 #include "cusz.h"
3 #include "context.h"
4
5 pszheader header; // declare metadata
6
7 uint8_t* ptr_compressed; // to access internal compressor buffer
8 size_t compressed_len; // to be changed by compressor
9
10 T *d_uncomp, *h_uncomp;
11 T *d_decomp, *h_decomp;
12
13 auto oribytes = sizeof(T) * len;
14 cudaMalloc(&d_uncomp, oribytes), cudaMallocHost(&h_uncomp, oribytes);
15 cudaMalloc(&d_decomp, oribytes), cudaMallocHost(&h_decomp, oribytes);
16
17 // User handles loading from filesystem & transferring to device.
18 // io::read_binary_to_array(fname, h_uncomp, len);
19 // cudaMemcpy(d_uncomp, h_uncomp, oribytes, cudaMemcpyHostToDevice);
20
21 // initialize CUDA stream
22 cudaStream_t stream; cudaStreamCreate(&stream);
23
24 // use default setup for cuSZ framework
25 pszframe* work = pszdefault_framework();
26
27 // alternative finer-grained setup
28 // pszframe* work = new pszframe{
29 //     .predictor = pszpredictor{.type = Lorenzo},
30 //     .quantizer = pszquantizer{.radius = 512},
31 //     .hfcoder = pszhfrc{.style = Coarse},
32 //     .max_outlier_percent = 20};
33
34 pszcompressor* comp = psz_create(work, F4);
35
36 // specify compression mode: Rel or Abs
37 // specify error bound: any fractional number
38 pszctx* ctx = new pszctx{.mode = Rel, .eb = 2.4e-4};
39 // x, y, z, w (w is a placeholder for future use)
40 pszlen uncomp_len = pszlen{3600, 1800, 1, 1};
41 pszlen decomp_len = uncomp_len;
42
43 // C++ time-record type
44 psz::TimeRecord compress_timerecord;
45 psz::TimeRecord decompress_timerecord;
46
47 { // compression scope
48 psz_compress_init(comp, uncomp_len, ctx);
49 psz_compress(comp, d_uncomp, uncomp_len, &ptr_compressed, &compressed_len, &header,
50     (void*)&compress_timerecord, stream);
51
52 // User can interpret the collected time information in other ways.
53 psz::TimeRecordViewer::view_compression(
54     &compress_timerecord, oribytes, compressed_len);

```

```

55
56 // verify header
57 printf("header.%-*s : %p\n", 12, "(addr)", (void*)&header);
58 printf("header.%-*s : %u, %u, %u\n", 12, "{x,y,z}", header.x, header.y,
59     header.z);
60 printf("header.%-*s : %lu\n", 12, "filesize", psz_utils::filesize(&header));
61 } // end of compression scope
62
63 // If needed, User should perform a memcpy to transfer `ptr_compressed`
64 // before `compressor` is destroyed.
65 uint8_t* compressed_buf;
66 cudaMalloc(&compressed_buf, compressed_len);
67 cudaMemcpy(compressed_buf, ptr_compressed, compressed_len, cudaMemcpyDeviceToDevice);
68
69 { // decompression scope
70 psz_decompress_init(comp, &header);
71 psz_decompress(
72     comp, ptr_compressed, compressed_len, d_decomp, decomp_len,
73     (void*)&decompress_timerecord, stream);
74 // print kernel time records
75 psz::TimeRecordViewer::view_decompression(
76     &decompress_timerecord, oribytes);
77 } // end of decompression scope
78
79 // demo: offline checking (de)compression quality.
80 psz::eval_dataquality_gpu(d_decomp, d_uncomp, len, compressed_len);
81
82 psz_release(comp);
83
84 cudaFree(compressed_buf);
85 cudaFree(d_uncomp), cudaFreeHost(h_uncomp);
86 cudaFree(d_decomp), cudaFreeHost(h_decomp);
87
88 cudaStreamDestroy(stream);

```


3 ZFP

contact: [Peter Lindstrom](#)

ZFP is a compressed format for representing multidimensional floating-point and integer arrays. ZFP provides compressed-array classes that support high throughput read and write random access to individual array elements. ZFP also supports serial and parallel compression of whole arrays, e.g., for applications that read and write large data sets to and from disk. For documentation, source code, and other resources, see <https://zfp.llnl.gov>.

3.1 CLI Interface

```
1 # print help for the zfp command
2 zfp
3
4 # load a 3d array of floats, use zfp with an absolute error bound of 1e-4,
5 # and collect and print various statistics
6 zfp -s -i /usecases/Hurricane/Uf48.dat -f -3 500 500 100 -a 1e-4
7
8 # use zfp's fixed-rate mode to compress to 13 bits/value
9 zfp -s -i /usecases/Hurricane/Uf48.dat -f -3 500 500 100 -r 13
10
11 # use zfp's fixed-precision mode with 10 bits of precision
12 zfp -s -i /usecases/Hurricane/Uf48.dat -f -3 500 500 100 -p 10
13
14 # use zfp's reversible (lossless) mode to compress to stdout, decompress from stdin,
15 # and encode the array size and compression mode in a header
16 zfp -h -i /usecases/Hurricane/Uf48.dat -f -3 500 500 100 -R -z - | \
17 zfp -h -z - -o /tmp/Uf48.out
```

3.2 C Interface

```
1 #include "zfp.h"
2
3 // input: (double* array, size_t nx, size_t ny, size_t nz, double tolerance)
4
5 // initialize metadata for the 3D array a[nz][ny][nx]
6 zfp_type type = zfp_type_double; // array scalar type
7 zfp_field* field = zfp_field_3d(array, type, nx, ny, nz); // array metadata
8
9 // initialize metadata for a compressed stream
10 zfp_stream* zfp = zfp_stream_open(NULL); // compressed stream and parameters
11 zfp_stream_set_accuracy(zfp, tolerance); // set tolerance for fixed-accuracy mode
12 // zfp_stream_set_precision(zfp, precision); // alternative: fixed-precision mode
13 // zfp_stream_set_rate(zfp, rate, type, 3, zfp_false); // alternative: fixed-rate mode
14 // zfp_stream_set_reversible(zfp); // alternative: lossless mode
15
16 // allocate buffer for compressed data
17 size_t bufsize = zfp_stream_maximum_size(zfp, field); // capacity of compressed buffer
18 void* buffer = malloc(bufsize); // storage for compressed stream
19
20 // associate bit stream with allocated buffer
21 bitstream* stream = stream_open(buffer, bufsize); // bit stream to compress to
22 zfp_stream_set_bit_stream(zfp, stream); // associate with compressed stream
23
```

```

24 // compress array
25 zfp_stream_rewind(zfp); // rewind stream to beginning
26 size_t zfp_size = zfp_compress(zfp, field); // return value is compressed byte size
27
28 // decompress array
29 zfp_stream_rewind(zfp); // rewind stream for decompression
30 assert(zfp_decompress(zfp, field) == zfp_size); // decompress stream

```

3.3 Python Interface

Install ZFP's Python bindings using `pip3 install --user zfp`.

```

1 import zfp
2 import numpy as np
3
4 # initialize 100 * 80 NumPy array to Hilbert matrix
5 my_array = np.array([1/(1+x+y) for y in range(100) for x in range(80)]).reshape((100, 80))
6
7 # compress array to within an absolute error tolerance of 1e-6
8 compressed_data = zfp.compress_numpy(my_array, tolerance=1e-6)
9
10 # decompress to another NumPy array
11 decompressed_array = zfp.decompress_numpy(compressed_data)
12
13 # verify that error tolerance is satisfied
14 np.testing.assert_allclose(my_array, decompressed_array, atol=1e-6)
15
16 # compression ratio
17 decompressed_array.nbytes / len(compressed_data)

```

3.4 C++ Interface

```

1 #include <climits>
2 #include <cmath>
3 #include <iostream>
4 #include "zfp/array2.hpp"
5 #include "zfp/codec/gencodec.hpp"
6
7 double f(size_t x, size_t y) { return 1. / (1 + x + y); }
8
9 template <class Array>
10 void evaluate(Array& a, std::string name)
11 {
12 // initialize array to Hilbert matrix
13 for (size_t y = 0; y < a.size_y(); y++)
14     for (size_t x = 0; x < a.size_x(); x++)
15         a(x, y) = f(x, y);
16
17 // optional: compress any cached data
18 a.flush_cache();
19
20 // compute error
21 double abserr = 0;
22 double relerr = 0;
23 for (size_t y = 0; y < a.size_y(); y++)

```

```

24     for (size_t x = 0; x < a.size_x(); x++) {
25         double d = std::abs(a(x, y) - f(x, y));
26         abserr += d;
27         relerr += d / std::abs(f(x, y));
28     }
29     abserr /= a.size();
30     relerr /= a.size();
31
32     // output error
33     std::cout << std::defaultfloat << a.rate() << "-bit " << name << " array" << std::endl;
34     std::cout << " mean absolute error = " << std::scientific << abserr << std::endl;
35     std::cout << " mean relative error = " << std::scientific << relerr << std::endl;
36 }
37
38 int main()
39 {
40     const size_t nx = 60; // array width
41     const size_t ny = 100; // array height
42
43     // scalar type for uncompressed array
44     typedef float real; // single precision
45     //typedef __fp16 real; // alternative: half precision
46     const double rate = sizeof(real) * CHAR_BIT; // storage size in bits/value
47
48     // uncompressed 2D array storing scalars of type real
49     zfp::array2< double, zfp::codec::generic2<double, real> > a(nx, ny, rate);
50     evaluate(a, "uncompressed");
51
52     // zfp compressed 2D array with equal storage size
53     zfp::array2<double> b(nx, ny, rate);
54     evaluate(b, "zfp");
55
56     return 0;
57 }

```

4 MGARD

Contact: [Xin Liang](#), [Jieyang Chen](#)

MGARD is a highly functional, performant, portable, and extendable compressor for scientific data. Built upon wavelet and finite element theories, MGARD features rigorous error controls on both raw data and derived quantities, diverse functionalities including both compression and refactoring, utility support for both uniform/non-uniform structured data and unstructured data, and unified interfaces on CPU and GPUs. For detailed documentation, publication, and source code, please refer to <https://github.com/CODARcode/MGARD>.

4.1 CLI Interface

```
1 # print help for the mgard-x command
2 mgard-x
3
4 Options
5     -z: compress data
6         -i <path to data file to be compressed>
7         -c <path to compressed file>
8         -t <s|d>: data type (s: single; d:double)
9         -n <ndim>: total number of dimensions
10            [dim1]: slowest dimension
11            [dim2]: 2nd slowest dimension
12            ...
13            [dimN]: fastest dimension
14     -u <path to coordinate file>
15     -m <abs|rel>: error bound mode (abs: absolute; rel: relative)
16     -e <error>: error bound
17     -s <smoothness>: smoothness parameter
18     -l choose lossless compressor (0:Huffman 1:Huffman+LZ4 2:Huffman+Zstd)
19     -d <auto|serial|openmp|cuda|hip|sycl>: device type
20     -v enable verbose (show timing and statistics)
21
22     -x: decompress data
23         -c <path to compressed file>
24         -o <path to decompressed file>
25         -d <auto|serial|cuda|hip>: device type
26         -v enable verbose (show timing and statistics)
27
28 # Example: compress a 3D array in MGARD with an absolute error bound of 1e-3
29 mgard-x -z -i /usecases/Hurricane/Uf48.dat -c /usecases/Hurricane/Uf48.dat.mgard -t s -n 3 100 500 500
30     -s 0 -l 2 -d auto -m abs -e 1e-3
31
32 # Example: decompress the 3D array
33 mgard-x -x -c /usecases/Hurricane/Uf48.dat.mgard -o /usecases/Hurricane/Uf48.dat.mgard.out -d auto
34
35 # Example: refactor a 3D array
36 mdr-x -z -i /usecases/Hurricane/Uf48.dat -c refactored.mgard -t s -n 3 100 500 500 -l 2 -h 1 -d auto
37
38 # Retrieve and reconstruct the data under absolute error bound 1e-3
39 mdr-x -x -c refactored.mgard -o /usecases/Hurricane/Uf48.dat.mgard.out -a /usecases/Hurricane/Uf48.dat
```

4.2 C/C++ Interface

MGARD uses namespace `mgard_x` with a sub-namespace `MDR` for refactoring APIs.

High-level APIs for error-controlled compression and decompression:

```
1  #include "mgard/compress_x.hpp"
2
3  // prepare data buffers
4  mgard_x::DIM num_dims = 3;
5  mgard_x::SIZE n1, n2, n3;
6  std::vector<mgard_x::SIZE> shape{n1, n2, n3};
7  mgard_x::SIZE in_byte = n1 * n2 * n3 * sizeof(double);
8  mgard_x::SIZE out_byte;
9  //... load data into in_array
10 double *in_array = ...;
11 void *compressed_array = NULL;
12 void *decompressed_array = NULL;
13 // tol: error tolerance
14 // s: smoothness parameter
15 double tol = 0.01, s = 0;
16
17 // MGARD config parameters
18 mgard_x::Config config;
19
20 // Compressing with high-level API
21 mgard_x::compress(num_dims, mgard_x::data_type::Double, shape, tol, s, mgard_x::error_bound_type::REL,
22
23 // Decompressing with high level API
24 mgard_x::decompress(compressed_array, out_byte, decompressed_array, config, false);
```

High-level APIs for refactoring and error-controlled progressive retrieval:

```
1  #include "mgard/mdr_x.hpp"
2  ...
3
4  // prepare data buffers
5  mgard_x::DIM num_dims = 3;
6  mgard_x::SIZE n1, n2, n3;
7  std::vector<mgard_x::SIZE> shape{n1, n2, n3};
8  mgard_x::SIZE in_byte = n1 * n2 * n3 * sizeof(double);
9  mgard_x::SIZE out_byte;
10 //... load data into in_array
11 double *in_array = ...;
12
13 mgard_x::Config config;
14 mgard_x::MDR::RefactoredMetadata refactored_metadata;
15 mgard_x::MDR::RefactoredData refactored_data;
16
17 // Refactor with high-level API
18 mgard_x::MDR::MDRefactor(D, mgard_x::data_type::Double, shape, in_array, refactored_metadata,
19     refactored_data, config, false);
20
21 // Save refactored_metadata and refactored_data to files
```

```
22 ...
23
24 mgard_x::MDR::ReconstructedData reconstructed_data;
25 // Read in refactored_metadata from file
26 ...
27 // Progressively reconstruct for each error bound
28 for (double tol : tolerances) {
29     // Specify error bound and smoothness parameter for each subdomain
30     for (auto &metadata : refactored_metadata.metadata) {
31         metadata.requested_tol = tol;
32         metadata.requested_s = s;
33     }
34     // Determine required data components for reconstruction
35     mgard_x::MDR::MDRequest(refactored_metadata, config);
36     // Read in required data components from files
37     ...
38     // Reconstruct with high-level API
39     mgard_x::MDR::MDReconstruct(refactored_metadata, refactored_data, reconstructed_data,
40         config, false, original_data);
41
42     // reconstructed_data now contains progressively reconstructed data
43     double out_data = reconstructed_data.data;
44 }
```

5 SPERR

Contact: [Samuel Li \(shaomeng@ucar.edu\)](mailto:shaomeng@ucar.edu)

Repo: github.com/NCAR/SPERR/

Wiki: github.com/NCAR/SPERR/wiki

SPERR uses *wavelet transforms* to decorrelate the data, encodes the quantized coefficients, and explicitly corrects any data point exceeding a prescribed point-wise error (PWE) tolerance. Most often, SPERR produces the smallest compressed bitstream honoring a PWE tolerance.

A SPERR bitstream can be used to reconstruct the data fields in two additional fashions: *flexible-rate* decoding and *multi-resolution* decoding.

- *Flexible-rate* decoding: any prefix of a SPERR bitstream (i.e., a sub-bitstream that starts from the very beginning) produced by a simple truncation is still valid to reconstruct the data field, though at a lower quality. This ability is useful for applications such as tiered storage and data sharing over slow connections, to name a few.
- *Multi-resolution* decoding: a hierarchy of the data field with coarsened resolutions can be obtained together with the native resolution. This ability is useful for quick analyses with limited resources.

On a Unix-like system with a working C++ compiler and CMake, SPERR can be built from source and made available to users in just a few commands. See this [README](#) for detail.

5.1 CLI Interface

Upon a successful build, four CLI utility programs are placed in the `./bin/` directory; three of them are most relevant here: `sperr2d`, `sperr3d`, and `sperr3d_trunc`. Each of them can be invoked with the `-h` option to display a help message.

5.1.1 `sperr2d`

`sperr2d` is responsible for compressing and decompressing a 2D data slice. Its help message contains all the options `sperr2d` takes:

```
1 $ ./bin/sperr2d -h
2
3 Usage: ./bin/sperr2d [OPTIONS] [filename]
4
5 Positionals:
6   filename TEXT:FILE           A data slice to be compressed, or
7                                 a bitstream to be decompressed.
8
9 Options:
10  -h,--help                     Print this help message and exit
11
12 Execution settings:
13  -c Excludes: -d               Perform a compression task.
14  -d Excludes: -c               Perform a decompression task.
15
16 Input properties:
17  --ftype UINT                  Specify the input float type in bits. Must be 32 or 64.
18  --dims [UINT,UINT]           Dimensions of the input slice. E.g., `--dims 128 128`
19                                 (The fastest-varying dimension appears first.)
20
21 Output settings:
22  --bitstream TEXT             Output compressed bitstream.
```

```

23  --decomp_f TEXT           Output decompressed slice in f32 precision.
24  --decomp_d TEXT           Output decompressed slice in f64 precision.
25  --decomp_lowres_f TEXT    Output lower resolutions of the decompressed slice in f32 precision.
26  --decomp_lowres_d TEXT    Output lower resolutions of the decompressed slice in f64 precision.
27  --print_stats Needs: -c   Show statistics measuring the compression quality.
28
29  Compression settings:
30  --pwe FLOAT Excludes: --psnr --bpp
31                          Maximum point-wise error (PWE) tolerance.
32  --psnr FLOAT Excludes: --pwe --bpp
33                          Target PSNR to achieve.
34  --bpp FLOAT:FLOAT in [0 - 64] Excludes: --pwe --psnr
35                          Target bit-per-pixel (bpp) to achieve.

```

Examples:

1. Compress a 2D slice in 512×512 dimension, single-precision floats with a PWE tolerance of 10^{-2} :
`./bin/sperr2d -c --ftype 32 --dims 512 512 --pwe 1e-2 \`
`--bitstream ./out.stream ./in.f32`
2. Perform the compression task described above, plus also write out the compress-decompressed slice, and finally print statistics measuring the compression quality:
`./bin/sperr2d -c --ftype 32 --dims 512 512 --pwe 1e-2 \`
`--decomp_f ./out.decomp --print_stats --bitstream ./out.stream ./in.f32`
3. Decompress from a SPERR bitstream, and write out the slice in native and coarsened resolutions:
`./bin/sperr2d -d --decomp_f ./out.decomp --decomp_lowres_f ./out.lowres ./sperr.stream`
 In this example, the output file `out.decomp` will be of the native resolution, and six other files (`out.lowres.256x256`, `out.lowres.128x128`, etc.) will also be produced with their filenames indicating the coarsened resolution.

5.1.2 sperr3d

`sperr3d` is responsible for compressing and decompressing a 2D data volume. Compared to `sperr2d` which compresses the input 2D slice as a whole, `sperr3d` divides an input 3D volume into smaller chunks, and then compresses each chunk individually. This treatment allows for compressing and decompressing all the small chunks in parallel. `sperr3d` uses 256^3 as the default chunk dimension, but any number in the order of low hundreds and preferably divides the full volume is a good number. Command line options `--chunks` and `--omp` control the chunking and parallel execution behavior respectively.

The help message of `sperr3d` details all the options it takes:

```

1  $ ./bin/sperr3d -h
2
3  Usage: ./bin/sperr3d [OPTIONS] [filename]
4
5  Positionals:
6  filename TEXT:FILE           A data volume to be compressed, or
7                                a bitstream to be decompressed.
8
9  Options:
10 -h,--help                    Print this help message and exit
11
12 Execution settings:
13 -c Excludes: -d              Perform a compression task.
14 -d Excludes: -c              Perform a decompression task.
15 --omp UINT                    Number of OpenMP threads to use. Default (or 0) to use all.

```



```

16
17 Input properties (for compression):
18   --ftype UINT          Specify the input float type in bits. Must be 32 or 64.
19   --dims [UINT,UINT,UINT] Dimensions of the input volume. E.g., `--dims 128 128 128`
20                        (The fastest-varying dimension appears first.)
21
22 Output settings:
23   --bitstream TEXT      Output compressed bitstream.
24   --decomp_f TEXT       Output decompressed volume in f32 precision.
25   --decomp_d TEXT       Output decompressed volume in f64 precision.
26   --decomp_lowres_f TEXT Output lower resolutions of the decompressed volume in f32 precision.
27   --decomp_lowres_d TEXT Output lower resolutions of the decompressed volume in f64 precision.
28   --print_stats Needs: -c Print statistics measuring the compression quality.
29
30 Compression settings:
31   --chunks [UINT,UINT,UINT] Dimensions of the preferred chunk size. Default: 256 256 256
32                        (Volume dims don't need to be divisible by these chunk dims.)
33   --pwe FLOAT Excludes: --psnr --bpp
34                        Maximum point-wise error (PWE) tolerance.
35   --psnr FLOAT Excludes: --pwe --bpp
36                        Target PSNR to achieve.
37   --bpp FLOAT:FLOAT in [0 - 64] Excludes: --pwe --psnr
38                        Target bit-per-pixel (bpp) to achieve.

```

Examples:

1. Compress a 3D volume in $384 \times 384 \times 256$ dimension, double-precision floats, using a PWE tolerance of 10^{-9} and chunks of $192 \times 192 \times 256$:

```
./bin/sperr3d -c --omp 4 --ftype 64 --dims 384 384 256 --chunks 192 192 256 \
--pwe 1e-9 --bitstream ./out.stream ./in.f64
```
2. Perform the compression task described above, plus also write out the compress-decompressed volume, and finally print statistics measuring the compression quality:

```
./bin/sperr3d -c --omp 4 --ftype 64 --dims 384 384 256 --chunks 192 192 256 \
--pwe 1e-9 --decomp_d ./out.decomp --print_stats --bitstream ./out.stream ./in.f64
```
3. Decompress from a SPERR bitstream, and write out the volume in native and coarsened resolutions:

```
./bin/sperr3d -d --decomp_d ./out.decomp --decomp_lowres_d ./out.lowres ./sperr.stream
```

In this example, the output file `out.decomp` will be of the native resolution, and five other files (`out.lowres.192x192x128`, `out.lowres.96x96x64`, etc.) will also be produced with their file-names indicating the coarsened resolution.

To support multi-resolution decoding in 3D cases, the individual chunks (`--chunks`) need to 1) approximate a cube, so that there are the same number of wavelet transforms performed on each dimension, and 2) divide the full volume in each dimension.

5.1.3 sperr3d_trunc

`sperr3d_trunc` is responsible for properly truncating a multi-chunked 3D bitstream, i.e., locating the offset of each chunk in a bitstream, and truncating each chunk individually.

The help message of `sperr3d_trunc` details its options:

```

1 $ ./bin/sperr3d_trunc -h
2
3 Usage: ./bin/sperr3d_trunc [OPTIONS] [filename]

```

```

4
5 Positionals:
6   filename TEXT:FILE           The original SPERR3D bitstream to be truncated.
7
8 Options:
9   -h,--help                   Print this help message and exit
10
11 Truncation settings:
12   --pct UINT REQUIRED          Percentage (1--100) of the original bitstream to truncate.
13   --omp UINT                  Number of OpenMP threads to use. Default (or 0) to use all.
14
15 Output settings:
16   -o TEXT                     Write out the truncated bitstream.
17
18 Input settings:
19   --orig32 TEXT              Original raw data in 32-bit precision to calculate compression
20                             quality using the truncated bitstream.
21   --orig64 TEXT              Original raw data in 64-bit precision to calculate compression
22                             quality using the truncated bitstream.

```

Examples:

1. Produce a truncated version of a bitstream using 10% of the original length:
`./bin/sperr3d_trunc --pct 10 -o ./stream.10 ./bitstream`
2. Perform the task above, plus evaluate compression quality using the truncated bitstream:
`./bin/sperr3d_trunc --pct 10 -o ./stream.10 --orig64 ./data.f64 ./bitstream`

SPERR bitstreams without using multi-chunks (i.e., `--dims` equals `--chunks` in 3D, and all 2D cases) can safely be truncated by any tool (e.g., `head` on Linux), not necessarily using `sperr3d_trunc`.

5.2 C++ Interface

5.2.1 2D Compression and Decompression

C++ class `sperr::SPECK2D_FLT` is responsible for 2D compression and decompression. The sample code walks through necessary steps to perform a compression and decompression task, and a more concrete example can be found [here](#).

```

1 //
2 // Example of using a sperr::SPECK2D_FLT() to compress a 2D slice.
3 // This is a 6-step process.
4 //
5 #include "SPECK2D_FLT.h"
6
7 // Step 1: create an encoder:
8 auto encoder = sperr::SPECK2D_FLT();
9
10 // Step 2: specify the 2D slice dimension (the third dimension is left with 1):
11 encoder.set_dims({128, 128, 1});
12
13 // Step 3: copy over the input data from a raw pointer (float* or double*):
14 encoder.copy_data(ptr, 16'384); // 16,384 is the number of values.
15 // Step 3 alternative: one can hand a memory buffer to the encoder to avoid a memory copy;
16 // use either version is cool.

```

```

17 encoder.take_data(std::move(input)); // input is of type std::vector<doubles>.
18
19 // Step 4: specify the compression quality measured in one of three metrics;
20 // only the last invoked quality metric is honored.
21 encoder.set_tolerance(1e-9); // PWE tolerance = 1e-9
22 encoder.set_bitrate(2.2); // Target bitrate = 2.2
23 encoder.set_psnr(102.2); // Target PSNR = 102.2
24
25 // Step 5: perform the compression task:
26 encoder.compress();
27
28 // Step 6: retrieve the compressed bitstream:
29 auto bitstream = std::vector<uint8_t>();
30 encoder.append_encoded_bitstream(bitstream);
31
32 //
33 // Example of using a sperr::SPECK2D_FLT() to decompress a bitstream.
34 // This is a 5-step process.
35 //
36 #include "SPECK2D_FLT.h"
37
38 // Step 1: create a decoder:
39 auto decoder = sperr::SPECK2D_FLT();
40
41 // Step 2: specify the 2D slice dimension (the third dimension is left with 1):
42 // This information is often saved once somewhere for many same-sized slices.
43 decoder.set_dims({128, 128, 1});
44
45 // Step 3: pass in the compressed bitstream as a raw pointer (uint8_t*):
46 decoder.use_bitstream(ptr, 16'384); // 16,384 is the length of the bitstream.
47
48 // Step 4: perform the decompression task:
49 decoder.decompress(multi_res); // a boolean, if to enable multi-resolution decoding
50
51 // Step 5: retrieve the decompressed volume:
52 std::vector<double> vol = decoder.view_decoded_data();
53 auto hierarchy = decoder.view_hierarchy(); // if multi-resolution was enabled
54 // Step 5 alternative: one can take ownership of the data buffer to avoid a memory copy.
55 std::vector<double> vol = decoder.release_decoded_data();
56 auto hierarchy = decoder.release_hierarchy(); // if multi-resolution was enabled

```

5.2.2 3D Compression and Decompression

C++ class `sperr::SPERR3D_OMP_C` is responsible for 3D compression, and `sperr::SPERR3D_OMP_D` is responsible for 3D decompression. The sample code walks through necessary steps to perform a compression and decompression task, and a more concrete example can be found [here](#).

```

1 //
2 // Example of using a sperr::SPERR3D_OMP_C() to compress a 3D volume.
3 // This is a 6-step process.
4 //
5 #include "SPERR3D_OMP_C.h"
6
7 // Step 1: create an encoder:
8 auto encoder = sperr::SPERR3D_OMP_C();

```

```

9
10 // Step 2: specify the volume and chunk dimensions, respectively:
11 encoder.set_dims_and_chunks({384, 384, 256}, {192, 192, 128});
12
13 // Step 3: specify the number of OpenMP threads to use:
14 encoder.set_num_threads(4);
15
16 // Step 4: specify the compression quality measured in one of three metrics;
17 // only the last invoked quality metric is honored.
18 encoder.set_tolerance(1e-9); // PWE tolerance = 1e-9
19 encoder.set_bitrate(2.2); // Target bitrate = 2.2
20 encoder.set_psnr(102.2); // Target PSNR = 102.2
21
22 // Step 5: perform the compression task:
23 // The input data is passed in in the form of a raw pointer (float* or double*),
24 // and the total number of values will be passed in here too.
25 encoder.compress(ptr, 384 * 384 * 256);
26
27 // Step 6: retrieve the compressed bitstream:
28 std::vector<uint8_t> stream = encoder.get_encoded_bitstream();
29
30 //
31 // Example of using a sperr::SPERR3D_OMP_D() to decompress a bitstream.
32 // This is a 5-step process.
33 //
34 #include "SPERR3D_OMP_D.h"
35
36 // Step 1: create a decoder:
37 auto decoder = sperr::SPERR3D_OMP_D();
38
39 // Step 2: specify the number of OpenMP threads to use:
40 decoder.set_num_threads(4);
41
42 // Step 3: pass in the compressed bitstream as a raw pointer (uint8_t*):
43 decoder.use_bitstream(ptr, 16*384); // 16,384 is the length of the bitstream.
44
45 // Step 4: perform the decompression task:
46 // Note that the pointer to the bitstream is passed in again!
47 decoder.decompress(ptr, multi_res); // a boolean, if to enable multi-resolution decoding
48
49 // Step 5: retrieve the decompressed volume:
50 auto [dimx, dimy, dimz] = decoder.get_dims(); // dimension of the volume
51 std::vector<double> vol = decoder.view_decoded_data();
52 auto hierarchy = decoder.view_hierarchy(); // if multi-resolution was enabled
53 // Step 5 alternative: one can take ownership of the data buffer to avoid memory copies.
54 std::vector<double> vol = decoder.release_decoded_data();
55 auto hierarchy = decoder.release_hierarchy(); // if multi-resolution was enabled

```

To achieve higher performance with repeated compression and decompression tasks, the encoder and decoder objects are better to be re-used rather than repeatedly destroyed and created.

5.3 C Interface

SPERR provides a C wrapper with a set of C functions. All of the C interface is in the header file `SPERR_C_API.h`, which itself documents the C functions and parameters, etc. The following example code walks through key steps to use the C API to perform compression and decompression, while more concrete examples are available for [2D](#) and [3D](#) cases.

5.3.1 Example: 2D

```
1  /*
2  * Example of using the SPERR C API to perform 2D compression and decompression tasks.
3  */
4  #include "SPERR_C_API.h"
5
6  /* Step 1: create variables to keep the output: */
7  void* stream = NULL; /* caller is responsible for free'ing it after use. */
8  size_t stream_len = 0;
9
10 /* Step 2: call the 2D compression function:
11  * Assume that we have a buffer of 128 * 128 floats (in float* type) to be compressed,
12  * using PWE tolerance = 1e-3.
13  */
14 int ret = sperr_comp_2d(ptr,          /* memory buffer containing the input */
15                        1,            /* the input is of type float; 0 means double. */
16                        128,          /* dimx */
17                        128,          /* dimy */
18                        3,            /* compression mode; 3 means fixed PWE */
19                        1e-3,         /* actual PWE tolerance */
20                        0,            /* not using a header for the output bitstream */
21                        &stream,      /* will hold the compressed bitstream */
22                        &stream_len); /* length of the compressed bitstream */
23 assert(ret == 0);
24
25 /*
26  * Now that the 2D compression is completed, one can decompress the bitstream to
27  * retrieve the raw values, as the rest of this example shows.
28  */
29
30 /* Step 3: create a pointer to hold the decompressed values: */
31 void* output = NULL; /* caller is responsible for free'ing it after use. */
32
33 /* Step 4: call the 2D decompression function: */
34 int ret2 = sperr_decomp_2d(stream,    /* compressed bitstream */
35                            stream_len, /* compressed bitstream length */
36                            1,          /* decompress to floats. 0 means to doubles. */
37                            128,        /* dimx */
38                            128,        /* dimy */
39                            &output);  /* decompressed data is stored here */
40 assert(ret2 == 0);
41 free(output); /* cleanup */
42 free(stream); /* cleanup */
```

5.3.2 Example: 3D

```
1  /*
2   * Example of using the SPERR C API to perform 3D compression and decompression tasks.
3   */
4  #include "SPERR_C_API.h"
5
6  /* Step 1: create variables to keep the output: */
7  void* stream = NULL; /* caller is responsible for free'ing it after use. */
8  size_t stream_len = 0;
9
10 /* Step 2: call the 3D compression function:
11  * Assume that we have a buffer of 256^3 floats (in float* type) to be compressed,
12  * using PWE tolerance = 1e-3 and chunk dimension of 128^3.
13  */
14 int ret = sperr_comp_3d(ptr,          /* memory buffer containing the input */
15                        1,            /* the input is of type float; 0 means double. */
16                        256,          /* dimx */
17                        256,          /* dimy */
18                        256,          /* dimz */
19                        128,          /* chunk_x */
20                        128,          /* chunk_y */
21                        128,          /* chunk_z */
22                        3,            /* compression mode; 3 means fixed PWE */
23                        1e-3,         /* actual PWE tolerance */
24                        4,            /* use 4 OpenMP threads */
25                        &stream,      /* will hold the compressed bitstream */
26                        &stream_len); /* length of the compressed bitstream */
27 assert(ret == 0);
28
29 /*
30  * Now that the 3D compression is completed, one can decompress the bitstream to
31  * retrieve the raw values, as the rest of this example shows.
32  */
33
34 /* Step 3: create a pointer to hold the decompressed values,
35  * and also variables to hold the volume dimensions.
36  */
37 void* output = NULL; /* caller is responsible for free'ing it after use. */
38 size_t dimx = 0, dimy = 0, dimz = 0;
39
40 /* Step 4: call the 3D decompression function: */
41 int ret2 = sperr_decomp_3d(stream,    /* compressed bitstream */
42                            stream_len, /* compressed bitstream length */
43                            1,         /* decompress to floats. 0 means to doubles. */
44                            4,         /* use 4 OpenMP threads */
45                            &dimx,    /* dimx of the decompressed volume */
46                            &dimy,    /* dimy of the decompressed volume */
47                            &dimz,    /* dimz of the decompressed volume */
48                            &output); /* decompressed data is stored here */
49 assert(ret2 == 0);
50 free(output); /* cleanup */
51 free(stream); /* cleanup */
```

6 LC Framework

contact: [Martin Burtscher](#)

LC is a framework for automatically generating high-speed lossless and guaranteed-error-bounded lossy data compression and decompression algorithms. It supports CPUs and GPUs.

The framework code and tutorial are available at <https://github.com/burtscher/LC-framework/>

6.1 Overview

LC consists of the following three parts:

- Component library
- Preprocessor library
- Framework

Both libraries contain separate encoders and decoders for CPU and GPU execution. The user can extend these libraries. The framework takes preprocessors and components from these libraries and chains them into a pipeline to build a compression algorithm. It similarly chains the corresponding decoders in the opposite order to build the matching decompression algorithm. Moreover, the framework can automatically search for effective algorithms for a given input file or set of files by testing user-selected sets of components in each pipeline stage.

6.2 Quick-Start Guide and Tutorial

6.2.1 Installation

To download LC, run the following Linux commands:

```
git clone https://github.com/burtscher/LC-framework.git
cd LC-framework/
```

If you want to run LC on the CPU, generate the framework as follows:

```
./generate_Host_LC-Framework.py
```

If, instead, you want to run LC on the GPU, generate the framework as follows:

```
./generate_Device_LC-Framework.py
```

In either case, run the printed command to compile the generated code. For the CPU, use:

```
g++ -O3 -march=native -fopenmp -DUSE_CPU -I. -std=c++17 -o lc lc.cpp
```

For the GPU, use:

```
nvcc -O3 -arch=sm_70 -DUSE_GPU -Xcompiler "-O3 -march=native -fopenmp" -I. -o lc lc.cu
```

You may have to adjust these commands and flags to your system and compiler. For instance, the `sm_70` should be changed to match your GPU's compute capability.

The `generate_Hybrid_LC-Framework.py` script is only for testing and should not be used as it generates slow code.

6.2.2 Usage Examples for Lossless Compression Algorithms

The following examples assume you have a file called *input.dat* in the current directory and want to find a good compression algorithm for it. See below for a description of the available preprocessors and components. Assume you believe that using bit shuffling (BIT) and run-length encoding (RLE) at 4-byte granularity make a good compressor. Then you can see how well it compresses by entering the following command (note the two pairs of quotes):

```
./lc input.dat CR "" "BIT_4 RLE_4"
```

This will produce output that lists the compression ratio at the end. If you want to see whether running the RLE component at 1-byte granularity performs better, try:

```
./lc input.dat CR "" "BIT_4 RLE_1"
```

To find out which components (and preprocessors) are available, simply run:

```
./lc
```

If you want to see more stats on the input and output data as well as throughput information in addition to the compression ratio, switch the mode from CR to AL:

```
./lc input.dat AL "" "BIT_4 RLE_1"
```

Note that using AL also turns on verification to make sure the decompressed data is bit-by-bit equivalent to the original data.

One of the key strengths of LC is its ability to automatically search for a good compression algorithm. For example, if you want LC to try all available components in the second stage, type:

```
./lc input.dat CR "" "BIT_4 .+"
```

The “+” is a regular expression that matches the names of all components in the library. You can use it to select any subset of the available components. Of course, you can also use a regular expression for the first pipeline stage by entering:

```
./lc input.dat CR "" ".+ .+"
```

This is not limited to two stages. To search for the best 3-stage pipeline, use:

```
./lc input.dat CR "" ".+ .+ .+"
```

Note that the search time increases exponentially with the number of stages. Before you perform a search, you can check the size of the search space using the PR mode as follows:

```
./lc input.dat PR "" ".+ .+ .+ .+"
```

The output lists the number of algorithms that will be tested as well as which components will be considered in each pipeline stage. If this number is too large, i.e., the search would take too long, try reducing the search space by limiting the number of components to be considered:

```
./lc input.dat CR "" "DIFF_4 .+ .+ R.+|C.+|H.+"
```

If available, we recommend using the GPU version of LC as it tends to be much faster than the CPU version. To further speed up the search, LC includes a genetic algorithm (GA) to quickly search for a good but not necessarily the best algorithm. If you want to run the GA to find a good pipeline with 5 stages, enter the following command:

```
./scripts/ga_search.py -s 5 input.dat
```

If you are interested in the throughput in addition to the compression ratio, use the EX mode of LC like this:

```
./lc input.dat EX "" ".+ .+"
```


The output includes the **Pareto front** at the end, allowing the user to pick the best algorithm for a given compression or decompression throughput. The six columns list the algorithm, the compression ratio, the CPU compression throughput, the CPU decompression throughput, the GPU compression throughput, and the GPU decompression throughput. The throughputs are given in gigabytes per second.

All CR and EX runs with more than one algorithm also write their results to a CSV file that can be opened with most spreadsheet applications to view and postprocess the results.

EX- and AL-runs with one algorithm write the compressed data to a file called LC.encoded and the decompressed data to a file called LC.decoded.

In summary, LC supports the following modes:

- AL: This mode provides the most detailed output but only works for a single pipeline.
- PR: This mode prints the search space and then quits.
- CR: This mode searches for the best compressing algorithm.
- EX: This mode searches for the algorithms on the Pareto front, taking into account both the compression ratio and the compression/decompression throughput.
- TS: This mode is for testing only and should not be used.

6.2.3 Usage Examples for Lossy Floating-Point Compression Algorithms

To generate lossy algorithms with LC, preprocessors are needed. They must be fully specified (no regular expressions are allowed) and cannot be searched for automatically as they require user-specified parameters such as the error bound.

To find a good lossy compression algorithm for IEEE-754 32-bit single-precision floating-point data that are quantized with a maximum point-wise absolute error bound of 0.01 and then losslessly compressed with three components, enter:

```
./lc input.dat CR "QUANT_ABS_0_f32(0.01)" ".+ .+ R.+|C.+|H.+"
```

To do the same with a point-wise relative error bound, use:

```
./lc input.dat CR "QUANT_REL_0_f32(0.01)" ".+ .+ R.+|C.+|H.+"
```

The preprocessors work with the CR, EX, and AL modes. However, since both EX and AL verify the result, the default lossless verification will likely fail for lossy compression. LC includes a set of verifiers that can be selected in lieu of the default verifier. For an point-wise absolute error bound of 0.001, use:

```
./lc input.dat EX "QUANT_ABS_0_f32(0.001)" ".+ R.+|C.+|H.+" "MAXABS_f32(0.001)"
```

See the `./verifiers/` directory for additional available verifiers or the description below.

These quantizers replace any lost bits with zeros. If you prefer those bits be replaced by random data to minimize autocorrelation, use:

```
./lc input.dat CR "QUANT_ABS_R_f32(0.01)" ".+ .+ R.+|C.+|H.+"
```

or

```
./lc input.dat CR "QUANT_REL_R_f32(0.01)" ".+ .+ R.+|C.+|H.+"
```

6.2.4 Standalone Compressor and Decompressor Generation

Once you have determined a good lossless or lossy compression algorithm (e.g., “TUPL4_1 RRE_1 CLOG_1”), you can generate a standalone compressor and a standalone decompressor that are optimized for this algorithm.

To generate the CPU version, run:

```
./generate_standalone_CPU_compressor_decompressor.py "" "TUPL4_1 RRE_1 CLOG_1"
```

To generate the GPU version, run:

```
./generate_standalone_GPU_compressor_decompressor.py "" "TUPL4_1 RRE_1 CLOG_1"
```

In either case, run the printed commands to compile the generated code. For the CPU, use:

```
g++ -O3 -march=native -fopenmp -I. -std=c++17 -o compress compressor-standalone.cpp
g++ -O3 -march=native -fopenmp -I. -std=c++17 -o decompress decompressor-standalone.cpp
```

For the GPU, use:

```
nvcc -O3 -arch=sm_70 -DUSE_GPU -Xcompiler "-march=native -fopenmp" -I. -o compress compressor-standalone.cpp
nvcc -O3 -arch=sm_70 -DUSE_GPU -Xcompiler "-march=native -fopenmp" -I. -o decompress decompressor-standalone.cpp
```

You may have to adjust these commands and flags to your system and compiler. For instance, the `sm_70` should be changed to match your GPU's compute capability.

At this point, you can compress files with:

```
./compress input_file_name compressed_file_name [y]
```

and decompress them with:

```
./decompress compressed_file_name decompressed_file_name [y]
```

Both commands accept an optional “y” parameter at the end. If it is specified, the compressor and decompressor will measure and print the throughput.

6.3 Available Components, Preprocessors, and Verifiers

All LC compression pipelines start with zero or more preprocessors and end with one or more components. The preprocessors require parentheses after their names and may take parameters. The components do not take any parameters and cannot have parentheses.

6.3.1 Available Components

The LC framework breaks the input data up into chunks of 16 kB, each of which is compressed independently and in parallel using the selected components. All components are lossless. Most components support different word sizes. The number at the end of their names indicates the word size in bytes. For example, “_4” means the word size is 4 bytes (e.g., ints or floats). To structure the description of the components, we group them into the following four categories: mutators, shufflers, predictors, and reducers. The goal of the first three types is to better expose patterns so that the reducer components can compress the data more effectively. Only reducer components make sense the the last pipeline stage.

Mutators computationally transform each value. This is done independently of other values and does not compress the data.

- **NUL**: This component performs the identity transformation, meaning it outputs the input verbatim. It is useful in that it allows longer pipelines to also cover algorithms with fewer stages.
- **TCMS**: This component converts each value from twos-complement to magnitude-sign representation, which is often easier to compress because it tends to yield more leading zero bits.
- **DBEFS**: This component operates on IEEE-754 floating-point values. It first de-biases the exponent and then rearranges the data fields from sign, exponent, fraction order to (de-biased) exponent, fraction, sign order.
- **DBESF**: This component operates on IEEE-754 floating-point values. It first de-biases the exponent and then rearranges the data fields from sign, exponent, fraction order to (de-biased) exponent, sign, fraction order.

Shufflers rearrange the order of the values but perform no computation on them. Some shufflers reorder the bits or bytes within a word. None of them compress the data.

- **BIT:** This component is often referred to as "bit shuffle" or "bit transpose". It takes the most significant bit of each value in the input and outputs them together, then it takes the second most significant bit of each value and outputs them, and so on down to the least significant bit. This improves compressibility if the values tend to have the same bits in certain positions.
- **TUPLk:** This component assumes the data to be a sequence of k-tuples, which it rearranges by listing all first tuple values, then all second tuple values, and so on. For example, a tuple size of $k = 3$ changes the linear sequence $x_1, y_1, z_1, x_2, y_2, z_2, x_3, y_3, z_3, x_4, y_4, z_4$ into $x_1, x_2, x_3, x_4, y_1, y_2, y_3, y_4, z_1, z_2, z_3, z_4$. This is beneficial as values belonging to the same "dimension" often correlate more with each other than with other values from within the same tuple.

Predictors guess the next value by extrapolating it from prior values and then subtracting the prediction from the actual value, which yields a residual sequence. If the predictions are accurate, the residuals cluster around zero, making them easier to compress than the original data. Predictors per se do not compress the data.

- **DIFF:** This component computes the difference sequence (also called "delta modulation") by subtracting the previous value from the current value and outputting the resulting difference. If neighboring values correlate with each other, this tends to produce a more compressible sequence.
- **DIFFMS:** This component computes the difference sequence like DIFF does but outputs the result in sign-magnitude format, which is often more compressible because it tends to produce values with many leading zero bits.

Reducers are the only components that can compress the data. They exploit various types of redundancies to do so.

- **CLOG:** This component breaks the data up into 32 subchunks, determines the smallest amount of leading zero bits of all values in a subchunk, records this count, and then stores only the remaining bits of each value. This compresses data with leading zero bits.
- **HCLOG:** This component works like CLOG except it first applies the TCMS transformation to all values in a subchunk that yield no leading zero bits when using CLOG.
- **RLE:** This component performs run-length encoding. It counts how many times a value appears in a row. Then it counts how many non-repeating values follow. Both counts are emitted and followed by a single instance of the repeating value as well as all non-repeating values.
- **RRE:** This component creates a bitmap in which each bit specifies whether the corresponding word in the input is a repetition of the prior word or not. It outputs the non-repeating words and a compressed version of the bitmap that is repeatedly compressed with the same algorithm.
- **RZE:** This component creates a bitmap in which each bit specifies whether the corresponding word in the input is zero or not. It outputs the non-zero words and a compressed version of the bitmap like RRE does.

6.3.2 Available Preprocessors

Preprocessors operate on the entire data (i.e., there is no chunking) and can be lossy or lossless. Some preprocessors support different data types. The end of their names indicates the data type for which they are designed. For example, "`__f32`" means the preprocessor targets 32-bit floating-point values. To structure the description of the preprocessors, we group them into lossy or lossless preprocessors.

Lossless

These quantizers support INFs and NaNs. The end of the quantizer name indicates the data type for which it is designed.

- NUL: This preprocessor performs the identity transformation, meaning it outputs the input verbatim. It takes no parameters.
- LOR1D: This preprocessor performs a 1-dimensional Lorenzo transformation, i.e., it computes a difference sequence.

6.3.3 Lossy Quantizers

All quantizers require a parameter that specifies the maximally allowed error bound EB. They take an optional second parameter specifying a threshold. Any value whose magnitude is at or above the threshold is compressed losslessly and not quantized. The quantizers support INFs and NaNs. The end of the quantizer name indicates the data type for which it is designed.

- QUANT_ABS_0: These preprocessors quantize 32- and 64-bit floating-point values based on the provided point-wise absolute error bound. All values that end up in the same quantization bin are decompressed to the same value. These preprocessors guarantee that the original value V is decoded to a value V' such that $V - EB \leq V' \leq V + EB$.
- QUANT_ABS_R: These preprocessors quantize 32- and 64-bit floating-point values based on the provided point-wise absolute error bound. Each value from the same quantization bin is decompressed to a random value within the provided error bound to minimize autocorrelation. These preprocessors guarantee that the original value V is decoded to a value V' such that $V - EB \leq V' \leq V + EB$.
- QUANT_R2R: These preprocessors quantize 32- and 64-bit floating-point values just like their QUANT_ABS counterparts except the provided error bound is first multiplied by the range of values occurring in the input, where the range is the maximum value minus the minimum value.
- QUANT_REL_0: These preprocessors quantize 32- and 64-bit floating-point values based on the provided point-wise relative error bound. All values that end up in the same quantization bin are decompressed to the same value. These preprocessors guarantee that the original value V is decoded to a value V' with the same sign such that $|V| / (1 + EB) \leq |V'| \leq |V| * (1 + EB)$.
- QUANT_REL_R: These preprocessors quantize 32- and 64-bit floating-point values based on the provided point-wise relative error bound. Each value from the same quantization bin is decompressed to a random value within the provided error bound to minimize autocorrelation. These preprocessors guarantee that the original value V is decoded to a value V' with the same sign such that $|V| / (1 + EB) \leq |V'| \leq |V| * (1 + EB)$.

6.3.4 Available Verifiers

Some verifiers support different data types. The end of their names indicates the data type for which they are designed.

- LOSSLESS: This verifier is the default. It passes verification if the decompressed output matches every bit of the original input.
- PASS: This verifier always passes verification and is only useful for debugging.
- MAXABS: This verifier takes a point-wise absolute error bound as parameter and only passes verification if every output value is within the specified error bound.
- MAXR2R: This verifier works like MAXABS except the provided error bound is first multiplied by the range of values occurring in the input, where the range is the maximum value minus the minimum value.
- MAXREL: This verifier takes a point-wise relative error bound as parameter and only passes verification if every output value is within the specified error bound.

- MSE: This verifier takes a mean squared error as parameter and only passes verification if the mean squared error of the output values is within the error bound.
- PSNR: This verifier takes a peak-signal-to-noise ratio (PSNR) as parameter and only passes verification if the PSNR of the output values is above the specified lower bound.

For learning how to add your own components and preprocessors, please see <https://github.com/burtscher/LC-framework/?tab=readme-ov-file#adding-and-removing-components>.

7 TEZip

TEZip is not installed in the software environment due to its comparatively complex dependencies and compatibility with the GPUs in the testing environment

contact: [Amarjit Singh](#)

TEZip(Time Evolutionary Zip) is developed in RIKEN R-CCS and designed to compress time evolutionary data by using a deep learning for prediction. TEZip compression/decompression procedures consist of three steps, model training, compression and decompression. For more details, please refer to the readthedocs document [full documentation](#)

7.1 CLI Interface

```
1 # 1. Creating training data
2 # Comandline options
3 # [Learning_image_directory]: The directory path containing the training images to be
4 #   dumped into the Hkl file (e.g., ./data)
5 # [Output_directory]: Directory path to output Hkl file (e.g., ./data_hkl).
6 # -v: Specifies the path of the directory used for verification among the directories
7 #   specified in the first argument. Without this option, randomly determined.
8 python train_data_create.py [Learning_image_directory] [Output_directory]
9
10 # 2. Model Training
11 # Comandline options
12 # -l: Execute learning mechanism.
13 #   [Output_directory] is "Path to the output directory of the model"
14 #   [Directory_for_training_data] is "Path to the training data directory(.hkl)"
15 # -f: Forced CPU mode flag
16 #   "-f" to the runtime will disable the GPU and force it to run on the CPU
17 # -v: Flag for screen output
18 #   "-v" to the runtime, the learning status, such as losses and epochs during
19 #   learning, will be output to the console.
20 python tezip.py -l [Output_directory] [Directory_for_training_data]
21
22 # 3. Compression
23 # Comandline options
24 # -c: Run the compression mechanism
25 #   [Model_directory] Path of the directory of trained models
26 #   [Directory_of_images_to_be_compressed] Directory path of the image to be
27 #   compressed
28 #   [Output_directory] Output directory path for compressed data
29 # -w: Specifying the keyframe switching criteria
30 #   SWP(Static Window-based Prediction)to specify how many keyframes to infer from
31 #   one keyframe of execution. If it is specified at the same time as -t it will
32 #   cause an error termination
33 # -t: Specify the criteria for keyframe switching
34 #   Specify the threshold value of MSE(Mean Square Error) for execution switching in
35 #   DWP(Dynamic Window-based Prediction). If it is specified at the same time as
36 #   -w it will cause an error termination
37 # -p: Number of images for warm-up
38 #   The more keyframes you specify for LSTM recording the larger the
39 #   "key_frame.dat" will be and the smaller the "entropy.dat" will likely be.
40 #   However, if you set the number of keyframes to 0 or 1 when running DWP the MSE
```

```

41 # will become larger and the final number of keyframes may become larger
42 # -m: Selecting an error-bound mechanism
43 # Select the error bouncing mechanism from the following four types
44 #     abs absolute error bound
45 #     rel relative bound ratio
46 #     absrel Do both of the above
47 #     pwrrel point wise relative error bound
48 # If you select multiple items or select non-existent items the #program will
49 # exit with an error
50 # -b: Threshold for error bouncing mechanism
51 # Specify the tolerance threshold of the error bouncing mechanism. If "-m" is
52 # specified as absrel, enter two values.
53 # If an inappropriate number of inputs are given for the one specified by -m
54 # the program will exit with an error.If the input contains 0 the error
55 # bouncing mechanism will not be executed and the data will be fully lossy
56 # compressed.
57 # -f: Forced CPU mode flag
58 # By adding -f to the runtime you can disable the GPU and force it to run on
59 # the CPU
60 # -v: Flag for screen output
61 # When -v is added at runtime the status during execution such as the value
62 # of MSE after inference and the time taken for the compression process will be
63 # output to the console
64 # -n: Flag to disable Entropy Coding for compression process
65 python tezip.py -c [Model_directory] [Directory_of_images_to_be_compressed]
66 [Output_directory] -p [Number_of_warm-up_sheets] -w or -t [-w
67 Number_of_inferences_to_be_made_from_a_single_keyframe , -t
68 MSE_threshold_for_keyframe_switching] -m [Error-bound_mechanism_name] -b
69 [Threshold_for_error_bouncing_mechanism]
70
71 # 4. Decompressoion
72 # Comandline options
73 # -u: Run the learning mechanism
74 #     [Model_directory] Directory path of trained models
75 #     [Directory_of_compressed_data] Directory path for compressed data (.dat), etc.
76 #     [Output_directory] Output directory path for unzipped data
77 # -f: Forced CPU mode flag
78 # By adding -f to the runtime, you can disable the GPU and force it to run on
79 # the CPU.
80 # -v: Flag for screen output
81 # By adding -v at runtime, the processing time during decompression is output
82 # to the console
83 python tezip.py -u [Model_directory] [Directory_of_compressed_data] [Output_directory]
84

```

7.2 C/C++ Interface

TEZip does not have its own C/C++ interface but may be called via LibPressio's C/C++.

8 DCTZ

contact: [Seung Woo Son](#)

DCTZ is a lossy compressor based on DCT (discrete cosine transform) tailored to work with floating point datasets, single- or double-precision, with an error-bound capability (primarily on relative error bound). The compression pipeline has been adapted from JPEG.

8.1 CLI Interface

DCTZ has a single test program that goes through compression/decompression as a single test scenario. There are two versions of test programs: `dctz-ec-test` and `dctz-qt-test`. `dctz-ec-test` is a more conservative compressor, whereas `dctz-qt-test` includes an extra quantization process to improve compression ratios while guaranteeing the user-defined error bounds. The test scripts, `test-dctz.sh` or `test-dctz-f.sh`, for a couple of sample datasets are available under the “tests” folder in the DCTZ repository.

```
1  # print help for dctz compressor
2  # both dctz-ec-test and dctz-qt-test has the same CLI
3  dctz-ec-test --help
4  Test case: ../dctz-ec-test -d|-f [err bound] [var name] [srcFilePath] [dimension sizes...]
5
6  # commandline options
7  # precision: -d | -f                double or float
8  # error bound: 1E-3, 1E-4, or 1E-5
9  # var name: unused but
10 # data file path
11 # dimensions list
12
13 # sample CLI for evaluating the CESM-ATM-taylor
14 # the error bound is set to 1E-3
15 dctz-qt-test -f 1E-3 var CESM-ATM-tylor/1800x3600/CLDHGH_1_1800_3600.dat 3600 1800
16
17 # The above command will generate a compressed file named CLDHGH_1_1800_3600.dat.z
18 # and the decompressed file named CLDHGH_1_1800_3600.dat.z.r.
19
```


9 LibPressio

contact: [Robert Underwood](#)

LibPressio provides a generic abstraction across the various compressors including ALL of the compressors listed above. It will eventually serve as the front end to FZ. In general, all features are supported from all interfaces.

The material presented here [adapts a subset of the LibPressio tutorial](#). The full tutorial covers GPU based compressors, using compressors with user provided metrics, usage with I/O libraries, automatic optimization of compressors with OptZConfig, and much more. You can also find [full documentation](#) and more [creatives uses of LibPressio](#).

Below we present the basics for the commandline, Python, C++, and C interfaces. We also have support for R, Julia, and Rust.

9.1 CLI interface

LibPressio allows you to rapid experiment with various compressors

```
1 #print help for the libpressio command
2 pressio
3
4 # load a 3d array of float, use SZ3 with an absolute error bound of 1e-4
5 # and collect various metrics and print them out
6 pressio \
7     -i /usecases/Hurricane/Uf48.dat -d 500 -d 500 -d 100 -t float
8     -b compressor=sz3 -o abs=1e-4
9     -m error_stat -m size -m time -M all
10
11 # same but with an HDF5 file
12 pressio \
13     -i /path/to/foo.hdf5 -I /some/dataset
14     -b compressor=sz3 -o abs=1e-4
15     -m error_stat -m size -m time -M all
16
17 # same but with an numpy file
18 pressio \
19     -i /path/to/foo.npy -T numpy
20     -b compressor=sz3 -o abs=1e-4
21     -m error_stat -m size -m time -M all
22
23
24 # same but with ZFP, and just error statistics
25 pressio \
26     -i /usecases/Hurricane/Uf48.dat -d 500 -d 500 -d 100 -t float
27     -b compressor=zfp -o abs=1e-4
28     -m error_stat -M all
29
30 # use ZFP's fixed rate mode
31 pressio \
32     -i /usecases/Hurricane/Uf48.dat -d 500 -d 500 -d 100 -t float
33     -b compressor=zfp -o zfp:rate=13
34     -m error_stat -M all
35
36 # SZ3 but with a REL error bound
37 pressio \
38     -i /usecases/Hurricane/Uf48.dat -d 500 -d 500 -d 100 -t float
```

```

39     -b compressor=zfp -o rel=1e-4
40     -m error_stat -M all
41
42     #print the builtin help for ZFP
43     pressio -a help -b compressor=zfp
44
45     #print the settings configured for ZFP
46     pressio \
47         -a settings -O all \
48         -b compressor=zfp -o rel=1e-4

```

9.2 Python Interface

You can use `pressio_new` client: `python sz3 /usecases/Hurricane/Uf48.dat` to generate a starter code. You can run `pressio_new` to see the list of available starter codes. The code that appears below is a commented version of what this tool outputs with one line added to compute error statistics:

```

1  import libpressio as lp
2  import numpy as np
3  from pprint import pprint
4
5  # read in the file as a binary array
6  input = np.fromfile("/usecases/Hurricane/Uf48.dat", dtype=np.float32).reshape(100, 500, 500)
7
8  # configure the compressor
9  compressor = lp.PressioCompressor.from_config({
10     "compressor_id": "pressio",
11     "early_config": {
12         "pressio:compressor": "sz3"
13     },
14     "compressor_config": {
15         "pressio:abs": 1e-5
16     }
17 })
18
19 # run the compressor
20 decompressed = input.copy()
21 compressed = compressor.encode(input)
22 decompressed = compressor.decode(compressed, decompressed)
23
24 # collect metrics results
25 pprint(compressor.get_metrics())

```

9.3 C++ Interface

You can use `pressio_new` client: `cpp sz3 /usecases/Hurricane/Uf48.dat` to generate a starter code. You can run `pressio_new` to see the list of available starter codes. The code that appears below is a commented version of what this tool outputs with one line added to compute error statistics:

```

1  #include <libpressio_ext/cpp/libpressio.h>
2  #include <libpressio_meta.h>
3  #include <iostream>
4  #include <string>
5
6  using namespace std::string_literals;

```

```

7
8 int main() {
9     //load 3rd party compressors definitions. These support compressors with licences that
10    //prevent it from being used with LibPressio directly (e.g. the GPL or a proritary licence)
11    pressio library;
12    libpressio_register_all();
13    pressio_compressor compressor = library.get_compressor("pressio");
14
15    //the data is a 3d array of type float. 100 is the slowest advancing dimension
16    //this is called column major order or fortran order
17    pressio_dtype type = pressio_float_dtype;
18    std::vector<size_t> dims {500,500,100};
19    pressio_data metadata = pressio_data::owning(type, dims);
20
21    // while using an I/O plugin is overkill here, we can change posix to "hdf5"
22    // to load an HDF5 dataset, or "numpy" to read a .npy file.
23    pressio_io io = library.get_io("posix");
24    if(!io) {
25        std::cerr << library.err_msg() << std::endl;
26        exit(library.err_code());
27    }
28    if(io->set_options({
29        {"io:path", "/usecases/Hurricane/Uf48.dat"s},
30    })) {
31        std::cerr << io->error_msg() << std::endl;
32        exit(io->error_code());
33    }
34    pressio_data* input = io->read(&metadata);
35    pressio_data compressed = pressio_data::empty(pressio_byte_dtype, {});
36    pressio_data output = pressio_data::owning(type, dims);
37
38    //configure the compressor to use SZ3 with an absolute error bound.
39    //we can change sz3 to zfp here and choose what compressor to use
40    if(compressor->set_options({
41        {"pressio:compressor", "sz3"},
42        {"pressio:abs", 1e-5}
43    })) {
44        std::cerr << io->error_msg() << std::endl;
45        exit(io->error_code());
46    }
47
48    //run the compressor
49    if(compressor->compress(input, &compressed) != 0) {
50        std::cerr << compressor->error_msg() << std::endl;
51        exit(compressor->error_code());
52    }
53    if(compressor->decompress(&compressed, &output) != 0) {
54        std::cerr << compressor->error_msg() << std::endl;
55        exit(compressor->error_code());
56    }
57
58    //collect the metrics results and print them them to the screen
59    std::cout << compressor->get_metrics_results() << std::endl;
60    delete input;

```

61 }

9.4 C interface

You can use `pressio_new client:c sz3 /usecases/Hurricane/Uf48.dat` to generate a starter code. You can run `pressio_new` to see the list of available starter codes. The code that appears below is a commented version of what this tool outputs with one line added to compute error statistics:

```
1  #include <stdlib.h>
2  #include <stdio.h>
3  #include <libpressio.h>
4  #include <libpressio_meta.h>
5
6  int main() {
7      struct pressio* library = pressio_instance();
8      //load 3rd party compressors definitions. These support compressors with licences that
9      //prevent it from being used with LibPressio directly (e.g. the GPL or a proritary licence)
10     libpressio_register_all();
11
12     //load the "pressio" meta compressor that provides some standardization about what
13     //compressors error bounds to support. For example, if a compressor support an ABS bound
14     //"pressio" will provide a REL bound for free
15     struct pressio_compressor* compressor = pressio_get_compressor(library, "pressio");
16     if(!compressor) {
17         fprintf(stderr, "%s\n", pressio_error_msg(library));
18         exit(pressio_error_code(library));
19     }
20
21     //the data is a 3d array of type float. 100 is the slowest advancing dimension
22     //this is called column major order or fortran order
23     enum pressio_dtype type = pressio_float_dtype;
24     size_t dims[] = {500,500,100};
25     struct pressio_data* metadata = pressio_data_new_owning(type,
26         sizeof(dims)/sizeof(dims[0]), dims);
27
28     // while using an I/O plugin is overkill here, we can change posix to "hdf5"
29     // to load an HDF5 dataset, or "numpy" to read a .npz file.
30     struct pressio_io* io = pressio_get_io(library, "posix");
31     if(!io) {
32         fprintf(stderr, "%s\n", pressio_error_msg(library));
33         exit(pressio_error_code(library));
34     }
35     struct pressio_options* io_opts = pressio_options_new();
36     pressio_options_set_string(io_opts, "io:path", "/usecases/Hurricane/Uf48.dat");
37     if(pressio_io_set_options(io, io_opts)) {
38         fprintf(stderr, "%s\n", pressio_io_error_msg(io));
39         exit(pressio_io_error_code(io));
40     }
41     pressio_options_free(io_opts);
42     struct pressio_data* input = pressio_io_read(io,metadata);
43     if(!input) {
44         fprintf(stderr, "%s\n", pressio_io_error_msg(io));
45         exit(pressio_io_error_code(io));
46     }
47     pressio_data_free(metadata);
```

```

48     pressio_io_free(io);
49
50     //setup output locations for the compressed and decompressed data
51     struct pressio_data* compressed = pressio_data_new_empty(pressio_byte_dtype, 0, NULL);
52     struct pressio_data* output = pressio_data_new_owing(type,
53         sizeof(dims)/sizeof(dims[0]), dims);
54
55     //configure the compressor to use SZ3 with an absolute error bound.
56     //we can change sz3 to zfp here and choose what compressor to use
57     struct pressio_options* compressor_opts = pressio_options_new();
58     pressio_options_set_string(compressor_opts, "pressio:compressor", "sz3");
59     //this line was added to compute error statistics when decompression is run
60     pressio_options_set_string(compressor_opts, "pressio:metric", "error_stat");
61     pressio_options_set_double(compressor_opts, "pressio:abs", 1e-5);
62     if(pressio_compressor_set_options(compressor, compressor_opts)) {
63         fprintf(stderr, "%s\n", pressio_compressor_error_msg(compressor));
64         exit(pressio_compressor_error_code(compressor));
65     }
66     pressio_options_free(compressor_opts);
67
68     //run the compressor
69     if(pressio_compressor_compress(compressor, input, compressed) != 0) {
70         fprintf(stderr, "%s\n", pressio_compressor_error_msg(compressor));
71         exit(pressio_compressor_error_code(compressor));
72     }
73     if(pressio_compressor_decompress(compressor, compressed, output) != 0) {
74         fprintf(stderr, "%s\n", pressio_compressor_error_msg(compressor));
75         exit(pressio_compressor_error_code(compressor));
76     }
77
78     //collect the metrics results and print them them to the screen
79     struct pressio_options* metrics = pressio_compressor_get_metrics_results(compressor);
80     char* metrics_str = pressio_options_to_string(metrics);
81     fprintf(stdout, "%s\n", metrics_str);
82     free(metrics_str);
83     pressio_options_free(metrics);
84
85     //free memory
86     pressio_data_free(input);
87     pressio_data_free(compressed);
88     pressio_data_free(output);
89     pressio_compressor_release(compressor);
90     pressio_release(library);
91 }

```