

# SPERR Interface and Examples

Samuel Li

February 15, 2024

## 1 Introduction

This document is the SPERR section of the handout for the [FZ workshop](#) hands-on session on February 15 2024, Sarasota, FL. It provides examples for the CLI interface, C++ interface, and C interface of SPERR (Section [2.1](#), [2.2](#), and [2.3](#), respectively).

## 2 SPERR

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

Repo: [github.com/NCAR/SPERR/](https://github.com/NCAR/SPERR/)

Wiki: [github.com/NCAR/SPERR/wiki](https://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.

### 2.1 CLI Interface

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

#### 2.1.1 `sperr2d`

`sperr2d` is responsible for compressing and decompressing a 2D data slice. In compression mode (`-c`), it takes as input a raw binary file consisting of single- or double-precision floating point values, and outputs a compressed bitstream. In decompression mode (`-d`), it takes as input a compressed bitstream, and outputs the decompressed binary file of single- or double-precision floating point values. 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 (choose one):
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  $512 \times 512$  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.

### 2.1.2 sperr3d

`sperr3d` is responsible for compressing and decompressing a 2D data volume. Similar to `sperr2d`, it operates in either compression (`-c`) or decompression (`-d`) mode, converting between raw binary floating-point values and compressed bitstreams. Different from `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 chunking step allows for compressing and decompressing all the small chunks individually and in parallel. `sperr3d` uses  $256^3$  as the default chunk size, but any number from dozens to low hundreds would work well. (Chunk dimensions that can divide the full volume are preferred, but not mandatory.) 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:



```
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

## 2.2 C++ Interface

### 2.2.1 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
```

```

7 // Step 1: create a decoder:
8 auto decoder = sperr::SPECK2D_FLT();
9
10 // Step 2: specify the 2D slice dimension (the third dimension is left with 1):
11 // This information is often saved once somewhere for many same-sized slices.
12 decoder.set_dims({128, 128, 1});
13
14 // Step 3: pass in the compressed bitstream as a raw pointer (uint8_t*):
15 decoder.use_bitstream(ptr, 16'384); // 16,384 is the length of the bitstream.
16
17 // Step 4: perform the decompression task:
18 decoder.decompress(multi_res); // a boolean, if to enable multi-resolution decoding
19
20 // Step 5: retrieve the decompressed volume:
21 std::vector<double> vol = decoder.view_decoded_data();
22 auto hierarchy = decoder.view_hierarchy(); // if multi-resolution was enabled
23 // Step 5 alternative: one can take ownership of the data buffer to avoid a memory copy.
24 std::vector<double> vol = decoder.release_decoded_data();
25 auto hierarchy = decoder.release_hierarchy(); // if multi-resolution was enabled

```

### 2.2.2 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();

```

```

1  //
2  // Example of using a sperr::SPERR3D_OMP_D() to decompress a bitstream.
3  // This is a 5-step process.
4  //
5  #include "SPERR3D_OMP_D.h"
6
7  // Step 1: create a decoder:
8  auto decoder = sperr::SPERR3D_OMP_D();
9
10 // Step 2: specify the number of OpenMP threads to use:
11 decoder.set_num_threads(4);
12
13 // Step 3: pass in the compressed bitstream as a raw pointer (uint8_t*):
14 decoder.use_bitstream(ptr, 16'384); // 16,384 is the length of the bitstream.
15
16 // Step 4: perform the decompression task:
17 // Note that the pointer to the bitstream is passed in again!
18 decoder.decompress(ptr, multi_res); // a boolean, if to enable multi-resolution decoding
19
20 // Step 5: retrieve the decompressed volume:
21 auto [dimx, dimy, dimz] = decoder.get_dims(); // dimension of the volume
22 std::vector<double> vol = decoder.view_decoded_data();
23 auto hierarchy = decoder.view_hierarchy(); // if multi-resolution was enabled
24 // Step 5 alternative: one can take ownership of the data buffer to avoid memory copies.
25 std::vector<double> vol = decoder.release_decoded_data();
26 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.

## 2.3 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.

### 2.3.1 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 */

```





```

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 */

```