
Using libRaptorQ library: RFC6330 interface

August 4, 2018

Abstract

libRaptorQ is a C++11 implementation of the RaptorQ Forward Error Correction. It includes two interfaces, one derived from the RFC6330 standard, and one RAW API for a more flexible and easy usage.

This document is only about the RFC6330 interface.

The implementation was started as a university laboratory project, and will be later used and included in Fenrir, the maintainer's master thesis project.

Contents

| | | |
|----------|--|-----------|
| 1 | Contacts | 3 |
| 2 | Build & install | 3 |
| 2.1 | Get the source code | 3 |
| 2.2 | Dependencies | 4 |
| 2.3 | Build & Install | 4 |
| 3 | Working with RaptorQ | 7 |
| 3.1 | Theory (you really need this) | 7 |
| 3.2 | RFC6330: Blocks & Symbols | 7 |
| 3.3 | C++ interface | 8 |
| 3.3.1 | Caching precomputations | 10 |
| 3.3.2 | Thread pool | 11 |
| 3.3.3 | The Encoder | 12 |
| 3.3.4 | Blocks | 14 |
| 3.3.5 | Symbols | 15 |
| 3.3.6 | The Decoder | 16 |
| 3.4 | C interface | 20 |
| 3.4.1 | Cache settings | 20 |
| 3.4.2 | C Constructors | 21 |
| 3.4.3 | Common functions for (de/en)coding | 22 |
| 3.4.4 | Encoding | 25 |
| 3.4.5 | Decoding | 27 |
| 4 | GNU Free Documentation License | 30 |

1 Contacts

The main development and discussions on the project, along with bug reporting, happens on the main website.

Mailing Lists Mailing lists are available at <https://fenrirproject.org/lists>
The two mailing lists are for development and announcements, but due to the low traffic of the development mailing list, it can also be used by users for questions on the project.

IRC Since there are not many developers for now, the main irc channel is **#fenrirproject** on freenode

2 Build & install

2.1 Get the source code

This document follows the 1.0 stable release.

You can get the tarballs from the main website here:

<https://fenrirproject.org/Luker/libRaptorQ/tags>

Or you can check out the repository (warning: development in progress):

```
$ git clone --recurse-submodules \
    https://fenrirproject.org/Luker/libRaptorQ.git
```

You can also get it from github:

```
$ git clone --recurse-submodules \
    https://github.com/LucaFulchir/libRaptorQ.git
```

GPG verification: Once you have cloned it, it's always a good thing to check the repository gpg signatures, so you can import my key with:

```
$ gpg --keyserver pgp.mit.edu --recv-key \
    7393DAD255BEB5751DBDA04AB11CD823BA278C85
```

Now you have the source, and the key, it's enough to check the signature of the last commit:

```
$ git log -n 1 --show-signature
```

The important part is that you get something like this:

```
gpg: Signature made Mon 11 Dec 2017 21:55:28 CET
gpg:                using RSA key 8F7474044095B405D0F042F0A2CCA134BC7C8572
gpg: Good signature from "Luca Fulchir <luker@fenrirproject.org>" [unknown]
gpg:                aka "Luca Fulchir <luca@fulchir.it>" [unknown]
gpg: WARNING: This key is not certified with a trusted signature!
gpg:                There is no indication that the signature belongs to the owner.
Primary key fingerprint: 7393 DAD2 55BE B575 1DBD A04A B11C D823 BA27 8C85
```

And as long as you got the right key, and you find the "**gpg: Good signature**" string, you can be sure you have the right code.

The main development happens in the *master* branch, while we keep one branch for each stable version.

2.2 Dependencies

libRaptorQ has 4 dependencies:

Eigen3 : This is used for matrix manipulation, which is a big part of RaptorQ.

lz4 : Used to compress cached precomputations.

git : This is used not only to get the source, but also by the build system. We get the last git commit id and feed it to clang or gcc as seed for their internal random number generator. This makes it possible to have reproducible builds.

optionparser : Library used to parse the command line easily in C++, without exceptions

All dependencies are included in the sources, so you do not need to download and compile them.

Eigen3 is only needed at build time, so there is no need to download it again. we use the 3.3.4 version.

LZ4 is included as a git submodule. if you do not have it, run:

```
$ git submodule init
$ git submodule update
```

and LZ4 will be built statically in the library (it will not be installed on the system)

2.3 Build & Install

libRaptorQ uses the cMake build system, so things are fairly standard:

```
$ cd libRaptorQ.git
$ mkdir build
$ cd build
$ cmake ../
$ make -j 4
```

By default, the libRaptorQ project tries to have deterministic builds. This means that if you compile things twice, or with two different computers, the hash of the resulting library will be the same, provided that the same compiler (clang, gcc 4.8, gcc 4.9 etc) was used. Currently the only exception is the clang compiler with the *PROFILING* option enabled, and this will not likely be solved.

There are lots of options, you can use in cmake. As always, you can change them by adding “**-Dcmake_option=cmake_value**” when calling cmake.

You can always use the cmake-gui or ccmake commands to have the list of possible options.

The ones we recognize are:

- LTO : ON/OFF. Default:**ON**. Enables *Link Time Optimizations* for clang or gcc. Makes libraries smaller and better optimized.
- PROFILING : ON/OFF. Default:**ON**. Profiling compiles everything once, then runs a test to see which code paths are used more, and then recompiles everything again, but making sure that the binary is optimized for those paths. Works with clang and gcc. Provides a slight speedup.
- CMAKE_C_COMPILER : gcc and clang are directly supported. other should work, too. This is only used if you want to build the C example.
- CMAKE_CXX_COMPILER : g++, clang++ are directly supported. other should work, too.
- CLANG_STDLIB : ON/OFF. Default:**OFF**.use clang’s libc++ standard library. Only available with clang.
- CMAKE_BUILD_TYPE : Type of build. you can choose between “Debug”, “Release”, “MinSizeRel” or “RelWithDebInfo”
- CMAKE_INSTALL_PREFIX : Default to */usr/local*. Change it to fit your distribution guidelines.
- RQ_ENDIANNESS : **Auto, BigEndian, LittleEndian** Force the build to use one endianness or the other
- RQ_LINKER : **Auto, gold, ld, bsd** Force the build to use one linker or the other
- USE_LZ4 : ON/OFF. Default:**ON**. compile with support for lz4

Then you can build everything by running:

```
$ make -j 4
```

Of course, you can configure the number of parallel jobs (the *-j* parameter) to be what you need.

Optional make targets: The following optional targets are also supported:

```
$ make docs tests examples  
$ make everything
```

The “docs” target builds this document, but you need latex with the refman style. The tests are only useful to check performance of rfc compliance right now. “examples” compiles the C and C++ examples, which will not be installed.

Install: The installation process is very simple:

```
$ make install DESTDIR=...
```

You can change the *DESTDIR* parameter to fit your distribution guidelines.

3 Working with RaptorQ

3.1 Theory (you really need this)

To be able to work with `libRaptorQ`, you must first understand how the RaptorQ algorithms works. We won't go into the details, but just what you need to be able to tune the algorithm to your needs.

Fountain codes: Fountain codes are a special *Forward-Error-Correcting* code class, which characteristic is simple: if you want to send K packets, you actually send $K + X$ packets, and the receiver only needs to get any K packets to be able to reconstruct your data. The number X of overhead packets can be as big as you need (theoretically infinite), so you can tune it to be slightly higher than the expected packet loss.

Systematic codes: RaptorQ is also a systematic code. This means that those first K packets are the input *as-is* (**source symbols**), and the X packets (**repair symbols**) have the information needed to recover any of the lost source packets. This has the big advantage of avoiding any kind of time and memory consuming decoding if there is no packet loss during the transmission.

Complexity: The RaptorQ algorithm is often presented as having a linear time encoder and decoder. This is both false and misleading. Generating the source or repair symbols from the intermediate symbols has linear complexity. Generating the intermediate symbols has cubic complexity on the number of symbols. Which is a completely different thing. It is still very quick. On a core i7, 2.4Ghz, you need to wait *0.4ms* for 10 symbols, *280ms* for 1.000 symbols, but it can take an hour for 27.000 symbols. RaptorQ handles up to 56.403 symbols.

Caching `libRaptorQ` can work with big matrices that take a lot of time to compute. For this reason the matrices can be saved once they have been computed the first time. `libRaptorQ` uses a **local** cache. The matrix can be compressed with LZ4.

3.2 RFC6330: Blocks & Symbols

To understand how to properly use and tune `libRaptorQ`, you first need to understand how RaptorQ and the **RFC6330** handles its inputs, outputs, and what the time and memory constraints are.

Input sequencing: The RFC6330 needs to have the whole input you want to send before it can start working.

This means that it might be a little more difficult to use in live-streaming contexts, or where you need real-time data, but for that you probably want to use the single-block API in the **RaptorQ** C++ namespace. There is a second document tracking that interface, this document tracks only the RFC6330 interface.

Once you have the whole input, RFC6330 divides it into **blocks**. Each block *is encoded and decoded independently* and will be divided into **symbols**. Each symbol *should* be transmitted separately in its own packet (if you are on the network).

Sizes: Each input can have up to *256 blocks*, each block can have up to *56.403 symbols*, and each symbol can have up to $2^{16} - 1$ *bytes* long. This gives a maximum files size of almost 881 GB (946270874880 bytes to be exact)

Interleaving: An other feature of RFC6330 is to automatically provide some interleaving of the input data before transmitting it. This means that one symbol will not represent one sequential chunk of your input data, but more likely it's the concatenation of different **sub-symbols**. The size of the subsymbol must thus be a fraction of the symbol size. *This feature is not used if you set the size of the subsymbol to the size of symbol.*

Memory and Time: Memory and time requirements are to be considered, though, as RFC6330 needs to run a cubic algorithm on matrix of size $K * K$, where K is the number of symbols in each block. The algorithm needs to keep in memory two of these matrices, although most of the work is done on only one. This is actually a lot. More benchmarks and optimizations will come later, for now remember that with 10 symbols it takes something like 0.4ms on a core i7 2.4GHZ, 280ms with 1000 symbols, and up to an hour with 27.000 symbols.

3.3 C++ interface

To use the C++ interface you only need to include the **RaptorQ/RFC6330_v1.hpp** header, and link **libRaptorQ** and your threading library (usually *libpthread*).

You can also use the library as a header-only include, by including **RaptorQ/RFC6330_v1_hdr.hpp**. All your code should keep working even if you switch between linked and header-only at a later date. If you use the header-only version, remember to have Eigen3 in your include path and to include the needed definitions, **RQ_LITTLE_ENDIAN** or **RQ_BIG_ENDIAN**

Namespace Since the libRaptorQ library now has two very different usages, everything was split into two namespaces: **RFC6330** and **RaptorQ**. The names should be explanatory: the first namespace (that we cover in this document) is for the RFC interface, while the second namespace is for the low-level interface.

Since we will cover only the **RFC6330** namespace here, remember that everything we talk about is in this namespace.

Both the RFC6330 and the RaptorQ namespace are versioned so you can keep using old APIs while we change stuff. just add `__v1`

to select the first version of the API. E.g: RFC6330__v1

Templates There are two main classes you will use:

```
template <typename Rnd_It, typename Fwd_It>
class Encoder
```

```
template <typename In_It, typename Fwd_It>
class Decoder
```

As you might guess, the classes for the encoder and decoder take two template parameters.

For the **Encoder**, the first parameter *MUST* be a *random access iterator* to the input data, and the second parameter is an *forward iterator*. The random access iterator will be used to scan your input data, and perform an interleaving (if you did not set the same size the symbol and to the subsymbol). The forward iterator will be used to write the data to your structure.

The same is done for the **Decoder**, but we do not need to do any interleaving on the input, so the first iterator can be just an input iterator, and nothing more.

3.3.1 Caching precomputations

libRaptorQ can now cache the most used matrices for a quicker reference. The cache can be scaled up or down dynamically.

```
local_cache_size : const uint64_t local_cache  
return: bool
```

Set the local cache size. returns false on error.

```
get_local_cache_size() : return: uint64_t  
get the size of our local cache
```

```
supported_compressions() : return: Compress  
Get the bitmask of all supported compression algorithms.  
currently only Compress::NONE and Compress::LZ4.
```

```
get_compression() : return: Compress  
Get the current compression used.
```

```
set_compression : const Compress compression  
return: bool  
Try to set a compression algorithm. LZ4 is not a mandatory  
dependency, so we might not have it
```

3.3.2 Thread pool

Both the encoder and the decoder can work with multiple blocks at the same time, so it makes sense to have a thread pool for such work.

`set_thread_pool` : **Input:** `const size_t threads`
 `. const uint16_t max_block_concurrency`
 `. const RaptorQ__v1::Work_State exit_type`
return: `bool`

set a pool of a certain number of threads. But the decoder is special: if one decoding fails we can always try again once we have at least one other symbol. The second parameter tells the pool how many times we can try to decode the same block concurrently. 1 is a safe option.

The final argument is used when downsizing the queue: what do we do if all the threads are running but you asked to have less threads?

```
enum class Work_State : uint8_t {  
    KEEP_WORKING  
    ABORT_COMPUTATION  
};
```

3.3.3 The Encoder

The constructor is the following:

```
Encoder (const Rnd_It data_from, const Rnd_It data_to,
         const uint16_t min_subsymbol_size,
         const uint16_t symbol_size,
         const size_t max_sub_block);
```

As we said, the RFC will decide by itself how many blocks each input will have, the actual size of the subsymbol, and the size of the blocks. Although needed for RFC compliance, you probably do not care about the *min_subsymbol_size*, in which case the safe choice is to put it at the same size of the symbol.

The *max_sub_block* the the maximum size of a sub-block in bytes. Again, unless you *really* understand the RFC, you should probably ignore this and put it to a high value. This parameter is used to limit the block size used by the RFC, and is the size (in bytes) of the biggest block you allow the RFC to use.

safe settings for max_sub_block: `sub_symbol_size * block_size`.

You can instantiate an encoder for example by doing:

```
std::vector<uint32_t> input, output;
...
using T_it = typename std::vector<uint32_t>::iterator;
RFC6330::Encoder<T_it, T_it> enc (input.begin(),
                                input.end(),
                                4, 1444, 10000)
```

This will create an Encoder that works on vectors of unsigned 32 bit integers for both input and output, that will create symbols of size 1444 bytes, interleaving your input every 4 bytes, and try to work with a big number of symbols per blocks.

The available methods for the encoder are the following:

`operator bool()` : **return:bool**
False if constructor parameters did not make sense. Else true.

`OTI_Common()` : **return: OTI_Common_Data**, aka *uint64_t*.
Keeps total **file size and symbol size**. You need to send this to the receiver, so that it will be able to properly decode the data.

`OTI_Scheme_Specific_Data()` : **return: OTI_Scheme_Specific_Data**, aka *uint32_t*.
Keeps number of **source blocks, sub blocks, and alignment**. As for the `OTI_Common_Data`, you need to send this to the receiver to be able to properly decode the data.

`compute` : **Input: const Compute flags**
return: std::future<std::pair<Error, uint8_t> >

For C++11 only (automatically disabled on C++98). Start the computation, return a future, which has a pair of Error and the number of blocks ready.

Compute flags:

```
enum class Compute : uint8_t {
    NONE
    // report every time there is more data, but
    // only
    // if that data starts at the beginning.
    PARTIAL_FROM_BEGINNING
    PARTIAL_ANY
    COMPLETE // only when all the data is available
    NO_BACKGROUND // only return after the
    // computation is finished
    NO_POOL // do not use the thread pool
    NO_RETRY // do not retry if we failed, even if
    // there is more data
};
```

precompute_max_memory : **return: size_t**

Each precomputation can take a lot of memory, depending on the configuration, so you might want to limit the number of precomputations run in parallel depending on the memory used. This method returns the amount of memory taken by **ONE** precomputation.

encode : **Input: Fwd_It &output, const Fwd_It end, const uint32_t esi, const uint8_t sbn.**
return: size_t.

Take as input the iterators to the data structure into where we have to save the encoded data, the **Encoding Symbol Id** and the **Source Block Number**. As you are writing in C++, you probably want to use the iterators begin/end, though. Returns the number of written iterators (**NOT** the bytes)

encode : **Input: Fwd_It &output, const Fwd_It end, const uint32_t id.**
return: size_t.

Exactly as before, but the **id** contains both the *source block number* and the *encoding symbol id*

begin() : **return: Block_Iterator<Rnd_It, Fwd_It>**

This returns an iterator to the blocks in which the RFC divided the input data. See later to understand how to use it.

end() : **return: const Block_Iterator<Rnd_It, Fwd_It>**

This returns an iterator to the end of the blocks in which the RFC divided the input data. See later to understand how to use it.

free : **Input: const uint8_t sbn**
return: void
Each block takes some memory, (a bit more than *symbols * symbol_size*), so once you are done sending source and repair symbols for one block, you might want to free the memory of that block.

blocks() : **return: uint8_t** The number of blocks.

block_size() : **Input: const uint8_t sbn**
return: uint32_t
The block size, in bytes. Each block can have different symbols and thus different size. The last block also might have more padding, and with this you get the real size of that block, without the padding.

symbol_size() : **return: uint16_t** The size of a symbol.

symbols : **Input: uint8_t sbn**
return: uint16_t
The number of symbols in a specific block. different blocks can have different symbols.

extended_symbols : **Input: uint8_t sbn**
return: uint16_t
The **real, internal** number of symbols in a specific block. This can be slightly higher than the `symbols()` call, and *you probably do not need this* for anything except if you are trying to know more about the RaptorQ algorithm

max_repair : **Input: const uint8_t sbn)**
return: uint32_t
The maximum amount of repair symbols that you can generate. Something less than 2^{24} , but the exact number depends on the number of symbols in a block

3.3.4 Blocks

With the *begin()/end()* calls you get *Input iterators* to the blocks. A Block has the following type:

```
template <typename Rnd_It, typename Fwd_It>
class Block
```

and exposes the following methods:

begin_source() : **return: Symbol_Iterator**
end_source() : **return: Symbol_Iterator**
begin_repair() : **return: Symbol_Iterator**

end_repair() : **Input:** `const uint32_t max_repair`
return: `Symbol_Iterator`

id() : **return:** `uint8_t`

max_repair() : **return:** `uint32_t`

symbols() : **return:** `uint16_t`

extended_symbols() : **return:** `uint16_t`
As before, this is the amount of internal symbols, that include padding symbols. You probably do not need this unless you are trying to know more about the RaptorQ algorithm

block_size() : **return:** `uint32_t`

As the names explain, you will get an iterator to the symbols in the block. As the number of repair symbols can vary, for now you get two separate begin/ends, so that you can check when you sent the source symbols, and how many repair symbols you send. The other functions are helpers for the details of the block.

3.3.5 Symbols

Finally, through the *Symbol_Iterator Input Iterator* we get the **Symbol** class:

```
template <typename Rnd_It, typename Fwd_It>
class Symbol
```

which exposes the following methods we need to get the symbol data:

operator() : **Input:** `Fwd_It &start, const Fwd_It end`
return: `uint64_t`
takes a forward iterator, and fill it with the symbol data.
returns the number of written iterators.

id() : **return:** `uint32_t`
return the id ($sbn + esi$) of this symbol, that you need to include in every packet you send, before the symbols.

block() : **return:** `uint8_t`
return the *block* of this symbol.

esi() : **return:** `uint32_t`
return the *esi* of this symbol.

3.3.6 The Decoder

The decoder is a bit simpler than the encoder.

There are two constructors for the Decoder:

```
std::vector<uint32_t> input, output;
using T_it = typename std::vector<uint32_t>::iterator;
RaptorQ::Decoder<T_it, T_it> dec (
    const OTI_Common_Data common,
    const OTI_Scheme_Specific_Data scheme)

RaptorQ::Decoder<T_it, T_it> dec (uint64_t size,
    uint16_t symbol_size,
    uint16_t sub_blocks,
    uint8_t blocks)
```

Which should be pretty self-explanatory, once you understand how the encoder works: you can either get the OTI parameters from the Encoder or you can provide the raw values by yourself (only if you know what you are doing)

The remaining methods are:

`begin()` : **return: Block_Iterator<In_It, Fwd_It>**
This returns an iterator to the blocks in which the RFC divided the input data.

`end()` : **return: const Block_Iterator<In_It, Fwd_It>**
This returns an iterator to the end of the blocks in which the RFC divided the input data.

`operator bool()` : **return: bool**
check if the Decoder was initialized properly

`compute` : **Input: const Compute flags**
return: std::future<std::pair<Error, uint8_t> >
Only enabled if you use C++11. Start the computation as soon as enough data is available.
Compute flags:

```
enum class Compute : uint8_t {
NONE
// report every time there is more data, but only
// if that data starts at the beginning.
PARTIAL_FROM_BEGINNING
PARTIAL_ANY
COMPLETE // only when all the data is available
NO_BACKGROUND // only return after the computation
                is finished
NO_POOL // do not use the thread pool
NO_RETRY // do not retry if we failed, even if there
                is more data
};
```

end_of_input : **Input:** **const Fill_With_Zeros fill**
. **const uint8_t block**
return: **std::vector<bool>**
Tell the decoder that we know that there will be no more data for this block. The first parameter is an **enum**, either **Fill_With_Zeros::YES** or **Fill_With_Zeros::NO** If you choose **YES** then the repair symbols will be discarded, the missing symbols will be filled with zeros and for output you will get the bitmask for which byte was received (true) or was filled with zeros (false).
If you choose **Fill_With_Zeros::NO** you will get *an empty vector*

end_of_input : **Input:** **const Fill_With_Zeros fill**
return: **std::vector<bool>**
Same as before, but we know that there will be no more inputs for any block.

decode_symbol : **Input:** **Fwd_It &start**
. **const Fwd_It end**
. **const uint16_t esi**
. **const uint8_t sbn** **return:** **uint64_t**
Decode a single symbol. you need to provide the block number and symbol identifier.

decode_bytes : **Input:** **Fwd_It &start**
. **const Fwd_It end**
. **const uint8_t skip**
return: **uint64_t**
Decode everything. returns the number of bytes written. You can start writing the output **skip** bytes after the first iterator in case the parameters you have used do not align (set to 0 for everyone else).

decode_bytes : **Input:** **Fwd_It &start**
. **const Fwd_It end**
. **const uint8_t skip**
. **const uint8_t sbn**
return: **uint64_t**
Same as before, but you can specify a **Single Block Number** instead of the whole output

`decode_aligned` : **Input:** `Fwd_It &start`
 . `const Fwd_It end`
 . `const uint8_t skip`
return: `aligned_res`
 Same as `decode_bytes`, but you can use this if for some weird reason the data does not align with the type of the output

```
struct aligned_res
{
  uint64_t written; // iterators written
  uint8_t offset; // bytes written in the last
                  // iterator if it was not full
};
```

`decode_block_aligned` : **Input:** `Fwd_It &start`
 . `const Fwd_It end`
 . `const uint8_t skip`
 . `const uint8_t sbn`
return: `aligned_res`
 Same as before, but now you can choose the block

`add_symbol` : **Input:** `In_It &start`
 . `const In_It end`
 . `const uint32_t esi`
 . `const uint8_t sbn`
return: `RFC6330::Error`
 Add one symbol, while explicitly specifying the symbol id and the block id.

`add_symbol` : **Input:** `In_It &start`
 . `const In_It end`
 . `const uint32_t id`
return: `RFC6330::Error`
 Same as before, but extract the block id and the symbol id from the `id` parameter

`blocks_ready()` : **return:** `uint8_t`
 return the number of the blocks that are ready for output

`is_ready()` : **return:** `bool`
 Whether everything has been decoded or not.

`is_block_ready` : **Input:** `const uint8_t block`
return: `bool`
 Test if the specified block has been decoded or not

`free` : **Input:** `const uint8_t sbn`
return: `void`
 You might have stopped using a block, but the memory is still there. Free it.

`bytes()` : **return: uint64_t**
The total bytes of the output

`blocks()` : **return: uint8_t**
The number of blocks.

`block_size()` : **Input: const uint8_t sbn**
return: uint32_t
The block size, in bytes. Each block can have different symbols and thus different size.

`symbol_size()` : **return: uint16_t** The size of a symbol.

`symbols` : **Input: uint8_t sbn**
return: uint16_t
The number of symbols in a specific block. different blocks can have different symbols.

3.4 C interface

The C interface looks a lot like the C++ one.

You need to include the **RaptorQ/RFC6330.h** header, and link the *libRaptorQ* library.

Static linking If you are working with the static version of libRaptorQ remember to link the C++ standard library used when compiling the library (*libstdc++* for gcc or maybe *libc++* for clang), your threading library (usually *libpthread*), and the C math library (*libm*).

API versioning To better support future changes in the API and to help you with future changes, there are only 2 API calls in this library:

RFC6330_api : **Input: uint32_t version**
return: struct RFC6330_base_api*

This will give you a pointer to a struct containing pointers to the calls you will use. The struct will change depending on which API you request (currently only number 1 is supported). Since this pointer works for every api, you have to cast it to the struct you expect, or you will not see the pointers to the function calls. Then you can use that pointer to access of the functions, like:

```
struct RFC6330_v1 *rfc = (struct RFC6330_v1*)
                        RFC6330_api (1);
rfc->set_compression (RQ_COMPRESS_LZ4);
```

RFC6330_free_api : **Input: struct RFC6330_base_api **api**
return: void

As the name implies, this frees the given structure.

Again, we only support API version 1 at the moment.

3.4.1 Cache settings

As for the C++11 interface, you can get and set the local memory cache.

supported_compressions : **return: RFC6330_Compress**
return the bitmask of all supported compressions. currently either `RQ_COMPRESS_NONE` or `RQ_COMPRESS_LZ4`

get_compression : **return: RFC6330_Compress**
Which compression are we currently using?

set_compression : **Input: const RFC6330_Compress compression**
return: bool
Set a compression method or another.

local_cache_size : **Input:** const size_t bytes
return: size_t
Set the maximum amount of cache usable. Default: 0

get_local_cache_size : **return:** size_t
get the amount of cache configured

3.4.2 C Constructors

The first thing you need is to allocate the encoder or the decoder.

The C interface is just a wrapper around the C++ code, so you still have to specify the same things as before. A quick glance at the constructors should give you all the information you need:

```
typedef enum { RQ_NONE, RQ_ENC_8, RQ_ENC_16,
              RQ_ENC_32, RQ_ENC_64, RQ_DEC_8,
              RQ_DEC_16, RQ_DEC_32, RQ_DEC_64}
              RFC6330_type;
```

Encoder : **Input:**const RFC6330_type type
. void *data_from
. const uint64_t size
. const uint16_t min_subsymbol_size
. const uint16_t symbol_size
. const size_t max_sub_block
return: RFC6330_ptr*
Allocate an encoder. The encoder will work with uint8_t, uint16_t, uint32_t or uint64_t depending on which RFC6330_type you choose. The best bet is usually to work with uint8_t

Decoder : **Input:**const RFC6330_type type
. const RFC6330_OTI_Common_Data common
. const RFC6330_OTI_Scheme_Specific_Data scheme
return: RFC6330_ptr*
The type parameter works the same as in the encoder. The other two are parameters you get from the encoder.

Decoder_raw : **Input:** RFC6330_type type
. const uint64_t size
. const uint16_t symbol_size
. const uint16_t sub_blocks
. const uint8_t blocks
. const uint8_t alignment
return: struct RFC6330_ptr*
Same as before, but here you can specify the parameters by hand.

initialized : **Input:** const struct RFC6330_ptr *ptr
return: bool
check if the encoder or decoder was initialized properly.

3.4.3 Common functions for (de/en)coding

These functions are used by both the encoder and the decoder, and will be helpful in tracking how much memory you will need to allocate, or in general in managing the encoder and decoder.

The **ptr** must be a valid encoder or decoder. The names for now are self-explanatory. Blocks can have different symbols, so the size of a block and the number of symbols in a block depend on which block we are talking about.

symbols, blocks, memory

set_thread_pool : **Input:** **const size_t threads**
 . **const uint16_t max_block_concurrency**
 . **const RFC6330_Work exit_type**
return: **bool**

Set the size of the thread pool for concurrent work. Since the decoder can fail, but also retry again if there is more data available, you can also specify how many times libRaptorQ will try to decode the same block concurrently. 1 is a safe option.

The last parameter, **RFC6330_Work** specify what to do when you are downsizing the thread pool and some threads are still working. The possible values are **RQ_WORK_KEEP_WORKING** and **RQ_WORK_ABORT_COMPUTATION**.

blocks : **Input:** **const struct RFC6330_ptr *ptr**
return: **uint8_t**
The number of blocks in this RFC6330 instance

symbols : **Input:** **const struct RFC6330_ptr *ptr**
 . **const uint8_t sbn**
return: **uint16_t**
The number of symbols in the specified block

extended_symbols : **Input:** **const struct RFC6330_ptr *ptr**
 . **const uint8_t sbn**
return: **uint16_t**
The the **real, internal** number of symbols in the specified block. You probably *do not want this* unless you are trying to understand more about the RaptorQ algorithm or you are benchmarking

symbol_size : **Input:** **const struct RFC6330_ptr *ptr**
return: **size_t**
The size of each symbol

free_block : **Input:** **struct RFC6330_ptr **ptr**
 . **const uint8_t sbn**
return: **void**
Free the allocated space for a single block

free : **Input:** struct RFC6330_ptr **ptr
return: void

Free the specified decoder or decoder and all its blocks.

Computation

compute : **Input:** const struct RFC6330_ptr *ptr
. const RFC6330_Compute flags
return: RFC6330_future*

Start the computation as described in the flags. Returns a wrapper to the C++11 futures. You can wait or poll the state of the future in different ways (see right below)

```
typedef enum {
RQ_COMPUTE_NONE
// report every time there is more data, but only
// if that data starts at the beginning.
RQ_COMPUTE_PARTIAL_FROM_BEGINNING
RQ_COMPUTE_PARTIAL_ANY
RQ_COMPUTE_COMPLETE // only when all the data
// is available
RQ_COMPUTE_NO_BACKGROUND // only return after
// the computation is finished
RQ_COMPUTE_NO_POOL // do not use the thread pool
RQ_COMPUTE_NO_RETRY // do not retry if we failed,
// even if there is more data
} RFC6330_Compute;
```

future_state : **Input:** const struct RFC6330_future *f
return: RFC6330_Error

poll the state of the future. possible results:

```
RQ_ERR_WRONG_INPUT // null pointer?
RQ_ERR_NOT_NEEDED // you have already run
// future_get() on this future
RQ_ERR_WORKING // still working on it
RQ_ERR_NONE // ready to run future_get()
```

`future_wait_for` : **Input:** `const struct RFC6330_future *f`
 . `const uint64_t time`
 . `const RFC6330_Unit_Time unit`
return: `RFC6330_Error`

Wait for a future to become ready for X amount of time. the unit is specified in the last parameter:

```
typedef enum {  
    RQ_TIME_NANOSEC,  
    RQ_TIME_MICROSEC,  
    RQ_TIME_MILLISEC,  
    RQ_TIME_SEC,  
    RQ_TIME_MIN,  
    RQ_TIME_HOUR  
} RFC6330_Unit_Time;
```

`future_wait` : **Input:** `const struct RFC6330_future *f`
return: `void`

Same as before, but wait indefinitely

`future_get` : **Input:** `struct RFC6330_future *future`
return: `struct RFC6330_Result`

Returns the state of the future. If the future is not ready it will wait indefinitely.

```
struct RAPTORQ_API RFC6330_Result {  
    RFC6330_Error error;  
    uint8_t sbn;  
};
```

The sbn changes meaning depending to be consistent with the flags you passed to the `compute()` call.

`future_free` : **Input:** `struct RFC6330_future **f`
return: `void`

Free the memory allocated for the future.

3.4.4 Encoding

Now we will look at the calls specific to the encoder (that is: calls that need an encoder as the first parameter, and don't make sense in a decoder)

OTI Data

First, we need to tell the receiver all the parameters that the encoder is using, and for that two functions are provided:

```
typedef uint64_t RaptorQ_OTI_Common_Data;  
typedef uint32_t RaptorQ_OTI_Scheme_Specific_Data;
```

OTI_Common : **Input:** struct RaptorQ_ptr *enc
return: RFC6330_OTI_Common_Data
Get the OTI_Common

OTI_Scheme_Specific : **Input:** struct RaptorQ_ptr *enc
return: RFC6330_OTI_Scheme_Specific_Data
Get the OTI_Scheme_Specific

Encoding

max_repair : **Input:** const struct RFC6330_ptr *enc
. const uint8_t sbn
return: uint32_t
return the max number of repair packets that block can have. This is around 2^{24} , so you probably always want to be much, much lower than this. The total number of symbols (source+repair) is limited in the RFC, and since the bigger blocks have more source symbols, the max number of repair symbols has to be lower in bigger blocks. Again, this is always in the ballpark of 2^{24} , you probably don't want to use this many repair symbols.

precompute_max_memory : **Input:** const struct RFC6330_ptr *enc
return: size_t
Get an estimate of how much memory (bytes) will cost the encoding of a single block. Not 100% accurate, but close enough.

encode_id : **Input:** const struct RFC6330_ptr *enc
. void **data
. const size_t size
. const uint32_t id
return: size_t
Get one encoded symbol, and store it in **data, which holds a pointer to "size" elements of type uint8_t, uint16_t, uint32_t or uint64_t depending on how the encoder was initialized. See RFC6330_type in the encoder initialization. Returns the number of elements written.

If you request repair symbols before the computation is finished, the call will block until the data is available. This call works for both source and repair symbols, just increase the id until you finish the source symbols

```
encode : Input: const struct RFC6330_ptr *enc  
        void **data  
        const size_t size  
        const uint32_t esi  
        const uint8_t sbn  
return: size_t
```

Same as before, but you can specify the symbol and block explicitly.

```
id : Input: const uint32_t esi  
     const uint8_t sbn  
return: uint32_t
```

combine the esi and the block number to form an id.

3.4.5 Decoding

As for the encoder, there are functions which are specific for a decoder. Let's look at them

General status calls

bytes : **Input: const struct RFC6330_ptr *dec**
return: uint64_t
Return the number of bytes that this decoder should output on complete, successful decoding.

blocks_ready : **Input: const struct RFC6330_ptr *dec**
return: uint8_t
Return the number of blocks ready to be returned. Useful if you don't want to wait for the whole data.
The number of blocks reported can mean either the number of sequential blocks ready, or the number of blocks regardless of order, depending on how the compute function was called (see "Computation" on 3.4.3)

is_ready : **Input: const struct RFC6330_ptr *dec**
return: bool
Test if the whole encoder is ready and all data decoded.

is_block_ready : **Input: const struct RFC6330_ptr *dec**
 const uint8_t block
return: bool
Test if a single block is ready and its data decoded.

Decoding calls

end_of_input : **Input: const struct RFC6330_ptr *dec**
 const RFC6330_Fill_With_Zeros fill
return: struct RFC6330_Byte_Tracker
Tell the decoder that there will be no more input for any block. The second parameter is an enum, it can be either **RQ_NO_FILL** or **RQ_FILL_WITH_ZEROS**.
If you use **RQ_NO_FILL** the structure returned will be empty, and the decoder will return error if there are not enough repair symbols.
If you use **RQ_FILL_WITH_ZEROS** the decoder will immediately stop decoding and all missing symbols will be filled with zeros. The returned structure will contain a bitmask so you know if each byte comes from the network (true) or it is one of the filled ones (false)

end_of_block_input : **Input: const struct RFC6330_ptr *dec**
 const RFC6330_Fill_With_Zeros fill
 const uint8_t block
return: struct RFC6330_Byte_Tracker
Same as before, but for a single block instead of the whole data.

add_symbol_id : **Input:** const struct RFC6330_ptr *dec
. const void **data
. const uint32_t size
. const uint32_t id
return: RFC6330_Error
Add a symbol to the decoder. The void **data will be cast to the correct type as specified in the decoder construction (RQ_DEC_8, RQ_DEC_16 etc.). Afterwards data will point **after** the symbol.

add_symbol_id : **Input:** const struct RFC6330_ptr *dec
. const void **data
. const uint32_t size
. const uint32_t esi
. const uint8_t sbn
return: RFC6330_Error
Same as with add_symbol_id, but now you can specify the esi and the block number manually

decode_aligned : **Input:** const struct RFC6330_ptr *dec
. const void **data
. const uint32_t size
. const uint8_t skip
return: RFC6330_Dec_Result
Decode everything as asked by the compute() call (see "Computation" on 3.4.3).
The void **data will be cast to the correct type as specified in the decoder construction (RQ_DEC_8, RQ_DEC_16 etc.). Afterwards data will point **after** the decoded data. You can skip "skip" bytes at the beginning of **data if the pointer is not aligned with your data.

```
struct RAPTORQ_API RFC6330_Dec_Result {
    uint64_t written;
    uint8_t offset;
};
```

The result is the number of written **elements** (size depending on RQ_DEC_8, RQ_DEC_16 etc).

The **offset** parameter is the amount of bytes written in the last element if the alignment used in the encoder does not fit the alignment used in the decoder. Usually you keep those synchronized, so it should be zero (element fully written)

`decode_block_aligned` : **Input:** `const struct RFC6330_ptr *dec`
 . `const void **data`
 . `const uint32_t size`
 . `const uint8_t skip`
 . `const uint8_t block`
return: `RFC6330_Dec_Result`
 Same as before, but now you can decode a single block.

`decode_symbol` : **Input:** `const struct RFC6330_ptr *dec`
 . `const void **data`
 . `const uint32_t size`
 . `const uint8_t skip`
 . `const uint32_t esi`
 . `const uint8_t sbn`
return: `uint64_t`
 As you can expect, decode a single symbol. The `void **data` will be cast to the correct type as specified in the decoder construction (`RQ_DEC_8`, `RQ_DEC_16` etc..). Afterwards data will point **after** the decoded data.
 Returns the number of bytes written

`decode_bytes` : **Input:** `const struct RFC6330_ptr *dec`
 . `const void **data`
 . `const uint32_t size`
 . `const uint8_t skip`
return: `uint64_t`
 Decode everything as asked by the `compute()` call (see "Computation" on 3.4.3).
 The `void **data` will be cast to the correct type as specified in the decoder construction (`RQ_DEC_8`, `RQ_DEC_16` etc..). Afterwards data will point **after** the decoded data. You can skip "skip" bytes at the beginning of `**data` if the pointer is not aligned with your data.
 Returns the number of written bytes.

`decode_block_bytes` : **Input:** `const struct RFC6330_ptr *dec`
 . `const void **data`
 . `const uint32_t size`
 . `const uint8_t skip`
 . `const uint8_t block`
return: `uint64_t`
 Same as before, but now decode a single block

4 GNU Free Documentation License

GNU Free Documentation License
Version 1.3, 3 November 2008

Copyright (C) 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.
<<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the

Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word

processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

The "publisher" means any person or entity that distributes copies of the Document to the public.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover

Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified

Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains

nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual

copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve

its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this

License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

"Massive Multiauthor Collaboration Site" (or "MMC Site") means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A "Massive Multiauthor Collaboration" (or "MMC") contained in the site means any set of copyrightable works thus published on the MMC site.

"CC-BY-SA" means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

"Incorporate" means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is "eligible for relicensing" if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (c) YEAR YOUR NAME.  
Permission is granted to copy, distribute and/or modify this document  
under the terms of the GNU Free Documentation License, Version 1.3  
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Index

Blocks

C++, 14

Cache

C, 20

C++, 10

cMake, 4

Contacts, 3

Decoder

C, 27

C++, 16

dependencies, 4

Encoder

C, 25

C++, 12

GPG, 3

Install, 5

Interface

C, 20

C++, 8

Interleaving, 8

Iterators, 10

Sequencing, 7

Sizes, 8

Static Linking, 20

Symbols

C++, 15

Theory, 7