Convert::Binary::C

Binary Data Conversion using C Types

Version 0.56

Marcus Holland-Moritz

mhx@cpan.org

September 18, 2004

Contents

Convert::Binary::C

Binary Data Conversion using C Types

1.1 Synopsis

1.1.1 Simple

```
use Convert::Binary::C;

#---------------------------------------------
# Create a new object and parse embedded code
#---------------------------------------------
my $c = Convert::Binary::C->new->parse( <<ENDC );

enum Month { JAN, FEB, MAR, APR, MAY, JUN,
             JUL, AUG, SEP, OCT, NOV, DEC };

struct Date {
  int        year;
  enum Month month;
  int        day;
};

ENDC

#---------------------------------------------
# Pack Perl data structure into a binary string
#---------------------------------------------
my $date = { year => 2002, month => 'DEC', day => 24 };

my $packed = $c->pack( 'Date', $date );
```

1.1.2 Advanced

```
use Convert::Binary::C;
use Data::Dumper;

#---------------------
# Create a new object
#---------------------
my $c = new Convert::Binary::C ByteOrder => 'BigEndian';
```

```
#------------------------------------------------------
# Add include paths and global preprocessor defines
#------------------------------------------------------
$c->Include( '/usr/lib/gcc-lib/i686-pc-linux-gnu/3.3.4/include',
             '/usr/include' )
  ->Define( qw( __USE_POSIX __USE_ISOC99=1 ) );


#---------------------------------
# Parse the 'time.h' header file
#---------------------------------
$c->parse_file( 'time.h' );


#--------------------------------------
# See which files the object depends on
#--------------------------------------
print Dumper( [$c->dependencies] );


#-----------------------------------------------------------
# See if struct timespec is defined and dump its definition
#-----------------------------------------------------------
if( $c->def( 'struct timespec' ) ) {
  print Dumper( $c->struct( 'timespec' ) );
}


#------------------------------
# Create some binary dummy data
#------------------------------
my $data = "binaryteststring";


#----------------------------------------------------------
# Unpack $data according to 'struct timespec' definition
#----------------------------------------------------------
if( length($data) >= $c->sizeof( 'timespec' ) ) {
  my $perl = $c->unpack( 'timespec', $data );
  print Dumper( $perl );
}


#----------------------------------------------------------
# See which member lies at offset 5 of 'struct timespec'
#----------------------------------------------------------
my $member = $c->member( 'timespec', 5 );
print "member( 'timespec', 5 ) = '$member'\n";
```

1.2 Description

Convert::Binary::C is a preprocessor and parser for C type definitions. It is highly configurable and should support arbitrarily complex data structures. Its object-oriented interface has `pack` and `unpack` methods that act as replacements for Perl's `pack` and `unpack` and allow to use the C types instead of a string representation of the data structure for conversion of binary data from and to Perl's complex data structures.

Actually, what Convert::Binary::C does is not very different from what a C compiler does, just that it doesn't compile the source code into an object file or executable, but only parses the code and allows Perl to use the enumerations, structs, unions and typedefs that have been defined within your C source for binary data conversion, similar to Perl's `pack` and `unpack`.

Beyond that, the module offers a lot of convenience methods to retrieve information about the C types that have been parsed.

1.2.1 Background and History

In late 2000 I wrote a real-time debugging interface for an embedded medical device that allowed me to send out data from that device over its integrated Ethernet adapter. The interface was `printf()`-like, so you could easily send out strings or numbers. But you could also send out what I called *arbitrary data*, which was intended for arbitrary blocks of the device's memory.

Another part of this real-time debugger was a Perl application running on my workstation that gathered all the messages that were sent out from the embedded device. It printed all the strings and numbers, and hex-dumped the arbitrary data. However, manually parsing a couple of 300 byte hex-dumps of a complex C structure is not only frustrating, but also error-prone and time consuming.

Using `unpack` to retrieve the contents of a C structure works fine for small structures and if you don't have to deal with struct member alignment. But otherwise, maintaining such code can be as awful as deciphering hex-dumps.

As I didn't find anything to solve my problem on the CPAN, I wrote a little module that translated simple C structs into `unpack` strings. It worked, but it was slow. And since it couldn't deal with struct member alignment, I soon found myself adding padding bytes everywhere. So again, I had to maintain two sources, and changing one of them forced me to touch the other one.

All in all, this little module seemed to make my task a bit easier, but it was far from being what I was thinking of:

- A module that could directly use the source I've been coding for the embedded device without any modifications.

- A module that could be configured to match the properties of the different compilers and target platforms I was using.

- A module that was fast enough to decode a great amount of binary data even on my slow workstation.

I didn't know how to accomplish these tasks until I read something about XS. At least, it seemed as if it could solve my performance problems. However, writing a C parser in C isn't easier than it is in Perl. But writing a C preprocessor from scratch is even worse.

Fortunately enough, after a few weeks of searching I found both, a lean, open-source C preprocessor library, and a reusable YACC grammar for ANSI-C. That was the beginning of the development of Convert::Binary::C in late 2001.

Now, I'm successfully using the module in my embedded environment since long before it appeared on CPAN. From my point of view, it is exactly what I had in mind. It's fast, flexible, easy to use and portable. It doesn't require external programs or other Perl modules.

1.2.2   About this document

This document describes how to use Convert::Binary::C. A lot of different features are presented, and the example code sometimes uses Perl's more advanced language elements. If your experience with Perl is rather limited, you should know how to use Perl's very good documentation system.

To look up one of the manpages, use the `perldoc` command. For example,

```
perldoc perl
```

will show you Perl's main manpage. To look up a specific Perl function, use `perldoc -f`:

```
perldoc -f map
```

gives you more information about the `map` function. You can also search the FAQ using `perldoc -q`:

```
perldoc -q array
```

will give you everything you ever wanted to know about Perl arrays. But now, let's go on with some real stuff!

1.2.3   Why use Convert::Binary::C?

Say you want to pack (or unpack) data according to the following C structure:

```
struct foo {
  char ary[3];
  unsigned short baz;
  int bar;
};
```

You could of course use Perl's `pack` and `unpack` functions:

```
@ary = (1, 2, 3);
$baz = 40000;
$bar = -4711;
$binary = pack 'c3 S i', @ary, $baz, $bar;
```

But this implies that the struct members are byte aligned. If they were long aligned (which is the default for most compilers), you'd have to write

```
$binary = pack 'c3 x S x2 i', @ary, $baz, $bar;
```

which doesn't really increase readability.

Now imagine that you need to pack the data for a completely different architecture with different byte order. You would look into the `pack` manpage again and perhaps come up with this:

```
$binary = pack 'c3 x n x2 N', @ary, $baz, $bar;
```

However, if you try to unpack `$foo` again, your signed values have turned into unsigned ones.

All this can still be managed with Perl. But imagine your structures get more complex? Imagine you need to support different platforms? Imagine you need to make changes to the structures? You'll not only have to change the C source but also dozens of `pack` strings in your Perl code. This is no fun. And Perl should be fun.

Now, wouldn't it be great if you could just read in the C source you've already written and use all the types defined there for packing and unpacking? That's what Convert::Binary::C does.

1.2.4   Creating a Convert::Binary::C object

To use Convert::Binary::C just say

```
use Convert::Binary::C;
```

to load the module. Its interface is completely object oriented, so it doesn't export any functions.

Next, you need to create a new Convert::Binary::C object. This can be done by either

```
$c = Convert::Binary::C->new;
```

or

```
$c = new Convert::Binary::C;
```

You can optionally pass configuration options to the constructor as described in the next section.

1.2.5   Configuring the object

To configure a Convert::Binary::C object, you can either call the `configure` method or directly pass the configuration options to the constructor. If you want to change byte order and alignment, you can use

```
$c->configure( ByteOrder => 'LittleEndian',
               Alignment => 2 );
```

or you can change the construction code to

```
$c = new Convert::Binary::C ByteOrder => 'LittleEndian',
                            Alignment => 2;
```

Either way, the object will now know that it should use little endian (Intel) byte order and 2-byte struct member alignment for packing and unpacking.

Alternatively, you can use the option names as names of methods to configure the object, like:

```
$c->ByteOrder( 'LittleEndian' );
```

You can also retrieve information about the current configuration of a Convert::Binary::C object. For details, see the section about the `configure` method.

1.2.6   Parsing C code

Convert::Binary::C allows two ways of parsing C source. Either by parsing external C header or C source files:

```
$c->parse_file( 'header.h' );
```

Or by parsing C code embedded in your script:

```
$c->parse( <<'CCODE' );
struct foo {
  char ary[3];
  unsigned short baz;
  int bar;
};
CCODE
```

Now the object `$c` will know everything about `struct foo`. The example above uses a so-called here-document. It allows to easily embed multi-line strings in your code. You can find more about here-documents in the *perldata* manpage or the *perlop* manpage.

Since the `parse` and `parse_file` methods throw an exception when a parse error occurs, you usually want to catch these in an `eval` block:

```
eval { $c->parse_file('header.h') };
if( $@ ) {
  # do something appropriate
}
```

Perl's special `$@` variable will contain an empty string (which evaluates to a false value in boolean context) on success or an error string on failure.

As another feature, `parse` and `parse_file` return a reference to their object on success, just like `configure` does when you're configuring the object. This will allow you to write constructs like this:

```
my $c = eval {
  Convert::Binary::C->new( Include => ['/usr/include'] )
                    ->parse_file( 'header.h' )
};
if( $@ ) {
  # do something appropriate
}
```

1.2.7   Packing and unpacking

Convert::Binary::C has two methods, `pack` and `unpack`, that act similar to the functions of same denominator in Perl. To perform the packing described in the example above, you could write:

```
$data = {
  ary => [1, 2, 3],
  baz => 40000,
  bar => -4711,
};
$binary = $c->pack( 'foo', $data );
```

Unpacking will work exactly the same way, just that the unpack method will take a byte string as its input and will return a reference to a (possibly very complex) Perl data structure.

```
$binary = from_memory();
$data = $c->unpack( 'foo', $binary );
```

You can now easily access all of the values:

```
print "foo.ary[1] = $data->{ary}[1]\n";
```

Or you can even more conveniently use the Data::Dumper module:

```
use Data::Dumper;
print Dumper( $data );
```

The output would look something like this:

```
$VAR1 = {
  'bar' => -271,
  'baz' => 5000,
  'ary' => [
    42,
    48,
    100
  ]
};
```

1.2.8   Preprocessor configuration

Convert::Binary::C uses Thomas Pornin's `ucpp` as an internal C preprocessor. It is compliant to ISO-C99, so you don't have to worry about using even weird preprocessor constructs in your code.

If your C source contains includes or depends upon preprocessor defines, you may need to configure the internal preprocessor. Use the `Include` and `Define` configuration options for that:

```
$c->configure( Include => ['/usr/include',
                           '/home/mhx/include'],
              Define  => [qw( NDEBUG FOO=42 )] );
```

If your code uses system includes, it is most likely that you will need to define the symbols that are usually defined by the compiler.

On some operating systems, the system includes require the preprocessor to predefine a certain set of assertions. Assertions are supported by `ucpp`, and you can define them either in the source code using `#assert` or as a property of the Convert::Binary::C object using `Assert`:

```
$c->configure( Assert => ['predicate(answer)'] );
```

1.2.9 Supported pragma directives

Convert::Binary::C supports the `pack` pragma to locally override struct member alignment. The supported syntax is as follows:

#pragma pack(ALIGN)

     Sets the new alignment to ALIGN.

#pragma pack

     Resets the alignment to its original value.

#pragma pack(push, ALIGN)

     Saves the current alignment on a stack and sets the new alignment to ALIGN.

#pragma pack(pop)

     Restores the alignment to the last value saved on the stack.

```
/*  Example assumes sizeof( short ) == 2, sizeof( long ) == 4.  */

#pragma pack(1)

struct nopad {
  char a;               /* no padding bytes between 'a' and 'b' */
  long b;
};

#pragma pack             /* reset to "native" alignment         */

#pragma pack( push, 2 )

struct pad {
  char    a;            /* one padding byte between 'a' and 'b' */
  long    b;

#pragma pack( push, 1 )

  struct {
    char  c;            /* no padding between 'c' and 'd'       */
    short d;
  }       e;            /* sizeof( e ) == 3                     */

#pragma pack( pop );    /* back to pack( 2 )                    */

  long    f;            /* one padding byte between 'e' and 'f' */
};

#pragma pack( pop );    /* back to "native"                     */
```

The `pack` pragma as it is currently implemented only affects the *maximum* struct member alignment. There are compilers that also allow to specify the *minimum* struct member alignment. This is not supported by Convert::Binary::C.

1.2.10 Automatic configuration using `ccconfig`

As there are over 20 different configuration options, setting all of them correctly can be a lengthy and tedious task.

The `ccconfig` script, which is bundled with this module, aims at automatically determining the correct compiler configuration by testing the compiler executable. It works for both, native and cross compilers.

1.3 Understanding Types

This section covers one of the fundamental features of Convert::Binary::C. It's how *type expressions*, referred to as TYPEs in the method reference, are handled by the module.

Many of the methods, namely `pack`, `unpack`, `sizeof`, `typeof`, `member` and `offsetof`, are passed a TYPE to operate on as their first argument.

1.3.1 Standard Types

These are trivial. Standard types are simply enum names, struct names, union names, or typedefs. Almost every method that wants a TYPE will accept a standard type.

For enums, structs and unions, the prefixes `enum`, `struct` and `union` are optional. However, if a typedef with the same name exists, like in

```
struct foo {
  int bar;
};

typedef int foo;
```

you will have to use the prefix to distinguish between the struct and the typedef. Otherwise, a typedef is always given preference.

1.3.2 Basic Types

Basic types, or atomic types, are `int` or `char`, for example. It's possible to use these basic types without having parsed any code. You can simply do

```
$c = new Convert::Binary::C;
$size = $c->sizeof( 'unsigned long' );
$data = $c->pack( 'short int', 42 );
```

Even though the above works fine, it is not possible to define more complex types on the fly, so

```
$size = $c->sizeof( 'struct { int a, b; }' );
```

will result in an error.

Basic types are not supported by all methods. For example, it makes no sense to use `member` or `offsetof` on a basic type. Using `typeof` isn't very useful, but supported.

1.3.3 Member Expressions

This is by far the most complex part, depending on the complexity of your data structures. Any standard type that defines a compound or an array may be followed by a member expression to select only a certain part of the data type. Say you have parsed the following C code:

```
struct foo {
  long type;
  struct {
    short x, y;
  } array[20];
};

typedef struct foo matrix[8][8];
```

You may want to know the size of the `array` member of `struct foo`. This is quite easy:

```
print $c->sizeof( 'foo.array' ), " bytes";
```

will print

```
80 bytes
```

depending of course on the `ShortSize` you configured.

If you wanted to unpack only a single column of `matrix`, that's easy as well (and of course it doesn't matter which index you use):

```
$column = $c->unpack( 'matrix[2]', $data );
```

Member expressions can be arbitrarily complex:

```
$type = $c->typeof( 'matrix[2][3].array[7].y' );
print "the type is $type";
```

will, for example, print

```
the type is short
```

Member expressions are also used as the second argument to `offsetof`.

1.3.4 Offsets

Members returned by the `member` method have an optional offset suffix to indicate that the given offset doesn't point to the start of that member. For example,

```
$member = $c->member( 'matrix', 1431 );
print $member;
```

will print

```
[2][1].type+3
```

If you would use this as a member expression, like in

```
$size = $c->sizeof( "matrix $member" );
```

the offset suffix will simply be ignored. Actually, it will be ignored for all methods if it's used in the first argument.

When used in the second argument to offsetof, it will usually do what you mean, i. e. the offset suffix, if present, will be considered when determining the offset. This behaviour ensures that

```
$member = $c->member( 'foo', 43 );
$offset = $c->offsetof( 'foo', $member );
print "'$member' is located at offset $offset of struct foo";
```

will always correctly set `$offset`:

```
'.array[9].y+1' is located at offset 43 of struct foo
```

If this is not what you mean, e. g. because you want to know the offset where the member returned by member starts, you just have to remove the suffix:

```
$member =~ s/\+\d+$//;
$offset = $c->offsetof( 'foo', $member );
print "'$member' starts at offset $offset of struct foo";
```

This would then print:

```
'.array[9].y' starts at offset 42 of struct foo
```

1.4   Using Hooks

Hooks are a new feature and in this early state are considered experimental. That means that both interface and behaviour may be subject to incompatible changes. Nevertheless hooks are quite interesting and can be extremely useful, so read on!

Using hooks, you can easily override the way pack and unpack handle data. If you define hooks for a certain data type, each time this data type is processed the corresponding hook will be called to allow you to modify the data.

1.4.1   Basic hooks

Here's an example. Let's assume the following C code has been parsed:

```
typedef unsigned long u_32;
typedef u_32 ProtoId;
typedef ProtoId MyProtoId;

struct MsgHeader {
  MyProtoId id;
  u_32      len;
};
```

```
struct String {
  u_32 len;
  char buf[];
};
```

You could now use the types above and, for example, unpack binary data representing a `MsgHeader` like this:

```
$msg_header = $c->unpack('MsgHeader', $data);
```

This would give you:

```
$msg_header = {
  'len' => 13,
  'id' => 42
};
```

Instead of dealing with `ProtoId`'s as integers, you would rather like to have them as clear text. You could provide subroutines to convert between clear text and integers:

```
%proto = (
  CATS      =>     1,
  DOGS      =>    42,
  HEDGEHOGS => 4711,
);

%rproto = reverse %proto;

sub ProtoId_unpack {
  $rproto{$_[0]} || 'unknown protocol'
}

sub ProtoId_pack {
  $proto{$_[0]} or die 'unknown protocol'
}
```

The hooks feature allows you to register these subroutines using the add_hooks method:

```
$c->add_hooks(ProtoId => { pack   => \&ProtoId_pack,
                           unpack => \&ProtoId_unpack });
```

Doing exactly the same unpack on `MsgHeader` again would now return:

```
$msg_header = {
  'len' => 13,
  'id' => 'DOGS'
};
```

Actually, if you don't need the reverse operation, you don't even have to register a `pack` hook. Or, even better, you can have a more intelligent `unpack` hook that creates a dual-typed variable:

```
use Scalar::Util qw(dualvar);
```

```
sub ProtoId_unpack2 {
  dualvar $_[0], $rproto{$_[0]} || 'unknown protocol'
}


$c->add_hooks(ProtoId => { unpack => \&ProtoId_unpack2 });


$msg_header = $c->unpack('MsgHeader', $data);
```

Just as before, this would print

```
$msg_header = {
  'len' => 13,
  'id' => 'DOGS'
};
```

but without requiring a `pack` hook for packing, at least as long as you keep the variable dual-typed.

Hooks are called with exactly one argument, which is the data that should be processed. They are called in scalar context and expected to return the processed data.

To get rid of registered hooks, you can use either the `delete_hooks` method

```
$c->delete_hooks('ProtoId');
```

or `delete_all_hooks`:

```
$c->delete_all_hooks;
```

Of course, hooks are not restricted to handling integer values. You could just as well add hooks for the `String` struct from the code above. A useful example would be to have these hooks:

```
sub string_unpack {
  my $s = shift;
  pack "c$s->{len}", @{$s->{buf}};
}

sub string_pack {
  my $s = shift;
  return {
    len => length $s,
    buf => [ unpack 'c*', $s ],
  }
}
```

While you would normally get the following output when unpacking a `String`

```
$string = {
  'len' => 12,
  'buf' => [
    72,
    101,
    108,
    108,
```

```
      111,
      32,
      87,
      111,
      114,
      108,
      100,
      33
    ]
};
```

you could just register the hooks using

```
$c->add_hooks(String => { pack   => \&string_pack,
                          unpack => \&string_unpack });
```

and you would get a nice human-readable string:

```
$string = 'Hello World!';
```

Packing a string turns out to be just as easy:

```
use Data::Hexdumper;

$data = $c->pack('String', 'Just another Perl hacker,');

print hexdump(data => $data);
```

This would print:

```
  0x0000 : 00 00 00 19 4A 75 73 74 20 61 6E 6F 74 68 65 72 : ....Just.another
  0x0010 : 20 50 65 72 6C 20 68 61 63 6B 65 72 2C          : .Perl.hacker,
```

1.4.2 Advanced hooks

But hooks are even more powerful. You can customize the arguments that are passed to your hooks and you can use arg to pass certain special arguments, such as the name of the type that is currently being processed by the hook.

The following example shows how it is easily possible to peek into the perl internals using hooks.

```
use Config;

$c = new Convert::Binary::C %CC, OrderMembers => 1;
$c->Include(["$Config{archlib}/CORE", @{$c->Include}]);
$c->parse(<<ENDC);
#include "EXTERN.h"
#include "perl.h"
ENDC

$c->add_hooks($_ => { unpack_ptr => [\&unpack_ptr,
                                     $c->arg(qw(SELF TYPE DATA))] })
    for qw( XPVAV XPVHV MAGIC MGVTBL HV );
```

First, we add the perl core include path and parse *perl.h*. Then, we add an `unpack_ptr` hook for a couple of the internal data types.

The `unpack_ptr` and `pack_ptr` hooks are called whenever a pointer to a certain data structure is processed. This is by far the most experimental part of the hooks feature, as this includes **any** kind of pointer. There's no way for the hook to know the difference between a plain pointer, or a pointer to a pointer, or a pointer to an array (this is because the difference doesn't matter anywhere else in Convert::Binary::C).

But the hook above makes use of another very interesting feature: It uses `arg` to pass special arguments to the hook subroutine. Usually, the hook subroutine is simply passed a single data argument. But using the above definition, it'll get a reference to the calling object (`SELF`), the name of the type being processed (`TYPE`) and the data (`DATA`).

But how does our hook look like?

```perl
sub unpack_ptr {
  my($self, $type, $ptr) = @_;
  $ptr or return '<NULL>';
  my $size = $self->sizeof($type);
  $self->unpack($type, unpack("P$size", pack('I', $ptr)));
}
```

As you can see, the hook is rather simple. First, it receives the arguments mentioned above. It performs a quick check if the pointer is `NULL` and shouldn't be processed any further. Next, it determines the size of the type being processed. And finally, it'll just use the Pn unpack template to read from that memory location and recursively call `unpack` to unpack the type. (And yes, this may of course again call other hooks.)

Now, let's test that:

```perl
my $ref = bless ["Boo!"], "Foo::Bar";
my $ptr = hex(("$ref" =~ /\(0x([[:xdigit:]]+)\)$/)[0]);

print Dumper(unpack_ptr($c, 'AV', $ptr));
```

Just for the fun of it, we create a blessed array reference. But how do we get a pointer to the corresponding `AV`? This is rather easy, as the address of the `AV` is just the hex value that appears when using the array reference in string context. So we just grab that and turn it into decimal. All that's left to do is just call our hook, as it can already handle `AV` pointers. And this is what we get:

```perl
$VAR1 = {
  'sv_any' => {
    'xav_array' => '137956856',
    'xav_fill' => '0',
    'xav_max' => '0',
    'xof_off' => '0',
    'xnv_nv' => '0',
    'xmg_magic' => '<NULL>',
    'xmg_stash' => {
      'sv_any' => {
        'xhv_array' => '0',
        'xhv_fill' => '0',
        'xhv_max' => '7',
        'xhv_keys' => '0',
        'xnv_nv' => '0',
```

```
        'xmg_magic' => '<NULL>',
        'xmg_stash' => '<NULL>',
        'xhv_riter' => '-1',
        'xhv_eiter' => '0',
        'xhv_pmroot' => '0',
        'xhv_name' => '138060760'
      },
      'sv_refcnt' => '2',
      'sv_flags' => '536870923'
    },
    'xav_alloc' => '137956856',
    'xav_arylen' => '0',
    'xav_flags' => '1'
  },
  'sv_refcnt' => '1',
  'sv_flags' => '4106'
};
```

Even though it is rather easy to do such stuff using `unpack_ptr` hooks, you should really know what you're doing and do it with extreme care because of the limitations mentioned above. It's really easy to run into segmentation faults when you're dereferencing pointers that are not yours.

1.4.3 Performance

Using hooks isn't for free. In performance-critical applications you have to keep in mind that hooks are actually perl subroutines and that they are called once for every value of a registered type that is being packed or unpacked. If only about 10% of the values require hooks to be called, you'll hardly notice the difference (if your hooks are implemented efficiently, that is). But if all values would require hooks to be called, that alone could easily make packing and unpacking very slow.

1.5 Methods

1.5.1 new

`new`

new OPTION1 => VALUE1, OPTION2 => VALUE2, ...

> The constructor is used to create a new Convert::Binary::C object. You can simply use

> $c = new Convert::Binary::C;

> without additional arguments to create an object, or you can optionally pass any arguments to the constructor that are described for the configure method.

1.5.2 configure

`configure`

configure OPTION

configure OPTION1 => VALUE1, OPTION2 => VALUE2, ...

This method can be used to configure an existing Convert::Binary::C object or to retrieve its current configuration.

To configure the object, the list of options consists of key and value pairs and must therefore contain an even number of elements. `configure` (and also `new` if used with configuration options) will throw an exception if you pass an odd number of elements. Configuration will normally look like this:

```
$c->configure( ByteOrder => 'BigEndian', IntSize => 2 );
```

To retrieve the current value of a configuration option, you must pass a single argument to `configure` that holds the name of the option, just like

```
$order = $c->configure( 'ByteOrder' );
```

If you want to get the values of all configuration options at once, you can call `configure` without any arguments and it will return a reference to a hash table that holds the whole object configuration. This can be conveniently used with the Data::Dumper module, for example:

```
use Convert::Binary::C;
use Data::Dumper;

$c = new Convert::Binary::C Define  => ['DEBUGGING', 'FOO=123'],
                            Include => ['/usr/include'];

print Dumper( $c->configure );
```

Which will print something like this:

```
$VAR1 = {
  'Define' => [
    'DEBUGGING',
    'FOO=123'
  ],
  'ByteOrder' => 'LittleEndian',
  'LongSize' => 4,
  'IntSize' => 4,
  'ShortSize' => 2,
  'HasMacroVAARGS' => 1,
  'Assert' => [],
  'UnsignedChars' => 0,
  'DoubleSize' => 8,
  'EnumType' => 'Integer',
  'PointerSize' => 4,
  'EnumSize' => 4,
  'DisabledKeywords' => [],
  'FloatSize' => 4,
  'LongLongSize' => 8,
  'Alignment' => 1,
  'LongDoubleSize' => 12,
  'KeywordMap' => {},
  'Include' => [
    '/usr/include'
  ],
```

```
    'HasCPPComments' => 1,
    'Warnings' => 0,
    'CompoundAlignment' => 1,
    'OrderMembers' => 0
};
```

Since you may not always want to write a `configure` call when you only want to change a single configuration item, you can use any configuration option name as a method name, like:

```
$c->ByteOrder( 'LittleEndian' ) if $c->IntSize < 4;
```

(Yes, the example doesn't make very much sense... ;-)

However, you should keep in mind that configuration methods that can take lists (namely `Include`, `Define` and `Assert`, but not `DisabledKeywords`) may behave slightly different than their `configure` equivalent. If you pass these methods a single argument that is an array reference, the current list will be **replaced** by the new one, which is just the behaviour of the corresponding `configure` call. So the following are equivalent:

```
$c->configure( Define => ['foo', 'bar=123'] );
$c->Define( ['foo', 'bar=123'] );
```

But if you pass a list of strings instead of an array reference (which cannot be done when using `configure`), the new list items are **appended** to the current list, so

```
$c = new Convert::Binary::C Include => ['/include'];
$c->Include( '/usr/include', '/usr/local/include' );
print Dumper( $c->Include );

$c->Include( ['/usr/local/include'] );
print Dumper( $c->Include );
```

will first print all three include paths, but finally only `/usr/local/include` will be configured:

```
$VAR1 = [
  '/include',
  '/usr/include',
  '/usr/local/include'
];
$VAR1 = [
  '/usr/local/include'
];
```

Furthermore, configuration methods can be chained together, as they return a reference to their object if called as a set method. So, if you like, you can configure your object like this:

```
$c = Convert::Binary::C->new( IntSize => 4 )
       ->Define( qw( __DEBUG__ DB_LEVEL=3 ) )
       ->ByteOrder( 'BigEndian' );

$c->configure( EnumType => 'Both', Alignment => 4 )
   ->Include( '/usr/include', '/usr/local/include' );
```

In the example above, `qw( ... )` is the word list quoting operator. It returns a list of all non-whitespace sequences, and is especially useful for configuring preprocessor defines or assertions. The following assignments are equivalent:

```
@array = ('one', 'two', 'three');
@array = qw(one two three);
```

You can configure the following options. Unknown options, as well as invalid values for an option, will cause the object to throw exceptions.

IntSize => 0 | 1 | 2 | 4 | 8

Set the number of bytes that are occupied by an integer. This is in most cases 2 or 4. If you set it to zero, the size of an integer on the host system will be used. This is also the default unless overridden by `CBC_DEFAULT_INT_SIZE` at compile time.

ShortSize => 0 | 1 | 2 | 4 | 8

Set the number of bytes that are occupied by a short integer. Although integers explicitly declared as `short` should be always 16 bit, there are compilers that make a short 8 bit wide. If you set it to zero, the size of a short integer on the host system will be used. This is also the default unless overridden by `CBC_DEFAULT_SHORT_SIZE` at compile time.

LongSize => 0 | 1 | 2 | 4 | 8

Set the number of bytes that are occupied by a long integer. If set to zero, the size of a long integer on the host system will be used. This is also the default unless overridden by `CBC_DEFAULT_LONG_SIZE` at compile time.

LongLongSize => 0 | 1 | 2 | 4 | 8

Set the number of bytes that are occupied by a long long integer. If set to zero, the size of a long long integer on the host system, or 8, will be used. This is also the default unless overridden by `CBC_DEFAULT_LONG_LONG_SIZE` at compile time.

FloatSize => 0 | 1 | 2 | 4 | 8 | 12 | 16

Set the number of bytes that are occupied by a single precision floating point value. If you set it to zero, the size of a `float` on the host system will be used. This is also the default unless overridden by `CBC_DEFAULT_FLOAT_SIZE` at compile time. For details on floating point support, see the section on Floating Point Values on page 65.

DoubleSize => 0 | 1 | 2 | 4 | 8 | 12 | 16

Set the number of bytes that are occupied by a double precision floating point value. If you set it to zero, the size of a `double` on the host system will be used. This is also the default unless overridden by `CBC_DEFAULT_DOUBLE_SIZE` at compile time. For details on floating point support, see the section on Floating Point Values on page 65.

LongDoubleSize => 0 | 1 | 2 | 4 | 8 | 12 | 16

Set the number of bytes that are occupied by a double precision floating point value. If you set it to zero, the size of a `long double` on the host system, or 12 will be used. This is also the default unless overridden by `CBC_DEFAULT_LONG_DOUBLE_SIZE` at compile time. For details on floating point support, see the section on Floating Point Values on page 65.

PointerSize => 0 | 1 | 2 | 4 | 8

Set the number of bytes that are occupied by a pointer. This is in most cases 2 or 4. If you set it to zero, the size of a pointer on the host system will be used. This is also the default unless overridden by `CBC_DEFAULT_PTR_SIZE` at compile time.

EnumSize => -1 | 0 | 1 | 2 | 4 | 8

Set the number of bytes that are occupied by an enumeration type. On most systems, this is equal to the size of an integer, which is also the default. However, for some compilers, the size of an enumeration type depends on the size occupied by the largest enumerator. So the size may vary between 1 and 8. If you have

```
enum foo {
  ONE = 100, TWO = 200
};
```

this will occupy one byte because the enum can be represented as an unsigned one-byte value. However,

```
enum foo {
  ONE = -100, TWO = 200
};
```

will occupy two bytes, because the -100 forces the type to be signed, and 200 doesn't fit into a signed one-byte value. Therefore, the type used is a signed two-byte value. If this is the behaviour you need, set the EnumSize to 0.

Some compilers try to follow this strategy, but don't care whether the enumeration has signed values or not. They always declare an enum as signed. On such a compiler, given

```
enum one { ONE = -100, TWO = 100 };
enum two { ONE =  100, TWO = 200 };
```

enum one will occupy only one byte, while enum two will occupy two bytes, even though it could be represented by a unsigned one-byte value. If this is the behaviour of your compiler, set EnumSize to -1.

Alignment => 0 | 1 | 2 | 4 | 8 | 16

Set the struct member alignment. This option controls where padding bytes are inserted between struct members. It globally sets the alignment for all structs/unions. However, this can be overridden from within the source code with the common `pack` pragma as explained in the section on Supported pragma directives on page 14. The default alignment is 1, which means no padding bytes are inserted. A setting of 0 means *native* alignment, i.e. the alignment of the system that Convert::Binary::C has been compiled on. You can determine the native properties using the `native` function.

The `Alignment` option is similar to the `-Zp[n]` option of the Intel compiler. It globally specifies the maximum boundary to which struct members are aligned. Consider the following structure and the sizes of `char`, `short`, `long` and `double` being 1, 2, 4 and 8, respectively.

```
struct align {
  char   a;
  short  b, c;
  long   d;
  double e;
};
```

With an alignment of 1 (the default), the struct members would be packed tightly:

```
0   1   2   3   4   5   6   7   8   9   10  11  12
+---+---+---+---+---+---+---+---+---+---+---+---+
| a |   b   |   c   |       d       |           ...
+---+---+---+---+---+---+---+---+---+---+---+---+

    12  13  14  15  16  17
     +---+---+---+---+---+
...     e               |
     +---+---+---+---+---+
```

With an alignment of 2, the struct members larger than one byte would be aligned to 2-byte boundaries, which results in a single padding byte between `a` and `b`.

```
0   1   2   3   4   5   6   7   8   9   10  11  12
+---+---+---+---+---+---+---+---+---+---+---+---+
| a | * |   b   |   c   |       d       |       ...
+---+---+---+---+---+---+---+---+---+---+---+---+
```

```
    12  13  14  15  16  17  18
     +---+---+---+---+---+---+
...          e              |
     +---+---+---+---+---+---+
```

With an alignment of 4, the struct members of size 2 would be aligned to 2-byte boundaries and larger struct members would be aligned to 4-byte boundaries:

```
0   1   2   3   4   5   6   7   8   9   10  11  12
+---+---+---+---+---+---+---+---+---+---+---+---+
| a | * |   b   |   c   | * | * |       d       | ...
+---+---+---+---+---+---+---+---+---+---+---+---+
    12  13  14  15  16  17  18  19  20
     +---+---+---+---+---+---+---+---+
... |           e               |
     +---+---+---+---+---+---+---+---+
```

This layout of the struct members allows the compiler to generate optimized code because aligned members can be accessed more easily by the underlying architecture.

Finally, setting the alignment to 8 will align `double`s to 8-byte boundaries:

```
0   1   2   3   4   5   6   7   8   9   10  11  12
+---+---+---+---+---+---+---+---+---+---+---+---+
| a | * |   b   |   c   | * | * |       d       | ...
+---+---+---+---+---+---+---+---+---+---+---+---+
    12  13  14  15  16  17  18  19  20  21  22  23  24
     +---+---+---+---+---+---+---+---+---+---+---+---+
... | * | * | * | * |               e               |
     +---+---+---+---+---+---+---+---+---+---+---+---+
```

Further increasing the alignment does not alter the layout of our structure, as only members larger that 8 bytes would be affected.

The alignment of a structure depends on its largest member and on the setting of the `Alignment` option. With `Alignment` set to 2, a structure holding a `long` would be aligned to a 2-byte boundary, while a structure containing only `char`s would have no alignment restrictions. (Unfortunately, that's not the whole story. See the `CompoundAlignment` option for details.)

Here's another example. Assuming 8-byte alignment, the following two structs will both have a size of 16 bytes:

```
struct one {
  char   c;
  double d;
};

struct two {
  double d;
  char   c;
};
```

This is clear for `struct one`, because the member `d` has to be aligned to an 8-byte boundary, and thus 7 padding bytes are inserted after `c`. But for `struct two`, the padding bytes are inserted *at the end* of the structure, which doesn't make much sense immediately. However, it makes perfect sense if you think about an array of `struct two`. Each `double` has to be aligned to an 8-byte boundary, an thus each array element would have to occupy 16 bytes. With that in mind, it would be strange if a `struct two` variable would have a different size. And it would make the widely used construct

```
struct two array[] = { {1.0, 0}, {2.0, 1} };
int elements = sizeof(array) / sizeof(struct two);
```

impossible.

The alignment behaviour described here seems to be common for all compilers. However, not all compilers have an option to configure their default alignment.

CompoundAlignment => 0 | 1 | 2 | 4 | 8 | 16

Usually, the alignment of a compound (i.e. a `struct` or a `union`) depends only on its largest member and on the setting of the `Alignment` option. There are, however, architectures and compilers where compounds can have different alignment constraints.

For most platforms and compilers, the alignment constraint for compounds is 1 byte. That is, on most platforms

```
struct onebyte {
  char byte;
};
```

will have an alignment of 1 and also a size of 1. But if you take an ARM architecture, the above `struct onebyte` will have an alignment of 4, and thus also a size of 4.

You can configure this by setting `CompoundAlignment` to 4. This will ensure that the alignment of compounds is always 4.

Setting `CompoundAlignment` to 0 means *native* compound alignment, i.e. the compound alignment of the system that Convert::Binary::C has been compiled on. You can determine the native properties using the native function.

There are also compilers for certain platforms that allow you to adjust the compound alignment. If you're not aware of the fact that your compiler/architecture has a compound alignment other than 1, strange things can happen. If, for example, the compound alignment is 2 and you have something like

```
typedef unsigned char U8;

struct msg_head {
  U8 cmd;
  struct {
    U8 hi;
    U8 low;
  } crc16;
  U8 len;
};
```

there will be one padding byte inserted before the embedded `crc16` struct and after the `len` member, which is most probably not what was intended:

```
0     1     2     3     4     5     6
+-----+-----+-----+-----+-----+-----+
| cmd |  *  | hi  | low | len |  *  |
+-----+-----+-----+-----+-----+-----+
```

Note that both `#pragma pack` and the `Alignment` option can override `CompoundAlignment`. If you set `CompoundAlignment` to 4, but `Alignment` to 2, compounds will actually be aligned on 2-byte boundaries.

ByteOrder => 'BigEndian' | 'LittleEndian'

Set the byte order for integers larger than a single byte. Little endian (Intel, least significant byte first) and big endian (Motorola, most significant byte first) byte order are supported. The default byte order is the same as the byte order of the host system unless overridden by `CBC_DEFAULT_BYTEORDER` at compile time.

EnumType => 'Integer' | 'String' | 'Both'

This option controls the type that enumeration constants will have in data structures returned by the unpack method. If you have the following definitions:

```
typedef enum {
   SUNDAY, MONDAY, TUESDAY, WEDNESDAY,
   THURSDAY, FRIDAY, SATURDAY
} Weekday;

typedef enum {
   JANUARY, FEBRUARY, MARCH, APRIL, MAY, JUNE, JULY,
   AUGUST, SEPTEMBER, OCTOBER, NOVEMBER, DECEMBER
} Month;

typedef struct {
   int      year;
   Month    month;
   int      day;
   Weekday  weekday;
} Date;
```

and a byte string that holds a packed Date struct, then you'll get the following results from a call to the unpack method.

Integer

Enumeration constants are returned as plain integers. This is fast, but may be not very useful. It is also the default.

```
$date = {
   'weekday' => 1,
   'month' => 0,
   'day' => 7,
   'year' => 2002
};
```

String

Enumeration constants are returned as strings. This will create a string constant for every unpacked enumeration constant and thus consumes more time and memory. However, the result may be more useful.

```
$date = {
   'weekday' => 'MONDAY',
   'month' => 'JANUARY',
   'day' => 7,
   'year' => 2002
};
```

Both

Enumeration constants are returned as double typed scalars. If evaluated in string context, the enumeration constant will be a string, if evaluated in numeric context, the enumeration constant will be an integer.

```
$date = $c->EnumType('Both')->unpack('Date', $binary);

printf "Weekday = %s (%d)\n\n", $date->{weekday},
                                $date->{weekday};

if( $date->{month} == 0 ) {
  print "It's $date->{month}, happy new year!\n\n";
}

print Dumper( $date );
```

This will print:

```
Weekday = MONDAY (1)
It's JANUARY, happy new year!
$VAR1 = {
  'weekday' => 'MONDAY',
  'month' => 'JANUARY',
  'day' => 7,
  'year' => 2002
};
```

DisabledKeywords => [KEYWORDS]

This option allows you to selectively deactivate certain keywords in the C parser. Some C compilers don't have the complete ANSI keyword set, i.e. they don't recognize the keywords `const` or `void`, for example. If you do

```
typedef int void;
```

on such a compiler, this will usually be ok. But if you parse this with an ANSI compiler, it will be a syntax error. To parse the above code correctly, you have to disable the `void` keyword in the Convert::Binary::C parser:

```
$c->DisabledKeywords( [qw( void )] );
```

By default, the Convert::Binary::C parser will recognize the keywords `inline` and `restrict`. If your compiler doesn't have these new keywords, it usually doesn't matter. Only if you're using the keywords as identifiers, like in

```
typedef struct inline {
  int a, b;
} restrict;
```

you'll have to disable these ISO-C99 keywords:

```
$c->DisabledKeywords( [qw( inline restrict )] );
```

The parser allows you to disable the following keywords:

```
asm
auto
const
double
enum
extern
float
inline
long
register
restrict
short
signed
static
unsigned
void
volatile
```

KeywordMap => { KEYWORD => TOKEN, ... }

This option allows you to add new keywords to the parser. These new keywords can either be mapped to existing tokens or simply ignored. For example, recent versions of the GNU

compiler recognize the keywords `__signed__` and `__extension__`. The first one obviously is a synonym for `signed`, while the second one is only a marker for a language extension.

Using the preprocessor, you could of course do the following:

```
$c->Define( qw( __signed__=signed __extension__= ) );
```

However, the preprocessor symbols could be undefined or redefined in the code, and

```
#ifdef __signed__
# undef __signed__
#endif

typedef __extension__ __signed__ long long s_quad;
```

would generate a parse error, because `__signed__` is an unexpected identifier.

Instead of utilizing the preprocessor, you'll have to create mappings for the new keywords directly in the parser using `KeywordMap`. In the above example, you want to map `__signed__` to the built-in C keyword `signed` and ignore `__extension__`. This could be done with the following code:

```
$c->KeywordMap( {
                __signed__    => 'signed',
                __extension__ => undef,
            } );
```

You can specify any valid identifier as hash key, and either a valid C keyword or `undef` as hash value. Having configured the object that way, you could parse even

```
#ifdef __signed__
# undef __signed__
#endif

typedef __extension__ __signed__ long long s_quad;
```

without problems.

Note that `KeywordMap` and `DisabledKeywords` perfectly work together. You could, for example, disable the `signed` keyword, but still have `__signed__` mapped to the original `signed` token:

```
$c->configure( DisabledKeywords => [ 'signed' ],
               KeywordMap       => { __signed__  => 'signed' } );
```

This would allow you to define

```
typedef __signed__ long signed;
```

which would normally be a syntax error because `signed` cannot be used as an identifier.

`UnsignedChars => 0 | 1`

> Use this boolean option if you want characters to be unsigned if specified without an explicit `signed` or `unsigned` type specifier. By default, characters are signed.

`Warnings => 0 | 1`

> Use this boolean option if you want warnings to be issued during the parsing of source code. Currently, warnings are only reported by the preprocessor, so don't expect the output to cover everything.
>
> By default, warnings are turned off and only errors will be reported. However, even these errors are turned off if you run without the `-w` flag.

`HasCPPComments => 0 | 1`

> Use this option to turn C++ comments on or off. By default, C++ comments are enabled. Disabling C++ comments may be necessary if your code includes strange things like:

```
one = 4 //* <- divide */ 4;
two = 2;
```

With C++ comments, the above will be interpreted as

```
one = 4
two = 2;
```

which will obviously be a syntax error, but without C++ comments, it will be interpreted as

```
one = 4 / 4;
two = 2;
```

which is correct.

HasMacroVAARGS => 0 | 1

Use this option to turn the `__VA_ARGS__` macro expansion on or off. If this is enabled (which is the default), you can use variable length argument lists in your preprocessor macros.

```
#define DEBUG( ... )  fprintf( stderr, __VA_ARGS__ )
```

There's normally no reason to turn that feature off.

Include => [INCLUDES]

Use this option to set the include path for the internal preprocessor. The option value is a reference to an array of strings, each string holding a directory that should be searched for includes.

Define => [DEFINES]

Use this option to define symbols in the preprocessor. The option value is, again, a reference to an array of strings. Each string can be either just a symbol or an assignment to a symbol. This is completely equivalent to what the `-D` option does for most preprocessors.

The following will define the symbol `FOO` and define `BAR` to be `12345`:

```
$c->configure( Define => [qw(FOO BAR=12345)] );
```

Assert => [ASSERTIONS]

Use this option to make assertions in the preprocessor. If you don't know what assertions are, don't be concerned, since they're deprecated anyway. They are, however, used in some system's include files. The value is an array reference, just like for the macro definitions. Only the way the assertions are defined is a bit different and mimics the way they are defined with the `#assert` directive:

```
$c->configure( Assert => ['foo(bar)'] );
```

OrderMembers => 0 | 1

When using unpack on compounds and iterating over the returned hash, the order of the compound members is generally not preserved due to the nature of hash tables. It is not even guaranteed that the order is the same between different runs of the same program. This can be very annoying if you simply use to dump your data structures and the compound members always show up in a different order.

By setting `OrderMembers` to a non-zero value, all hashes returned by unpack are tied to a class that preserves the order of the hash keys. This way, all compound members will be returned in the correct order just as they are defined in your C code.

```
use Convert::Binary::C;
use Data::Dumper;

$c = Convert::Binary::C->new->parse( <<'ENDC' );
struct test {
  char one;
```

```
      char two;
      struct {
        char never;
        char change;
        char this;
        char order;
      } three;
      char four;
    };
    ENDC

    $data = "Convert";

    $u1 = $c->unpack( 'test', $data );
    $c->OrderMembers(1);
    $u2 = $c->unpack( 'test', $data );

    print Data::Dumper->Dump( [$u1, $u2], [qw(u1 u2)] );
```

This will print something like:

```
    $u1 = {
      'three' => {
        'change' => 118,
        'order' => 114,
        'this' => 101,
        'never' => 110
      },
      'one' => 67,
      'two' => 111,
      'four' => 116
    };
    $u2 = {
      'one' => '67',
      'two' => '111',
      'three' => {
        'never' => '110',
        'change' => '118',
        'this' => '101',
        'order' => '114'
      },
      'four' => '116'
    };
```

To be able to use this option, you have to install either the Tie::Hash::Indexed or the Tie::IxHash module. If both are installed, Convert::Binary::C will give preference to Tie::Hash::Indexed because it's faster.

When using this option, you should keep in mind that tied hashes are significantly slower and consume more memory than ordinary hashes, even when the class they're tied to is implemented efficiently. So don't turn this option on if you don't have to.

You can also influence hash member ordering by using the CBC_ORDER_MEMBERS environment variable.

You can reconfigure all options even after you have parsed some code. The changes will be applied to the already parsed definitions. This works as long as array lengths are not affected by the changes. If you have Alignment and IntSize set to 4 and parse code like this

```
typedef struct {
  char abc;
  int  day;
} foo;

struct bar {
  foo  zap[2*sizeof(foo)];
};
```

the array `zap` in `struct bar` will obviously have 16 elements. If you reconfigure the alignment to 1 now, the size of `foo` is now 5 instead of 8. While the alignment is adjusted correctly, the number of elements in array `zap` will still be 16 and will not be changed to 10.

1.5.3   parse

`parse` **CODE**

Parses a string of valid C code. All enumeration, compound and type definitions are extracted. You can call the `parse` and `parse_file` methods as often as you like to add further definitions to the Convert::Binary::C object.

`parse` will throw an exception if an error occurs. On success, the method returns a reference to its object.

See the section on Parsing C code on page 12 for an example.

1.5.4   parse_file

`parse_file` **FILE**

Parses a C source file. All enumeration, compound and type definitions are extracted. You can call the `parse` and `parse_file` methods as often as you like to add further definitions to the Convert::Binary::C object.

`parse_file` will throw an exception if an error occurs. On success, the method returns a reference to its object.

See the section on Parsing C code on page 12 for an example.

You must be aware that the preprocessor is reset with every call to `parse` or `parse_file`. Also, you may use types previously defined, but you are not allowed to redefine types.

When you're parsing C source files instead of C header files, note that local definitions are ignored. This means that type definitions hidden within functions will not be recognized by Convert::Binary::C. This is necessary because different functions (even different blocks within the same function) can define types with the same name:

```
void my_func( int i )
{
  if( i < 10 ) {
    enum digit { ONE, TWO, THREE } x = ONE;
    printf("%d, %d\n", i, x);
  }
  else {
    enum digit { THREE, TWO, ONE } x = ONE;
    printf("%d, %d\n", i, x);
  }
}
```

The above is a valid piece of C code, but it's not possible for Convert::Binary::C to distinguish between the different definitions of `enum digit`, as they're only defined locally within the corresponding block.

1.5.5 clean

`clean`

Clears all information that has been collected during previous calls to `parse` or `parse_file`. You can use this method if you want to parse some entirely different code, but with the same configuration.

The `clean` method returns a reference to its object.

1.5.6 clone

`clone`

Makes the object return an exact independent copy of itself.

```
$c = new Convert::Binary::C Include => ['/usr/include'];
$c->parse_file( 'definitions.c' );
$clone = $c->clone;
```

The above code is technically equivalent (Mostly. Actually, using `sourcify` and `parse` might alter the order of the parsed data, which would make methods such as `compound` return the definitions in a different order.) to:

```
$c = new Convert::Binary::C Include => ['/usr/include'];
$c->parse_file( 'definitions.c' );
$clone = new Convert::Binary::C %{$c->configure};
$clone->parse( $c->sourcify );
```

Using `clone` is just a lot faster.

1.5.7 def

`def` **NAME**

`def` **TYPE**

If you need to know if a definition for a certain type name exists, use this method. You pass it the name of an enum, struct, union or typedef, and it will return a non-empty string being either `"enum"`, `"struct"`, `"union"`, or `"typedef"` if there's a definition for the type in question, an empty string if there's no such definition, or `undef` if the name is completely unknown. If the type can be interpreted as a basic type, `"basic"` will be returned.

If you pass in a TYPE, the output will be slightly different. If the specified member exists, the `def` method will return `"member"`. If the member doesn't exist, or if the type cannot have members, the empty string will be returned. Again, if the name of the type is completely unknown, `undef` will be returned. This may be useful if you want to check if a certain member exists within a compound, for example.

```
use Convert::Binary::C;

my $c = Convert::Binary::C->new->parse( <<'ENDC' );
```

```
  typedef struct __not  not;
  typedef struct __not *ptr;

  struct foo {
    enum bar *xxx;
  };

  typedef int quad[4];

  ENDC

  for my $type ( qw( not ptr foo bar xxx foo.xxx foo.abc
                    xxx.yyy quad quad[3] quad[4] short[1] ),
               'unsigned long' )
  {
    my $def = $c->def( $type );
    printf "%-14s  =>  %s\n", $type, defined $def
                                   ? "'$def'" : 'undef';
  }
```

The following would be returned by the def method:

```
not            =>  ''
ptr            =>  'typedef'
foo            =>  'struct'
bar            =>  ''
xxx            =>  undef
foo.xxx        =>  'member'
foo.abc        =>  ''
xxx.yyy        =>  undef
quad           =>  'typedef'
quad[3]        =>  'member'
quad[4]        =>  ''
short[1]       =>  undef
unsigned long  =>  'basic'
```

So, if def returns a non-empty string, you can safely use any other method with that type's name or with that member expression.

In cases where the typedef namespace overlaps with the namespace of enums/structs/unions, the def method will give preference to the typedef and will thus return the string "typedef". You could however force interpretation as an enum, struct or union by putting "enum", "struct" or "union" in front of the type's name.

1.5.8 pack

pack **TYPE**

pack **TYPE, DATA**

pack **TYPE, DATA, STRING**

Use this method to pack a complex data structure into a binary string according to a type definition that has been previously parsed. DATA must be a scalar matching the type definition. C structures and unions are represented by references to Perl hashes, C arrays by references to Perl arrays.

```
use Convert::Binary::C;
use Data::Dumper;
use Data::Hexdumper;

$c = Convert::Binary::C->new( ByteOrder => 'BigEndian',
                              LongSize  => 4,
                              ShortSize => 2 )
                     ->parse( <<'ENDC' );
struct test {
  char    ary[3];
  union {
    short word[2];
    long  quad;
  }       uni;
};
ENDC
```

Hashes don't have to contain a key for each compound member and arrays may be truncated:

```
$binary = $c->pack( 'test', { ary => [1, 2], uni => { quad => 42 } } );
```

Elements not defined in the Perl data structure will be set to zero in the packed byte string. If you pass `undef` as or simply omit the second parameter, the whole string will be initialized with zero bytes. On success, the packed byte string is returned.

```
print hexdump( data => $binary );
```

The above code would print:

```
0x0000 : 01 02 00 00 00 00 2A                             : ......*
```

You could also use unpack and dump the data structure.

```
$unpacked = $c->unpack( 'test', $binary );
print Data::Dumper->Dump( [$unpacked], ['unpacked'] );
```

This would print:

```
$unpacked = {
  'uni' => {
    'word' => [
      0,
      42
    ],
    'quad' => 42
  },
  'ary' => [
    1,
    2,
    0
  ]
};
```

If TYPE refers to a compound object, you may pack any member of that compound object. Simply add a member expression to the type name, just as you would access the member in C:

```
$array = $c->pack( 'test.ary', [1, 2, 3] );
print hexdump( data => $array );

$value = $c->pack( 'test.uni.word[1]', 2 );
print hexdump( data => $value );
```

This would give you:

```
0x0000 : 01 02 03                                        : ...
0x0000 : 00 02                                           : ..
```

Call pack with the optional STRING argument if you want to use an existing binary string to insert the data. If called in a void context, pack will directly modify the string you passed as the third argument. Otherwise, a copy of the string is created, and pack will modify and return the copy, so the original string will remain unchanged.

The 3-argument version may be useful if you want to change only a few members of a complex data structure without having to unpack everything, change the members, and then pack again (which could waste lots of memory and CPU cycles). So, instead of doing something like

```
$test = $c->unpack( 'test', $binary );
$test->{uni}{quad} = 4711;
$new = $c->pack( 'test', $test );
```

to change the `uni.quad` member of `$packed`, you could simply do either

```
$new = $c->pack( 'test', { uni => { quad => 4711 } }, $binary );
```

or

```
$c->pack( 'test', { uni => { quad => 4711 } }, $binary );
```

while the latter would directly modify `$packed`. Besides this code being a lot shorter (and perhaps even more readable), it can be significantly faster if you're dealing with really big data blocks.

If the length of the input string is less than the size required by the type, the string (or its copy) is extended and the extended part is initialized to zero. If the length is more than the size required by the type, the string is kept at that length, and also a copy would be an exact copy of that string.

```
$too_short = pack "C*", (1 .. 4);
$too_long  = pack "C*", (1 .. 20);

$c->pack( 'test', { uni => { quad => 0x4711 } }, $too_short );
print "too_short:\n", hexdump( data => $too_short );

$copy = $c->pack( 'test', { uni => { quad => 0x4711 } }, $too_long );
print "\ncopy:\n", hexdump( data => $copy );
```

This would print:

```
too_short:
  0x0000 : 01 02 03 00 00 47 11                          : .....G.

copy:
  0x0000 : 01 02 03 00 00 47 11 08 09 0A 0B 0C 0D 0E 0F 10 : .....G..........
  0x0010 : 11 12 13 14                                    : ....
```

1.5.9 unpack

unpack **TYPE, STRING**

> Use this method to unpack a binary string and create an arbitrarily complex Perl data structure based on a previously parsed type definition.

```
use Convert::Binary::C;
use Data::Dumper;

$c = Convert::Binary::C->new( ByteOrder => 'BigEndian',
                              LongSize  => 4,
                              ShortSize => 2 )
                    ->parse( <<'ENDC' );
struct test {
  char    ary[3];
  union {
    short word[2];
    long *quad;
  }       uni;
};
ENDC

# Generate some binary dummy data
$binary = pack "C*", (1 .. $c->sizeof('test'));
```

> On failure, e.g. if the specified type cannot be found, the method will throw an exception. On success, a reference to a complex Perl data structure is returned, which can directly be dumped using the Data::Dumper module:

```
$unpacked = $c->unpack( 'test', $binary );
print Dumper( $unpacked );
```

> This would print:

```
$VAR1 = {
  'uni' => {
    'word' => [
      1029,
      1543
    ],
    'quad' => 67438087
  },
  'ary' => [
    1,
    2,
    3
  ]
};
```

> If TYPE refers to a compound object, you may unpack any member of that compound object. Simply add a member expression to the type name, just as you would access the member in C:

```
$binary2 = substr $binary, $c->offsetof('test', 'uni.word');
```

```
$unpack1 = $unpacked->{uni}{word};
$unpack2 = $c->unpack( 'test.uni.word', $binary2 );

print Data::Dumper->Dump( [$unpack1, $unpack2], [qw(unpack1 unpack2)] );
```

You will find that the output is exactly the same for both `$unpack1` and `$unpack2`:

```
$unpack1 = [
  1029,
  1543
];
$unpack2 = [
  1029,
  1543
];
```

When unpack is called in list context, it will unpack as many elements as possible from STRING, including zero if STRING is not long enough.

1.5.10   sizeof

`sizeof` **TYPE**

This method will return the size of a C type in bytes. If it cannot find the type, it will throw an exception.

If the type defines some kind of compound object, you may ask for the size of a member of that compound object:

```
$size = $c->sizeof( 'test.uni.word[1]' );
```

This would set `$size` to 2.

1.5.11   typeof

`typeof` **TYPE**

This method will return the type of a C member. While this only makes sense for compound types, it's legal to also use it for non-compound types. If it cannot find the type, it will throw an exception.

The typeof method can be used on any valid member, even on arrays or unnamed types. It will always return a string that holds the name (or in case of unnamed types only the class) of the type, optionally followed by a '`*`' character to indicate it's a pointer type, and optionally followed by one or more array dimensions if it's an array type. If the type is a bitfield, the type name is followed by a colon and the number of bits.

```
struct test {
  char    ary[3];
  union {
    short word[2];
    long *quad;
  }       uni;
  struct {
    unsigned short six:6;
    unsigned short ten:10;
  }       bits;
};
```

Given the above C code has been parsed, calls to typeof would return the following values:

```
$c->typeof('test')          => 'struct test'
$c->typeof('test.ary')      => 'char [3]'
$c->typeof('test.uni')      => 'union'
$c->typeof('test.uni.quad')  => 'long *'
$c->typeof('test.uni.word')  => 'short [2]'
$c->typeof('test.uni.word[1]') => 'short'
$c->typeof('test.bits')     => 'struct'
$c->typeof('test.bits.six')  => 'unsigned short :6'
$c->typeof('test.bits.ten')  => 'unsigned short :10'
```

1.5.12   offsetof

offsetof **TYPE, MEMBER**

You can use offsetof just like the C macro of same denominator. It will simply return the offset (in bytes) of MEMBER relative to TYPE.

```
use Convert::Binary::C;

$c = Convert::Binary::C->new( Alignment   => 4
                            , LongSize    => 4
                            , PointerSize => 4
                            )
                   ->parse( <<'ENDC' );
typedef struct {
  char abc;
  long day;
  int *ptr;
} week;

struct test {
  week zap[8];
};
ENDC

@args = (
  ['test',       'zap[5].day'  ],
  ['test.zap[2]', 'day'        ],
  ['test',       'zap[5].day+1'],
);

for( @args ) {
  my $offset = eval { $c->offsetof( @$_ ) };
  printf "\$c->offsetof( '%s', '%s' ) => $offset\n", @$_;
}
```

The final loop will print:

```
$c->offsetof( 'test', 'zap[5].day' ) => 64
$c->offsetof( 'test.zap[2]', 'day' ) => 4
$c->offsetof( 'test', 'zap[5].day+1' ) => 65
```

- The first iteration simply shows that the offset of `zap[5].day` is 64 relative to the beginning of `struct test`.

- You may additionally specify a member for the type passed as the first argument, as shown in the second iteration.

- Even the offset suffix is supported by `offsetof`, so the third iteration will correctly print 65.

Unlike the C macro, `offsetof` also works on array types.

```
$offset = $c->offsetof( 'test.zap', '[3].ptr+2' );
print "offset = $offset";
```

This will print:

```
offset = 46
```

If TYPE is a compound, MEMBER may optionally be prefixed with a dot, so

```
printf "offset = %d\n", $c->offsetof( 'week', 'day' );
printf "offset = %d\n", $c->offsetof( 'week', '.day' );
```

are both equivalent and will print

```
offset = 4
offset = 4
```

This allows to

- use the C macro style, without a leading dot, and

- directly use the output of the `member` method, which includes a leading dot for compound types, as input for the MEMBER argument.

1.5.13   member

`member` **TYPE**

`member` **TYPE, OFFSET**

You can think of `member` as being the reverse of the `offsetof` method. However, as this is more complex, there's no equivalent to `member` in the C language.

Usually this method is used if you want to retrieve the name of the member that is located at a specific offset of a previously parsed type.

```
use Convert::Binary::C;

$c = Convert::Binary::C->new( Alignment   => 4
                            , LongSize    => 4
                            , PointerSize => 4
                            )
                    ->parse( <<'ENDC' );
typedef struct {
  char abc;
  long day;
  int *ptr;
} week;
```

```
struct test {
  week zap[8];
};
ENDC

for my $offset ( 24, 39, 69, 99 ) {
  print "\$c->member( 'test', $offset )";
  my $member = eval { $c->member( 'test', $offset ) };
  print $@ ? "\n  exception: $@" : " => '$member'\n";
}
```

This will print:

```
$c->member( 'test', 24 ) => '.zap[2].abc'
$c->member( 'test', 39 ) => '.zap[3]+3'
$c->member( 'test', 69 ) => '.zap[5].ptr+1'
$c->member( 'test', 99 )
  exception: Offset 99 out of range (0 <= offset < 96)
```

- The output of the first iteration is obvious. The member `zap[2].abc` is located at offset 24 of `struct test`.

- In the second iteration, the offset points into a region of padding bytes and thus no member of `week` can be named. Instead of a member name the offset relative to `zap[3]` is appended.

- In the third iteration, the offset points to `zap[5].ptr`. However, `zap[5].ptr` is located at 68, not at 69, and thus the remaining offset of 1 is also appended.

- The last iteration causes an exception because the offset of 99 is not valid for `struct test` since the size of `struct test` is only 96.

You can additionally specify a member for the type passed as the first argument:

```
$member = $c->member('test.zap[2]', 6);
print $member;
```

This will print:

```
.day+2
```

Like offsetof, member also works on array types:

```
$member = $c->member('test.zap', 42);
print $member;
```

This will print:

```
[3].day+2
```

While the behaviour for `struct`s is quite obvious, the behaviour for `union`s is rather tricky. As a single offset usually references more than one member of a union, there are certain rules that the algorithm uses for determining the *best* member.

- The first non-compound member that is referenced without an offset has the highest priority.

- If no member is referenced without an offset, the first non-compound member that is referenced with an offset will be returned.

- Otherwise the first padding region that is encountered will be taken.

As an example, given 4-byte-alignment and the union

```
union choice {
  struct {
    char  color[2];
    long  size;
    char  taste;
  }      apple;
  char   grape[3];
  struct {
    long  weight;
    short price[3];
  }      melon;
};
```

the `member` method would return what is shown in the *Member* column of the following table. The *Type* column shows the result of the `typeof` method when passing the corresponding member.

```
Offset    Member                Type
-------------------------------------
   0      .apple.color[0]       'char'
   1      .apple.color[1]       'char'
   2      .grape[2]             'char'
   3      .melon.weight+3       'long'
   4      .apple.size           'long'
   5      .apple.size+1         'long'
   6      .melon.price[1]       'short'
   7      .apple.size+3         'long'
   8      .apple.taste          'char'
   9      .melon.price[2]+1     'short'
  10      .apple+10             'struct'
  11      .apple+11             'struct'
```

It's like having a stack of all the union members and looking through the stack for the shiniest piece you can see. The beginning of a member (denoted by uppercase letters) is always shinier than the rest of a member, while padding regions (denoted by dashes) aren't shiny at all.

```
Offset  0   1   2   3   4   5   6   7   8   9   10  11
------------------------------------------------------
apple  (C) (C)  -   -  (S) (s)  s  (s) (T)  -  (-) (-)
grape   G   G  (G)
melon   W   w   w  (w)  P   p  (P)  p   P  (p)  -   -
```

If you look through that stack from top to bottom, you'll end up at the parenthesized members.

Alternatively, if you're not only interested in the *best* member, you can call `member` in list context, which makes it return *all* members referenced by the given offset.

```
Offset    Member                Type
-------------------------------------
   0      .apple.color[0]       'char'
          .grape[0]             'char'
          .melon.weight         'long'
   1      .apple.color[1]       'char'
          .grape[1]             'char'
```

```
            .melon.weight+1      'long'
    2       .grape[2]            'char'
            .melon.weight+2      'long'
            .apple+2             'struct'
    3       .melon.weight+3      'long'
            .apple+3             'struct'
    4       .apple.size          'long'
            .melon.price[0]      'short'
    5       .apple.size+1        'long'
            .melon.price[0]+1    'short'
    6       .melon.price[1]      'short'
            .apple.size+2        'long'
    7       .apple.size+3        'long'
            .melon.price[1]+1    'short'
    8       .apple.taste         'char'
            .melon.price[2]      'short'
    9       .melon.price[2]+1    'short'
            .apple+9             'struct'
   10       .apple+10            'struct'
            .melon+10            'struct'
   11       .apple+11            'struct'
            .melon+11            'struct'
```

The first member returned is always the *best* member. The other members are sorted according to the rules given above. This means that members referenced without an offset are followed by members referenced with an offset. Padding regions will be at the end.

If OFFSET is not given in the method call, `member` will return a list of *all* possible members of TYPE.

```
  print "$_\n" for $c->member( 'choice' );
```

This will print:

```
  .apple.color[0]
  .apple.color[1]
  .apple.size
  .apple.taste
  .grape[0]
  .grape[1]
  .grape[2]
  .melon.weight
  .melon.price[0]
  .melon.price[1]
  .melon.price[2]
```

In scalar context, the number of possible members is returned.

1.5.14 initializer

`initializer` **TYPE**

`initializer` **TYPE, DATA**

The `initializer` method can be used retrieve an initializer string for a certain TYPE. This can be useful if you have to initialize only a couple of members in a huge compound type or if you simply want to generate initializers automatically.

```
struct date {
  unsigned year : 12;
  unsigned month:  4;
  unsigned day  :  5;
  unsigned hour :  5;
  unsigned min  :  6;
};

typedef struct {
  enum { DATE, QWORD } type;
  short number;
  union {
    struct date   date;
    unsigned long qword;
  } choice;
} data;
```

Given the above code has been parsed

```
$init = $c->initializer( 'data' );
print "data x = $init;\n";
```

would print the following:

```
data x = {
     0,
     0,
     {
          {
                0,
                0,
                0,
                0,
                0
          }
     }
};
```

You could directly put that into a C program, although it probably isn't very useful yet. It becomes more useful if you actually specify how you want to initialize the type:

```
$data = {
  type   => 'QWORD',
  choice => {
    date  => { month => 12, day => 24 },
    qword => 4711,
  },
  stuff => 'yes?',
};

$init = $c->initializer( 'data', $data );
print "data x = $init;\n";
```

This would print the following:

```
data x = {
      QWORD,
      0,
      {
            {
                  0,
                  12,
                  24,
                  0,
                  0
            }
      }
};
```

As only the first member of a `union` can be initialized, `choice.qword` is ignored. You will not be warned about the fact that you probably tried to initialize a member other than the first. This is considered a feature, because it allows you to use unpack to generate the initializer data:

```
$data = $c->unpack( 'data', $binary );
$init = $c->initializer( 'data', $data );
```

Since unpack unpacks all union members, you would otherwise have to delete all but the first one previous to feeding it into `initializer`.

Also, `stuff` is ignored, because it actually isn't a member of `data`. You won't be warned about that either.

1.5.15   dependencies

`dependencies`

> After some code has been parsed using either the parse or parse_file methods, the dependencies method can be used to retrieve information about all files that the object depends on, i.e. all files that have been parsed.
>
> In scalar context, the method returns a hash reference. Each key is the name of a file. The values are again hash references, each of which holds the size, modification time (mtime), and change time (ctime) of the file at the moment it was parsed.

```
use Convert::Binary::C;
use Data::Dumper;

#-----------------------------------------------------------
# Create object, set include path, parse 'string.h' header
#-----------------------------------------------------------
my $c = Convert::Binary::C->new
        ->Include( '/usr/lib/gcc-lib/i686-pc-linux-gnu/3.3.4/include',
                   '/usr/include' )
        ->parse_file( 'string.h' );

#-----------------------------------------------------------
# Get dependencies of the object, extract dependency files
#-----------------------------------------------------------
my $depend = $c->dependencies;
my @files  = keys %$depend;
```

```
#----------------------------
# Dump dependencies and files
#----------------------------
print Data::Dumper->Dump( [$depend, \@files],
                          [qw( depend  *files )] );
```

The above code would print something like this:

```
$depend = {
  '/usr/include/features.h' => {
    'ctime' => 1092288224,
    'mtime' => 1092288215,
    'size' => 10792
  },
  '/usr/include/sys/cdefs.h' => {
    'ctime' => 1092288222,
    'mtime' => 1092288215,
    'size' => 8600
  },
  '/usr/include/gnu/stubs.h' => {
    'ctime' => 1092288221,
    'mtime' => 1092288215,
    'size' => 818
  },
  '/usr/include/string.h' => {
    'ctime' => 1092288224,
    'mtime' => 1092288215,
    'size' => 15011
  },
  '/usr/lib/gcc-lib/i686-pc-linux-gnu/3.3.4/include/stddef.h' => {
    'ctime' => 1093624078,
    'mtime' => 1093624077,
    'size' => 12695
  }
};
@files = (
  '/usr/include/features.h',
  '/usr/include/sys/cdefs.h',
  '/usr/include/gnu/stubs.h',
  '/usr/include/string.h',
  '/usr/lib/gcc-lib/i686-pc-linux-gnu/3.3.4/include/stddef.h'
);
```

In list context, the method returns the names of all files that have been parsed, i.e. the following lines are equivalent:

```
@files = keys %{$c->dependencies};
@files = $c->dependencies;
```

1.5.16 sourcify

```
sourcify
```

sourcify `CONFIG`

Returns a string that holds the C code necessary to represent all parsed C data structures.

```
use Convert::Binary::C;

$c = new Convert::Binary::C;
$c->parse( <<'END' );

#define NUMBER 42

typedef struct _mytype mytype;

struct _mytype {
  union {
    int          iCount;
    enum count *pCount;
  } counter;
#pragma pack( push, 1 )
  struct {
    char string[NUMBER];
    int  array[NUMBER/sizeof(int)];
  } storage;
#pragma pack( pop )
  mytype *next;
};

enum count { ZERO, ONE, TWO, THREE };

END

print $c->sourcify;
```

The above code would print something like this:

```
/* typedef predeclarations */

typedef struct _mytype mytype;

/* defined enums */

enum count
{
    ZERO,
    ONE,
    TWO,
    THREE
};

/* defined structs and unions */

struct _mytype
{
    union
    {
```

```
            int iCount;
            enum count *pCount;
      } counter;
#pragma pack( push, 1 )
      struct
      {
            char string[42];
            int array[10];
      } storage;
#pragma pack( pop )
      mytype *next;
};
```

The purpose of the `sourcify` method is to enable some kind of platform-independent caching. The C code generated by `sourcify` can be parsed by a standard C compiler, as well as of course the Convert::Binary::C parser. However, it might be significantly shorter than the code that has originally been parsed. When parsing a typical header file, it's easily possible that you need to open dozens of other files that are included from that file, and end up parsing several hundred kilobytes of C code. Since most of it is usually preprocessor directives, function prototypes and comments, the `sourcify` function strips this down to a few kilobytes. Saving the `sourcify` string and parsing it next time instead of the original code may be a lot faster.

The `sourcify` method takes a hash reference as an optional argument. It can be used to tweak the method's output. The following options can be configured.

`Context => 0 | 1`

> Turns preprocessor context information on or off. If this is turned on, `sourcify` will insert `#line` preprocessor directives in its output. So in the above example

```
print $c->sourcify( { Context => 1 } );
```

> would print:

```
/* typedef predeclarations */

typedef struct _mytype mytype;

/* defined enums */

#line 20 "[buffer]"
enum count
{
     ZERO,
     ONE,
     TWO,
     THREE
};

/* defined structs and unions */

#line 6 "[buffer]"
struct _mytype
{
#line 7 "[buffer]"
     union
     {
          int iCount;
          enum count *pCount;
     } counter;
```

```
#pragma pack( push, 1 )
#line 12 "[buffer]"
      struct
      {
            char string[42];
            int array[10];
      } storage;
#pragma pack( pop )
      mytype *next;
};
```

Note that `"[buffer]"` refers to the here-doc buffer when using parse.

The following methods can be used to retrieve information about the definitions that have been parsed. The examples given in the description for enum, compound and typedef all assume this piece of C code has been parsed:

```
typedef unsigned long U32;
typedef void *any;

enum __socket_type
{
  SOCK_STREAM    = 1,
  SOCK_DGRAM     = 2,
  SOCK_RAW       = 3,
  SOCK_RDM       = 4,
  SOCK_SEQPACKET = 5,
  SOCK_PACKET    = 10
};

struct STRUCT_SV {
  void *sv_any;
  U32   sv_refcnt;
  U32   sv_flags;
};

typedef union {
  int abc[2];
  struct xxx {
    int a;
    int b;
  }   ab[3][4];
  any ptr;
} test;
```

1.5.17   enum_names

`enum_names`

Returns a list of identifiers of all defined enumeration objects. Enumeration objects don't necessarily have an identifier, so something like

```
enum { A, B, C };
```

will obviously not appear in the list returned by the enum_names method. Also, enumerations that are not defined within the source code - like in

```
struct foo {
  enum weekday *pWeekday;
  unsigned long year;
};
```

where only a pointer to the weekday enumeration object is used - will not be returned, even though they have an identifier. So for the above two enumerations, enum_names will return an empty list:

```
@names = $c->enum_names;
```

The only way to retrieve a list of all enumeration identifiers is to use the enum method without additional arguments. You can get a list of all enumeration objects that have an identifier by using

```
@enums = map { $_->{identifier} || () } $c->enum;
```

but these may not have a definition. Thus, the two arrays would look like this:

```
@names = ();
@enums = ('weekday');
```

The def method returns a true value for all identifiers returned by enum_names.

1.5.18   enum

enum

enum LIST

Returns a list of references to hashes containing detailed information about all enumerations that have been parsed.

If a list of enumeration identifiers is passed to the method, the returned list will only contain hash references for those enumerations. The enumeration identifiers may optionally be prefixed by enum.

If an enumeration identifier cannot be found, a warning is issued and the returned list will contain an undefined value at that position.

In scalar context, the number of enumerations will be returned as long as the number of arguments to the method call is not 1. In the latter case, a hash reference holding information for the enumeration will be returned.

The list returned by the enum method looks similar to this:

```
@enum = (
  {
    'enumerators' => {
      'SOCK_STREAM' => 1,
      'SOCK_RAW' => 3,
      'SOCK_SEQPACKET' => 5,
      'SOCK_RDM' => 4,
      'SOCK_PACKET' => 10,
      'SOCK_DGRAM' => 2
    },
    'identifier' => '__socket_type',
```

```
      'context' => 'definitions.c(4)',
      'sign' => 0
    }
  );
```

identifier

holds the enumeration identifier. This key is not present if the enumeration has no identifier.

context

is the context in which the enumeration is defined. This is the filename followed by the line number in parentheses.

enumerators

is a reference to a hash table that holds all enumerators of the enumeration.

sign

is a boolean indicating if the enumeration is signed (i.e. has negative values).

One useful application may be to create a hash table that holds all enumerators of all defined enumerations:

```
%enum = map %{ $_->{enumerators} || {} }, $c->enum;
```

The `%enum` hash table would then be:

```
%enum = (
  'SOCK_STREAM' => 1,
  'SOCK_RAW' => 3,
  'SOCK_SEQPACKET' => 5,
  'SOCK_RDM' => 4,
  'SOCK_DGRAM' => 2,
  'SOCK_PACKET' => 10
);
```

1.5.19 compound_names

`compound_names`

Returns a list of identifiers of all structs and unions (compound data structures) that are defined in the parsed source code. Like enumerations, compounds don't need to have an identifier, nor do they need to be defined.

Again, the only way to retrieve information about all struct and union objects is to use the `compound` method and don't pass it any arguments. If you should need a list of all struct and union identifiers, you can use:

```
@compound = map { $_->{identifier} || () } $c->compound;
```

The `def` method returns a true value for all identifiers returned by `compound_names`.

If you need the names of only the structs or only the unions, use the `struct_names` and `union_names` methods respectively.

1.5.20 compound

compound

compound **LIST**

> Returns a list of references to hashes containing detailed information about all compounds (structs and unions) that have been parsed.
>
> If a list of struct/union identifiers is passed to the method, the returned list will only contain hash references for those compounds. The identifiers may optionally be prefixed by `struct` or `union`, which limits the search to the specified kind of compound.
>
> If an identifier cannot be found, a warning is issued and the returned list will contain an undefined value at that position.
>
> In scalar context, the number of compounds will be returned as long as the number of arguments to the method call is not 1. In the latter case, a hash reference holding information for the compound will be returned.
>
> The list returned by the compound method looks similar to this:

```
@compound = (
  {
    'identifier' => 'STRUCT_SV',
    'align' => 1,
    'context' => 'definitions.c(14)',
    'pack' => 0,
    'type' => 'struct',
    'declarations' => [
      {
        'declarators' => [
          {
            'declarator' => '*sv_any',
            'size' => 4,
            'offset' => 0
          }
        ],
        'type' => 'void'
      },
      {
        'declarators' => [
          {
            'declarator' => 'sv_refcnt',
            'size' => 4,
            'offset' => 4
          }
        ],
        'type' => 'U32'
      },
      {
        'declarators' => [
          {
            'declarator' => 'sv_flags',
            'size' => 4,
            'offset' => 8
          }
        ],
```

```
          'type' => 'U32'
        }
      ],
      'size' => 12
},
{
    'identifier' => 'xxx',
    'align' => 1,
    'context' => 'definitions.c(22)',
    'pack' => 0,
    'type' => 'struct',
    'declarations' => [
      {
        'declarators' => [
          {
            'declarator' => 'a',
            'size' => 4,
            'offset' => 0
          }
        ],
        'type' => 'int'
      },
      {
        'declarators' => [
          {
            'declarator' => 'b',
            'size' => 4,
            'offset' => 4
          }
        ],
        'type' => 'int'
      }
    ],
    'size' => 8
},
{
    'align' => 1,
    'context' => 'definitions.c(20)',
    'pack' => 0,
    'type' => 'union',
    'declarations' => [
      {
        'declarators' => [
          {
            'declarator' => 'abc[2]',
            'size' => 8,
            'offset' => 0
          }
        ],
        'type' => 'int'
      },
      {
        'declarators' => [
```

```
                   {
                     'declarator' => 'ab[3][4]',
                     'size' => 96,
                     'offset' => 0
                   }
                 ],
                 'type' => 'struct xxx'
               },
               {
                 'declarators' => [
                   {
                     'declarator' => 'ptr',
                     'size' => 4,
                     'offset' => 0
                   }
                 ],
                 'type' => 'any'
               }
             ],
             'size' => 96
           }
         );
```

identifier
> holds the struct or union identifier. This key is not present if the compound has no identifier.

context
> is the context in which the struct or union is defined. This is the filename followed by the line number in parentheses.

type
> is either 'struct' or 'union'.

size
> is the size of the struct or union.

align
> is the alignment of the struct or union.

pack
> is the struct member alignment if the compound is packed, or zero otherwise.

declarations
> is an array of hash references describing each struct declaration:
>
> **type**
> > is the type of the struct declaration. This may be a string or a reference to a hash describing the type.
>
> **declarators**
> > is an array of hashes describing each declarator:
> >
> > **declarator**
> > > is a string representation of the declarator.
> >
> > **offset**
> > > is the offset of the struct member represented by the current declarator relative to the beginning of the struct or union.
> >
> > **size**
> > > is the size occupied by the struct member represented by the current declarator.

It may be useful to have separate lists for structs and unions. One way to retrieve such lists would be to use

```
push @{$_->{type} eq 'union' ? \@unions : \@structs}, $_
    for $c->compound;
```

However, you should use the struct and union methods, which is a lot simpler:

```
@structs = $c->struct;
@unions  = $c->union;
```

1.5.21   struct_names

`struct_names`

> Returns a list of all defined struct identifiers. This is equivalent to calling `compound_names`, just that it only returns the names of the struct identifiers and doesn't return the names of the union identifiers.

1.5.22   struct

`struct`

`struct` **LIST**

> Like the `compound` method, but only allows for structs.

1.5.23   union_names

`union_names`

> Returns a list of all defined union identifiers. This is equivalent to calling `compound_names`, just that it only returns the names of the union identifiers and doesn't return the names of the struct identifiers.

1.5.24   union

`union`

`union` **LIST**

> Like the `compound` method, but only allows for unions.

1.5.25   typedef_names

`typedef_names`

> Returns a list of all defined typedef identifiers. Typedefs that do not specify a type that you could actually work with will not be returned.
>
> The `def` method returns a true value for all identifiers returned by `typedef_names`.

1.5.26 **typedef**

typedef

typedef **LIST**

> Returns a list of references to hashes containing detailed information about all typedefs that have been parsed.
>
> If a list of typedef identifiers is passed to the method, the returned list will only contain hash references for those typedefs.
>
> If an identifier cannot be found, a warning is issued and the returned list will contain an undefined value at that position.
>
> In scalar context, the number of typedefs will be returned as long as the number of arguments to the method call is not 1. In the latter case, a hash reference holding information for the typedef will be returned.
>
> The list returned by the typedef method looks similar to this:

```
  @typedef = (
    {
      'declarator' => 'U32',
      'type' => 'unsigned long'
    },
    {
      'declarator' => '*any',
      'type' => 'void'
    },
    {
      'declarator' => 'test',
      'type' => {
        'align' => 1,
        'context' => 'definitions.c(20)',
        'pack' => 0,
        'type' => 'union',
        'declarations' => [
          {
            'declarators' => [
              {
                'declarator' => 'abc[2]',
                'size' => 8,
                'offset' => 0
              }
            ],
            'type' => 'int'
          },
          {
            'declarators' => [
              {
                'declarator' => 'ab[3][4]',
                'size' => 96,
                'offset' => 0
              }
            ],
            'type' => 'struct xxx'
          },
```

```
              {
                'declarators' => [
                  {
                    'declarator' => 'ptr',
                    'size' => 4,
                    'offset' => 0
                  }
                ],
                'type' => 'any'
              }
            ],
            'size' => 96
          }
        }
    );
```

declarator
>   is the type declarator.

type
>   is the type specification. This may be a string or a reference to a hash describing the type.
>   See enum and compound for a description on how to interpret this hash.

1.5.27   add_hooks

add_hooks **TYPE => { HOOK => SUB, HOOK => [SUB, ARGS], ... }, ...**

>   This method allows you to register subroutines as hooks. Hooks are called whenever a certain TYPE
>   is packed or unpacked. Hooks are currently considered an experimental feature.

>   You can register hooks for all types except for basic types. This means you cannot register a hook
>   for int. TYPE is the type you want to register a hook for. HOOK can be one of the following:

```
  pack
  unpack
  pack_ptr
  unpack_ptr
```

>   'pack' and 'unpack' hooks are called when processing their TYPE, while 'pack_ptr' and
>   'unpack_ptr' hooks are called when processing pointers to their TYPE.

>   SUB is a reference to a subroutine that takes one input argument, processes it and returns one
>   output argument. Alternatively, you can pass custom arguments to the hook by passing an array
>   reference instead of SUB that holds the subroutine reference in the first element and the arguments
>   to be passed to the subroutine as the other elements. With the latter alternative, you can even
>   pass special arguments to the hook using arg.

>   You can register hooks for as many types as you like with one call to add_hooks:

```
    $c->add_hooks(ObjectType => { pack   => \&obj_pack,
                                  unpack => \&obj_unpack },
                  ProtocolId => { unpack => sub {
                                      $protos[$_[0]]
                                  },
                                  unpack_ptr => [sub {
                                    sprintf "$_[0]:{0x%X}", $_[1]
                                  }, $c->arg('TYPE', 'DATA')] });
```

To remove registered hooks, use either `delete_hooks` or `delete_all_hooks`.

You can also use `add_hooks` to update hooks for types already registered. To remove only a single subroutine, pass `undef` instead of a subroutine reference.

```
$c->add_hooks(ObjectType => { pack => undef });
```

If all subroutines are removed, the whole hook is removed.

`add_hooks` will throw an exception if an error occurs. On success, the method returns a reference to its object.

See the section on Using Hooks on page 17 for examples on how to use hooks.

1.5.28 delete_hooks

`delete_hooks` **LIST**

Deletes the hooks for all types passed in.

```
$c->delete_hooks(qw(ObjectType ProtocolId));
```

`delete_hooks` will throw an exception if an error occurs. On success, the method returns a reference to its object.

1.5.29 delete_all_hooks

`delete_all_hooks`

Removes all registered hooks and returns a reference to its object.

1.5.30 arg

`arg 'ARG', ...`

Creates placeholders for special arguments to be passed to hook subroutines. These arguments are currently:

SELF

A reference to the calling Convert::Binary::C object. This may be useful if you need to work with the object inside the hook.

TYPE

The name of the type that is currently being processed by the hook.

DATA

The data argument that is usually passed to the hook.

HOOK

The type of the hook as which the hook has been called, e.g. `pack` or `unpack_ptr`.

`arg` will return a placeholder for each argument it is being passed.

1.6 Functions

You can alternatively call the following functions as methods on Convert::Binary::C objects.

1.6.1 feature

`feature` **STRING**

> Checks if Convert::Binary::C was built with certain features. For example,

```
print "debugging version"
    if Convert::Binary::C::feature('debug');
```

will check if Convert::Binary::C was built with debugging support enabled. The `feature` function returns `1` if the feature is enabled, `0` if the feature is disabled, and `undef` if the feature is unknown. Currently the only features that can be checked are `ieeefp`, `debug` and `threads`.

You can enable or disable certain features at compile time of the module by using the

```
perl Makefile.PL enable-feature disable-feature
```

syntax.

1.6.2 native

`native`

`native` **STRING**

> Returns the value of a property of the native system that Convert::Binary::C was built on. For example,

```
$size = Convert::Binary::C::native('IntSize');
```

will fetch the size of an `int` on the native system. The following properties can be queried:

```
Alignment
ByteOrder
CompoundAlignment
DoubleSize
EnumSize
FloatSize
IntSize
LongDoubleSize
LongLongSize
LongSize
PointerSize
ShortSize
```

You can also call `native` without arguments, in which case it will return a reference to a hash with all properties, like:

```
$native = {
  'ByteOrder' => 'LittleEndian',
  'LongSize' => 4,
  'IntSize' => 4,
  'ShortSize' => 2,
  'FloatSize' => 4,
  'Alignment' => 4,
  'LongLongSize' => 8,
```

```
        'LongDoubleSize' => 12,
        'DoubleSize' => 8,
        'EnumSize' => 4,
        'CompoundAlignment' => 1,
        'PointerSize' => 4
      };
```

The contents of that hash are suitable for passing them to the `configure` method.

1.7 Debugging

Like perl itself, Convert::Binary::C can be compiled with debugging support that can then be selectively enabled at runtime. You can specify whether you like to build Convert::Binary::C with debugging support or not by explicitly giving an argument to *Makefile.PL*. Use

```
perl Makefile.PL enable-debug
```

to enable debugging, or

```
perl Makefile.PL disable-debug
```

to disable debugging. The default will depend on how your perl binary was built. If it was built with `-DDEBUGGING`, Convert::Binary::C will be built with debugging support, too.

Once you have built Convert::Binary::C with debugging support, you can use the following syntax to enable debug output. Instead of

```
use Convert::Binary::C;
```

you simply say

```
use Convert::Binary::C debug => 'all';
```

which will enable all debug output. However, I don't recommend to enable all debug output, because that can be a fairly large amount.

1.7.1 Debugging options

Instead of saying `all`, you can pass a string that consists of one or more of the following characters:

```
m   enable memory allocation tracing
M   enable memory allocation & assertion tracing

h   enable hash table debugging
H   enable hash table dumps

d   enable debug output from the XS module
c   enable debug output from the ctlib
t   enable debug output about type objects
```

```
l    enable debug output from the C lexer
p    enable debug output from the C parser
r    enable debug output from the #pragma parser


y    enable debug output from yacc (bison)
```

So the following might give you a brief overview of what's going on inside Convert::Binary::C:

```
use Convert::Binary::C debug => 'dct';
```

When you want to debug memory allocation using

```
use Convert::Binary::C debug => 'm';
```

you can use the Perl script *check˙alloc.pl* that resides in the *ctlib/util/tool* directory to extract statistics about memory usage and information about memory leaks from the resulting debug output.

1.7.2   Redirecting debug output

By default, all debug output is written to `stderr`. You can, however, redirect the debug output to a file with the `debugfile` option:

```
use Convert::Binary::C debug     => 'dcthHm',
                       debugfile => './debug.out';
```

If the file cannot be opened, you'll receive a warning and the output will go the `stderr` way again.

Alternatively, you can use the environment variables CBC_DEBUG_OPT and CBC_DEBUG_FILE to turn on debug output.

If Convert::Binary::C is built without debugging support, passing the `debug` or `debugfile` options will cause a warning to be issued. The corresponding environment variables will simply be ignored.

1.8   Environment

1.8.1   CBC_ORDER_MEMBERS

Setting this variable to a non-zero value will globally turn on hash key ordering for compound members. Have a look at the `OrderMembers` option for details.

Setting the variable to the name of a perl module will additionally use this module instead of the predefined modules for member ordering to tie the hashes to.

1.8.2   CBC_DEBUG_OPT

If Convert::Binary::C is built with debugging support, you can use this variable to specify the debugging options.

1.8.3 CBC_DEBUG_FILE

If Convert::Binary::C is built with debugging support, you can use this variable to redirect the debug output to a file.

1.8.4 CBC_DISABLE_PARSER

This variable is intended purely for development. Setting it to a non-zero value disables the Convert::Binary::C parser, which means that no information is collected from the file or code that is parsed. However, the preprocessor will run, which is useful for benchmarking the preprocessor.

1.9 Flexible Array Members And Incomplete Types

Flexible array members are a feature introduced with ISO-C99. It's a common problem that you have a variable length data field at the end of a structure, for example an array of characters at the end of a message struct. ISO-C99 allows you to write this as:

```
struct message {
  long header;
  char data[];
};
```

The advantage is that you clearly indicate that the size of the appended data is variable, and that the `data` member doesn't contribute to the size of the `message` structure.

When packing or unpacking data, Convert::Binary::C deals with flexible array members as if their length was adjustable. For example, unpack will adapt the length of the array depending on the input string:

```
$msg1 = $c->unpack('message', 'abcdefg');
$msg2 = $c->unpack('message', 'abcdefghijkl');
```

The following data is unpacked:

```
$msg1 = {
  'data' => [
    101,
    102,
    103
  ],
  'header' => 1633837924
};
$msg2 = {
  'data' => [
    101,
    102,
    103,
    104,
    105,
    106,
    107,
    108
```

```
    ],
    'header' => 1633837924
  };
```

Similarly, pack will adjust the length of the output string according to the data you feed in:

```
  use Data::Hexdumper;

  $msg = {
    header => 4711,
    data   => [0x10, 0x20, 0x30, 0x40, 0x77..0x88],
  };

  $data = $c->pack('message', $msg);

  print hexdump(data => $data);
```

This would print:

```
    0x0000 : 00 00 12 67 10 20 30 40 77 78 79 7A 7B 7C 7D 7E : ...g..0@wxyz{|}~
    0x0010 : 7F 80 81 82 83 84 85 86 87 88                   : ..........
```

Incomplete types such as

```
  typedef unsigned long array[];
```

are handled in exactly the same way. Thus, you can easily

```
  $array = $c->unpack('array', '?'x20);
```

which will unpack the following array:

```
  $array = [
    1061109567,
    1061109567,
    1061109567,
    1061109567,
    1061109567
  ];
```

1.10   Floating Point Values

When using Convert::Binary::C to handle floating point values, you have to be aware of some limitations.

You're usually safe if all your platforms are using the IEEE floating point format. During the Convert::Binary::C build process, the `ieeefp` feature will automatically be enabled if the host is using IEEE floating point. You can check for this feature at runtime using the `feature` function:

```
  if (Convert::Binary::C::feature('ieeefp')) {
    # do something
  }
```

When IEEE floating point support is enabled, the module can also handle floating point values of a different byteorder.

If your host platform is not using IEEE floating point, the `ieeefp` feature will be disabled. Convert::Binary::C then will be more restrictive, refusing to handle any non-native floating point values.

However, Convert::Binary::C cannot detect the floating point format used by your target platform. It can only try to prevent problems in obvious cases. If you know your target platform has a completely different floating point format, don't use floating point conversion at all.

Whenever Convert::Binary::C detects that it cannot properly do floating point value conversion, it will issue a warning and will not attempt to convert the floating point value.

1.11   Bitfields

Bitfields are currently not supported by Convert::Binary::C, because I generally don't use them. I plan to support them in a later release, when I will have found an easy way of integrating them into the module.

Whenever a method has to deal with bitfields, it will issue a warning message that bitfields are unsupported. Thus, you may use bitfields in your C source code, but you won't be annoyed with warning messages unless you really use a type that actually contains bitfields in a method call like `sizeof` or `pack`.

While bitfields are not appropriately handled by the conversion routines yet, they are already parsed correctly. This means that you can reliably use the declarator fields as returned by the `struct` or `typedef` methods. Given the following source

```
struct bitfield {
  int seven:7;
  int :1;
  int four:4, :0;
  int integer;
};
```

a call to `struct` will return

```
@struct = (
  {
    'identifier' => 'bitfield',
    'align' => 1,
    'context' => 'bitfields.c(1)',
    'pack' => 0,
    'type' => 'struct',
    'declarations' => [
      {
        'declarators' => [
          {
            'declarator' => 'seven:7'
          }
        ],
        'type' => 'int'
      },
      {
        'declarators' => [
```

```
          {
            'declarator' => ':1'
          }
        ],
        'type' => 'int'
      },
      {
        'declarators' => [
          {
            'declarator' => 'four:4'
          },
          {
            'declarator' => ':0'
          }
        ],
        'type' => 'int'
      },
      {
        'declarators' => [
          {
            'declarator' => 'integer',
            'size' => 4,
            'offset' => 0
          }
        ],
        'type' => 'int'
      }
    ],
    'size' => 4
  }
);
```

No size/offset keys will be returned for bitfield entries. Also, the size of a structure containing bitfields is not valid, as bitfields internally do not increase the size of a structure yet.

1.12   Multithreading

Convert::Binary::C was designed to be thread-safe.

1.13   Inheritance

If you wish to derive a new class from Convert::Binary::C, this is relatively easy. Despite their XS implementation, Convert::Binary::C objects are actually blessed hash references.

The XS data is stored in a read-only hash value for the key that is the empty string. So it is safe to use any non-empty hash key when deriving your own class. In addition, Convert::Binary::C does quite a lot of checks to detect corruption in the object hash.

If you store private data in the hash, you should override the `clone` method and provide the necessary code to clone your private data. You'll have to call `SUPER::clone`, but this will only clone the Convert::Binary::C part of the object.

For an example of a derived class, you can have a look at Convert::Binary::C::Cached.

1.14 Portability

Convert::Binary::C should build and run on most of the platforms that Perl runs on:

- Various Linux systems

- Various BSD systems

- HP-UX

- Compaq/HP Tru64 Unix

- Mac-OS X

- Cygwin

- Windows 98/NT/2000/XP

Also, many architectures are supported:

- Various Intel Pentium and Itanium systems

- Various Alpha systems

- HP PA-RISC

- Power-PC

- StrongARM

The module should build with any perl binary from 5.5.3 up to the latest development version.

1.15 Comparison With Similar Modules

Most of the time when you're really looking for Convert::Binary::C you'll actually end up finding one of the following modules. Some of them have different goals, so it's probably worth pointing out the differences.

1.15.1 C::Include

Like Convert::Binary::C, this module aims at doing conversion from and to binary data based on C types. However, its configurability is very limited compared to Convert::Binary::C. Also, it does not parse all C code correctly. It's slower than Convert::Binary::C, doesn't have a preprocessor. On the plus side, it's written in pure Perl.

1.15.2 C::DynaLib::Struct

This module doesn't allow you to reuse your C source code. One main goal of Convert::Binary::C was to avoid code duplication or, even worse, having to maintain different representations of your data structures. Like C::Include, C::DynaLib::Struct is rather limited in its configurability.

1.15.3   Win32::API::Struct

This module has a special purpose. It aims at building structs for interfacing Perl code with Windows API code.

1.16   Credits

- My love Jennifer for always being there, for filling my life with joy and last but not least for proofreading the documentation.

- Alain Barbet <alian@cpan.org> for testing and debugging support.

- Mitchell N. Charity for giving me pointers into various interesting directions.

- Alexis Denis for making me improve (externally) and simplify (internally) floating point support. He can also be blamed (indirectly) for the `initializer` method, as I need it in my effort to support bitfields some day.

- Michael J. Hohmann <mjh@scientist.de> for endless discussions on our way to and back home from work, and for making me think about supporting `pack` and `unpack` for compound members.

- Thorsten Jens <thojens@gmx.de> for testing the package on various platforms.

- Mark Overmeer <mark@overmeer.net> for suggesting the module name and giving invaluable feedback.

- Thomas Pornin <pornin@bolet.org> for his excellent `ucpp` preprocessor library.

- Marc Rosenthal for his suggestions and support.

- James Roskind, as his C parser was a great starting point to fix all the problems I had with my original parser based only on the ANSI ruleset.

- Gisbert W. Selke for spotting some interesting bugs and providing extensive reports.

- Steffen Zimmermann for a prolific discussion on the cloning algorithm.

1.17   Bugs

I'm sure there are still lots of bugs in the code for this module. If you find any bugs, Convert::Binary::C doesn't seem to build on your system or any of its tests fail, please use the CPAN Request Tracker at http://rt.cpan.org/ to create a ticket for the module. Alternatively, just send a mail to <mhx@cpan.org>.

1.18   Todo

If you're interested in what I currently plan to improve (or fix), have a look at the *TODO* file.

1.19 Postcards

If you're using my module and like it, you can show your appreciation by sending me a postcard from where you live. I won't urge you to do it, it's completely up to you. To me, this is just a very nice way of receiving feedback about my work. Please send your postcard to:

```
Marcus Holland-Moritz
Kuppinger Weg 28
71116 Gaertringen
GERMANY
```

If you feel that sending a postcard is too much effort, you maybe want to rate the module at http://cpanratings.perl.org/.

1.20 Copyright

Copyright (c) 2002-2004 Marcus Holland-Moritz. All rights reserved. This program is free software; you can redistribute it and/or modify it under the same terms as Perl itself.

The `ucpp` library is (c) 1998-2002 Thomas Pornin. For license and redistribution details refer to *ctlib/ucpp/README*.

Portions copyright (c) 1989, 1990 James A. Roskind.

The include files located in *t/include/include*, which are used in some of the test scripts are (c) 1991-1999, 2000, 2001 Free Software Foundation, Inc. They are neither required to create the binary nor linked to the source code of this module in any other way.

1.21 See Also

See the *ccconfig* manpage, the *perl* manpage, the *perldata* manpage, the *perlop* manpage, the *perlvar* manpage, the *Data::Dumper* manpage and the *Scalar::Util* manpage.

Convert::Binary::C::Cached

Caching for Convert::Binary::C

2.1 Synopsis

```
use Convert::Binary::C::Cached;
use Data::Dumper;

#------------------------
# Create a cached object
#------------------------
$c = new Convert::Binary::C::Cached
         Cache   => '/tmp/cache.c',
         Include => ['include']
       ;

#--------------------------------------------------
# Parse 'stdio.h' and dump the definition of FILE
#--------------------------------------------------
$c->parse_file( 'stdio.h' );

print Dumper( $c->typedef( 'FILE' ) );
```

2.2 Description

Convert::Binary::C::Cached simply adds caching capability to Convert::Binary::C. You can use it in just the same way that you would use Convert::Binary::C. The interface is exactly the same.

To use the caching capability, you must pass the `Cache` option to the constructor. If you don't pass it, you will receive an ordinary Convert::Binary::C object. The argument to the `Cache` option is the file that is used for caching this object.

The caching algorithm automatically detects when the cache file cannot be used and the original code has to be parsed. In that case, the cache file is updated. An update of the cache file can be triggered by one or more of the following factors:

- The cache file doesn't exist, which is obvious.

- The cache file is corrupt, i.e. cannot be parsed.

- The object's configuration has changed.

- The embedded code for a `parse` method call has changed.

- At least one of the files that the object depends on does not exist or has a different size or a different modification or change timestamp.

2.3 Limitations

You cannot call `parse` or `parse_file` more that once when using a Convert::Binary::C::Cached object. This isn't a big problem, as you usually don't call them multiple times.

If a dependency file changes, but the change affects neither the size nor the timestamps of that file, the caching algorithm cannot detect that an update is required.

2.4 Copyright

Copyright (c) 2002-2004 Marcus Holland-Moritz. All rights reserved. This program is free software; you can redistribute it and/or modify it under the same terms as Perl itself.

2.5 See Also

See the *Convert::Binary::C* manpage.

ccconfig

Get Convert::Binary::C configuration for a compiler

3.1 Synopsis

ccconfig *options* [– compiler-options]

options:

```
 -c
 --cc       compiler     compiler executable to test
                         default: auto-determined


 -p
 --ppout    flag         compiler option for sending
                         preprocessor output to stdout
                         default: -E


 -t
 --temp     file         name of the temporary test file
                         default: _t_e_s_t.c


 --nodelete              don't delete temporary files


 --norun                 don't try to run executables


 --quiet                 don't display anything
 --nostatus              don't display status indicator


 --version               print version number


 --debug                 debug mode
```

3.2 Description

ccconfig will try to determine a usable configuration for Convert::Binary::C from testing a compiler executable. It is not necessary that the binaries generated by the compiler can be executed, so **ccconfig** can be used for cross-compilers.

The tool is still experimental, and you should neither rely on its results without checking, nor expect it to work in your environment.

3.3 Options

3.3.1 `--cc` **compiler**

This option allows you to explicitly specify a compiler executable. This is especially useful if you don't want to use your system compiler.

3.3.2 `--ppout` **flag**

This option tells `ccconfig` which flag must be used to make the compiler write the preprocessor output to standard output. The default is `-E`, which is correct for many compilers.

3.3.3 `--temp` **file**

Allows you to change the name of the temporary test file.

3.3.4 `--nodelete`

Don't attempt to delete temporary files that have been created by the compiler. Normally, `ccconfig` will look for all files with the same basename as the temporary test file and delete them.

3.3.5 `--norun`

You can specify this option if the executables generated by your compiler cannot be run on your machine, i.e. if you have a cross-compiler. However, `ccconfig` will automatically find out that it cannot run the executables.

When this option is set, a different set of algorithms is used to determine a couple of configuration settings. These algorithms are all based upon placing a special signature in the object file. They are less reliable that the standard algorithms, so you shouldn't use them unless you have to.

3.3.6 `--quiet`

Don't display anything except for the final configuration.

3.3.7 `--nostatus`

Hide the status indicator. Recommended if you want to redirect the script output to a file:

```
ccconfig --nostatus >config.pl 2>ccconfig.log
```

3.3.8 `--version`

Writes the program name, version and path to standard output.

3.3.9 `--debug`

Generate tons of debug output. Don't use unless you know what you're doing.

3.4 Examples

Normally, a simple

```
ccconfig
```

without arguments is enough if you want the configuration for your system compiler. While `ccconfig` is running, it will write lots of status information to `stderr`. When it's done, it will write a Perl hash table to `stdout` which can be directly used as a configuration for Convert::Binary::C.

If you want the configuration for a different compiler, or `ccconfig` cannot determine your system compiler automatically, use

```
ccconfig -c gcc32
```

if your compiler's name is `gcc32`.

If you want to pass additional options to the compiler, you can do so after a double-dash on the command line:

```
ccconfig -- -g -DDEBUGGING
```

or

```
ccconfig -c gcc32 -- -ansi -fshort-enums
```

3.5 Copyright

Copyright (c) 2002-2004 Marcus Holland-Moritz. All rights reserved. This program is free software; you can redistribute it and/or modify it under the same terms as Perl itself.

3.6 See Also

See the *Convert::Binary::C* manpage.